

Cloud Search Service

User Guide

Issue	01
Date	2025-09-05



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

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

Trademarks and Permissions



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

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

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

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

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

Contents

1 CSS Service Permission Management.....	1
1.1 Creating IAM Users and Granting Them Permissions to Use CSS.....	1
1.2 Example: Granting IAM Users the Permission to Create CSS Clusters.....	4
2 Using Elasticsearch for Data Search.....	11
2.1 Procedure for Using Elasticsearch.....	11
2.2 Elasticsearch Cluster Planning Suggestions.....	15
2.2.1 Planning Cluster AZs and HA.....	15
2.2.2 Planning the Cluster Version and Security Mode.....	18
2.2.3 Planning Node Types and Specifications.....	22
2.3 Creating an Elasticsearch Cluster (New Version).....	35
2.4 Creating an Elasticsearch Cluster (Old Version).....	54
2.5 Accessing Elasticsearch Clusters.....	66
2.5.1 Elasticsearch Cluster Access Methods.....	67
2.5.2 Logging In to an Elasticsearch Cluster Using Kibana.....	70
2.5.3 Logging In to an Elasticsearch Cluster Through Cerebro.....	76
2.5.4 Accessing an Elasticsearch Cluster Using cURL Commands.....	77
2.5.5 Accessing an Elasticsearch Cluster Using Java.....	78
2.5.5.1 Accessing an Elasticsearch Cluster Through the High Level REST Client.....	78
2.5.5.2 Accessing an Elasticsearch Cluster Through the Low Level REST Client.....	87
2.5.5.3 Accessing an Elasticsearch Cluster Through the Transport Client.....	103
2.5.5.4 Accessing an Elasticsearch Cluster Using Spring Boot.....	104
2.5.6 Accessing an Elasticsearch Cluster Using Python.....	110
2.5.7 Accessing an Elasticsearch Cluster Through the MRS Hive Client.....	113
2.5.8 Accessing an Elasticsearch Cluster Using Go.....	119
2.5.9 Accessing an Elasticsearch Cluster Using LDAP.....	122
2.6 Importing Data to an Elasticsearch Cluster.....	126
2.6.1 Different Ways to Ingest Data into an Elasticsearch Cluster.....	127
2.6.2 Using In-house Built Logstash to Import Data to Elasticsearch.....	129
2.6.3 Using Open Source Elasticsearch APIs to Import Data to Elasticsearch.....	136
2.6.4 Using CDM to Import Data to Elasticsearch.....	140
2.6.5 Enhancing the Data Ingestion Performance of Elasticsearch Clusters.....	143
2.7 Searching Data in an Elasticsearch Cluster.....	146
2.7.1 Using DSL to Search for Data in Elasticsearch.....	146

2.7.2 Using SQL to Search for Data in Elasticsearch.....	148
2.8 Enhancing Search Capabilities for Elasticsearch Clusters.....	152
2.8.1 Search Enhancement Features for Elasticsearch Clusters.....	152
2.8.2 Configuring Vector Search for Elasticsearch Clusters.....	156
2.8.2.1 About CSS Vector Search.....	156
2.8.2.2 Creating an Elasticsearch Vector Cluster.....	157
2.8.2.3 Creating a Vector Index.....	160
2.8.2.4 Importing Vector Data.....	170
2.8.2.5 Performing Vector Search.....	172
2.8.2.6 Optimizing Vector Cluster Performance.....	181
2.8.2.7 Managing the Vector Index Cache.....	185
2.8.2.8 Using Nested Fields for Vector Search.....	186
2.8.2.9 Sample Python Code for Vector Search.....	189
2.8.2.10 Sample Java Code for Vector Search.....	191
2.8.3 Configuring Decoupled Storage and Compute for an Elasticsearch Cluster.....	195
2.8.4 Configuring Flow Control 2.0 for an Elasticsearch Cluster.....	208
2.8.5 Configuring Flow Control 1.0 for an Elasticsearch Cluster.....	220
2.8.6 Configuring Large Query Isolation for an Elasticsearch Cluster.....	236
2.8.7 Configuring Enhanced Aggregation for an Elasticsearch Cluster.....	242
2.8.8 Configuring Read/Write Splitting Between Two Elasticsearch Clusters.....	252
2.8.9 Switching Between Hot and Cold Storage for an Elasticsearch Cluster.....	261
2.8.10 Configuring an Index Recycle Bin for an Elasticsearch Cluster.....	264
2.9 Configuring Networking for an Elasticsearch Cluster.....	267
2.9.1 Configuring Routes for an Elasticsearch Cluster.....	267
2.9.2 Configuring Public Network Access for an Elasticsearch Cluster.....	269
2.9.3 Configuring VPC Endpoints for an Elasticsearch Cluster.....	273
2.9.4 Configuring a Dedicated Load Balancer for an Elasticsearch Cluster.....	277
2.10 Backing up and Restoring the Data of an Elasticsearch Cluster.....	289
2.10.1 Creating Snapshots to Back Up the Data of an Elasticsearch Cluster.....	289
2.10.2 Restoring the Data of an Elasticsearch Cluster Using a Snapshot.....	299
2.10.3 Deleting an Elasticsearch Cluster Snapshot.....	303
2.11 Scaling an Elasticsearch Cluster.....	304
2.11.1 Scaling Out/Up an Elasticsearch Cluster.....	304
2.11.2 Scaling In an Elasticsearch Cluster.....	310
2.12 Upgrading the Version of an Elasticsearch Cluster.....	316
2.13 Managing Elasticsearch Clusters.....	327
2.13.1 Viewing Elasticsearch Cluster Information.....	327
2.13.2 Creating Users for an Elasticsearch Cluster and Granting Cluster Access.....	335
2.13.3 Setting Tags for an Elasticsearch Cluster.....	340
2.13.4 Configuring Default Parameters in the .yaml Configuration File of an Elasticsearch Cluster.....	341
2.13.5 Associating an Elasticsearch Cluster with an Enterprise Project.....	343
2.13.6 Changing the Node Specifications and Storage Type of an Elasticsearch Cluster.....	345

2.13.7 Replacing a Specified Node for an Elasticsearch Cluster.....	352
2.13.8 Changing the Security Mode of an Elasticsearch Cluster.....	356
2.13.9 Switching AZs for an Elasticsearch Cluster.....	363
2.13.10 Configuring Custom Word Dictionaries for an Elasticsearch Cluster.....	368
2.13.10.1 Introduction to Word Dictionaries for Elasticsearch Clusters.....	368
2.13.10.2 Configuring and Using Custom Word Dictionaries for an Elasticsearch Cluster.....	371
2.13.11 Switching Between Simplified and Traditional Chinese for Data Search in an Elasticsearch Cluster.....	380
2.13.12 Restarting an Elasticsearch Cluster.....	382
2.13.13 Deleting an Elasticsearch Cluster.....	384
2.14 Managing Index Policies for an Elasticsearch Cluster.....	385
2.14.1 Creating and Managing Index Policies for an Elasticsearch Cluster.....	385
2.14.2 Automating Index Rollover in an Elasticsearch Cluster Through Index Lifecycle Management.....	389
2.14.3 Decoupling Index Storage and Compute in an Elasticsearch Cluster Through Index Lifecycle Management.....	393
2.15 Elasticsearch Cluster Monitoring and Log Management.....	396
2.15.1 Elasticsearch Cluster Monitoring Metrics.....	396
2.15.1.1 Monitoring Metrics for Elasticsearch Clusters in Cloud Eye.....	396
2.15.1.2 Elasticsearch Cluster Monitoring Metrics Supported by the Kernel.....	454
2.15.2 Configuring Elasticsearch Cluster Monitoring.....	454
2.15.2.1 Using Cloud Eye to Monitor Elasticsearch Clusters.....	455
2.15.2.2 Configuring Kernel Monitoring for an Elasticsearch Cluster.....	457
2.15.2.3 Configuring Index Monitoring for an Elasticsearch Cluster.....	460
2.15.3 Setting Alarm Alerting via SMN for an Elasticsearch Cluster.....	471
2.15.4 Synchronizing Elasticsearch Alerts to Prometheus.....	476
2.15.5 Intelligent Risk Detection for Elasticsearch Clusters.....	481
2.15.6 Managing Elasticsearch Cluster Logs.....	484
2.16 Viewing Elasticsearch Cluster Audit Logs.....	495
3 Using OpenSearch for Data Search.....	500
3.1 Procedure for Using OpenSearch.....	500
3.2 OpenSearch Cluster Planning Suggestions.....	503
3.2.1 Planning Cluster AZs and HA.....	503
3.2.2 Planning the Cluster Version and Security Mode.....	507
3.2.3 Planning Node Types and Specifications.....	510
3.3 Creating an OpenSearch Cluster (New Version).....	524
3.4 Creating an OpenSearch Cluster (Old Version).....	544
3.5 Accessing an OpenSearch Cluster.....	556
3.5.1 OpenSearch Cluster Access Methods.....	557
3.5.2 Logging In to an OpenSearch Cluster Through OpenSearch Dashboards.....	560
3.5.3 Logging In to an OpenSearch Cluster Through Cerebro.....	565
3.5.4 Accessing an OpenSearch Cluster Using cURL Commands.....	566
3.5.5 Accessing an OpenSearch Cluster Using Java.....	567
3.5.5.1 Accessing an OpenSearch Cluster Using the High Level REST Client.....	567

3.5.5.2 Accessing an OpenSearch Cluster Using the Low Level REST Client.....	584
3.5.6 Accessing an OpenSearch Cluster with the Python Client.....	619
3.5.7 Accessing an OpenSearch Cluster Using LDAP.....	625
3.6 Importing Data to an OpenSearch Cluster.....	629
3.6.1 Different Ways to Ingest Data into an OpenSearch Cluster.....	630
3.6.2 Using In-house Built Logstash to Ingest Data into an OpenSearch Cluster.....	632
3.6.3 Using Open Source OpenSearch APIs to Import Data to an OpenSearch Cluster.....	639
3.6.4 Using CDM to Import Data to an OpenSearch Cluster.....	642
3.6.5 Enhancing the Data Ingestion Performance of OpenSearch Clusters.....	644
3.7 Searching Data in an OpenSearch Cluster.....	648
3.7.1 Using DSL to Search for Data in OpenSearch.....	648
3.7.2 Using SQL to Search for Data in OpenSearch.....	650
3.8 Enhancing Search Capabilities for OpenSearch Clusters.....	654
3.8.1 Search Enhancement Features for OpenSearch Clusters.....	654
3.8.2 Configuring Vector Search for OpenSearch Clusters.....	656
3.8.2.1 About CSS Vector Search.....	656
3.8.2.2 Creating an OpenSearch Vector Cluster.....	658
3.8.2.3 Creating a Vector Index.....	661
3.8.2.4 Importing Vector Data.....	670
3.8.2.5 Performing Vector Search.....	672
3.8.2.6 Optimizing Vector Cluster Performance.....	681
3.8.2.7 Managing the Vector Index Cache.....	685
3.8.2.8 Using Nested Fields for Vector Search.....	686
3.8.2.9 Sample Python Code for Vector Search.....	690
3.8.2.10 Sample Java Code for Vector Search.....	692
3.8.3 Configuring Decoupled Storage and Compute for an OpenSearch Cluster.....	695
3.8.4 Configuring Flow Control for an OpenSearch Cluster.....	709
3.8.5 Configuring Large Query Isolation for an OpenSearch Cluster.....	722
3.8.6 Configuring Enhanced Aggregation for an OpenSearch Cluster.....	728
3.8.7 Configuring Read/Write Splitting Between Two OpenSearch Clusters.....	736
3.8.8 Switching Between Hot and Cold Storage for an OpenSearch Cluster.....	744
3.8.9 Configuring an Index Recycle Bin for an OpenSearch Cluster.....	748
3.9 Configuring Networking for an OpenSearch Cluster.....	751
3.9.1 Configuring Routes for an OpenSearch Cluster.....	751
3.9.2 Configuring Public Network Access for an OpenSearch Cluster.....	753
3.9.3 Configuring VPC Endpoint Service for an OpenSearch Cluster.....	757
3.9.4 Configuring a Dedicated Load Balancer for an OpenSearch Cluster.....	761
3.10 Backing up and Restoring the Data of an OpenSearch Cluster.....	773
3.10.1 Creating Snapshots to Back Up the Data of an OpenSearch Cluster.....	773
3.10.2 Restoring a Snapshot of an OpenSearch Cluster.....	783
3.10.3 Deleting an OpenSearch Cluster Snapshot.....	787
3.11 Scaling an OpenSearch Cluster.....	788

3.11.1 Scaling Out/Up an OpenSearch Cluster.....	788
3.11.2 Scaling In an OpenSearch Cluster.....	793
3.12 Upgrading the Version of an OpenSearch Cluster.....	800
3.13 Managing OpenSearch Clusters.....	811
3.13.1 Viewing OpenSearch Cluster Information.....	811
3.13.2 Creating Users for an OpenSearch Cluster and Granting Cluster Access.....	819
3.13.3 Setting Tags for an OpenSearch Cluster.....	823
3.13.4 Configuring Default Parameters in the .yaml Configuration File of an OpenSearch Cluster.....	824
3.13.5 Associating an OpenSearch Cluster with an Enterprise Project.....	827
3.13.6 Changing the Node Specifications of an OpenSearch Cluster.....	828
3.13.7 Replacing Specified Nodes for an OpenSearch Cluster.....	835
3.13.8 Changing the Security Mode of an OpenSearch Cluster.....	839
3.13.9 Switching AZs for an OpenSearch Cluster.....	846
3.13.10 Configuring Custom Word Dictionaries for an OpenSearch Cluster.....	851
3.13.10.1 Introduction to Word Dictionaries for OpenSearch Clusters.....	852
3.13.10.2 Configuring and Using Custom Word Dictionaries for an OpenSearch Cluster.....	854
3.13.11 Switching Between Simplified and Traditional Chinese for Data Search in an OpenSearch Cluster.....	861
3.13.12 Restarting an OpenSearch Cluster.....	863
3.13.13 Deleting an OpenSearch Cluster.....	865
3.14 Managing Index Policies for OpenSearch Clusters.....	866
3.14.1 Creating and Managing OpenSearch Index Policies.....	866
3.14.2 Automating Index Rollover in an OpenSearch Cluster Through Index Lifecycle Management.....	868
3.14.3 Decoupling Index Storage and Compute in an OpenSearch Cluster Through Index Lifecycle Management.....	873
3.15 OpenSearch Cluster Monitoring and Log Management.....	875
3.15.1 Monitoring Metrics for OpenSearch Clusters in Cloud Eye.....	875
3.15.2 Using Cloud Eye to Monitor OpenSearch Clusters.....	917
3.15.3 Configuring Index Monitoring for an OpenSearch Cluster.....	920
3.15.4 Setting Alarm Alerting via SMN for an OpenSearch Cluster.....	930
3.15.5 Synchronizing OpenSearch Alerts to Prometheus.....	941
3.15.6 Intelligent Risk Detection for OpenSearch Clusters.....	947
3.15.7 Managing OpenSearch Cluster Logs.....	950
3.16 Viewing OpenSearch Cluster Audit Logs.....	960
4 Using Logstash for Data Migration.....	966
4.1 Procedure for Using Logstash.....	966
4.2 Logstash Cluster Planning Suggestions.....	967
4.2.1 Planning Cluster AZs and HA.....	968
4.2.2 Planning Node Specifications and Capacity.....	969
4.3 Creating a Logstash Cluster (New Version).....	974
4.4 Creating a Logstash Cluster (Old Version).....	980
4.5 Configuring Networking for a Logstash Cluster.....	984
4.5.1 Configuring Routes for a Logstash Cluster.....	984

4.5.2 Configuring Public Network Access for a Logstash Cluster..... 988

4.6 Configuring a Logstash Migration Task..... 991

4.6.1 Creating a Logstash Configuration File..... 991

4.6.2 Starting a Logstash Migration Task..... 994

4.6.3 Stopping Logstash Migration Tasks..... 996

4.6.4 Logstash Configuration File Templates..... 997

4.7 Scaling a Logstash Cluster..... 1007

4.7.1 Scaling Out a Logstash Cluster..... 1007

4.7.2 Scaling In a Logstash Cluster..... 1009

4.8 Managing Logstash Clusters..... 1012

4.8.1 Viewing Logstash Cluster Information..... 1012

4.8.2 Setting Tags for a Logstash Cluster..... 1015

4.8.3 Associating a Logstash Cluster with an Enterprise Project..... 1016

4.8.4 Changing the Node Specifications of a Logstash Cluster..... 1018

4.8.5 Managing CA Certificates for a Logstash Cluster..... 1021

4.8.6 Forcibly Restarting a Logstash Cluster..... 1022

4.8.7 Deleting a Logstash Cluster..... 1022

4.9 Logstash Cluster Monitoring and Log Management..... 1023

4.9.1 Monitoring Metrics for Logstash Clusters in Cloud Eye..... 1023

4.9.2 Using Cloud Eye to Monitor Logstash Clusters..... 1030

4.9.3 Managing Logstash Cluster Logs..... 1033

4.10 Viewing Logstash Cluster Audit Logs..... 1040

5 CSS Resource Monitoring..... 1044

1 CSS Service Permission Management

1.1 Creating IAM Users and Granting Them Permissions to Use CSS

You can use [Identity and Access Management \(IAM\)](#) for fine-grained permissions control for CSS. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user has their own identity credentials for accessing CSS resources.
- Grant users only the permissions required to perform a given task based on their job responsibilities.
- Entrust an account or a cloud service to perform efficient O&M on your CSS resources.

If your Huawei Cloud account does not require individual IAM users, you may skip this section.


This section describes the procedure for granting permissions.

Prerequisites

You know what permissions to assign to which users. CSS supports administrator permissions and read-only permissions. For details, see [Permissions Management](#).

Process Flow

Figure 1-1 Process of granting CSS permissions



1. **Create a user group and assign permissions.**
Create a user group on the IAM console, and assign the **CSS ReadOnlyAccess** permission to the group.
2. **Create an IAM user and add it to a user group.**
Create a user on the IAM console and add it to the user group created in 1.
By adding a user to a user group, you assign all the permissions that this user group has to that user. All users in a group have the same permissions.
3. **Log in** and verify permissions.
Log in to the console as the created user, switch to the authorized region, and verify the permissions.
 - Choose **Service List > Cloud Search Service**. Then click **Create Cluster** on the CSS console. If the cluster cannot be purchased (assuming that the current permissions include only **CSS ReadOnlyAccess**), the **CSS ReadOnlyAccess** policy has already taken effect.
 - Choose any other service from **Service List**. (Assume that the current policy contains only **CSS ReadOnlyAccess**.) If a message appears indicating insufficient permissions to access the service, the **CSS ReadOnlyAccess** policy has already taken effect.

Examples of CSS Custom Policies

Custom policies can be created to supplement the system-defined policies of CSS. For the actions supported for custom policies, see [Permissions Policies and Supported Actions](#).

You can create custom policies in either of the following ways:

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.
- JSON: Create a JSON policy or edit an existing one.

For details about how to create custom policies, see [Creating a Custom Policy](#). The following provides examples of custom CSS policies.

NOTE

To be compatible with the open-source ecosystem, the IAM permissions and data plane cluster permissions of CSS are managed separately. To enhance data-plane security, you need to use the security mode.

To let an IAM user access an OBS bucket, you need to grant the **GetBucketStoragePolicy**, **GetBucketLocation**, **ListBucket**, and **ListAllMyBuckets** permissions to the user.

- Example 1: Allowing users to create CSS clusters

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Action": [
        "css:cluster:create",
        "vpc:securityGroups:get",
        "vpc:securityGroups:create",
        "vpc:securityGroups:delete",
        "vpc:securityGroupRules:get",
        "vpc:securityGroupRules:create",
        "vpc:securityGroupRules:delete",
        "vpc:vpcs:list",
        "vpc:privateIps:list",
        "vpc:ports:get",
        "vpc:ports:create",
        "vpc:ports:update",
        "vpc:ports:delete",
        "vpc:quotas:list",
        "vpc:subnets:get",
        "ecs:cloudServerFlavors:get",
        "ecs:serverInterfaces:use",
        "ecs:cloudServers:addNics",
        "ecs:quotas:get",
        "evs:types:get",
        "evs:quotas:get"
      ],
      "Effect": "Allow"
    }
  ]
}
```

- Example 2: Denying cluster deletion

A policy with only **Deny** permissions must be used in conjunction with other policies for it to take effect. If the permissions assigned to a user contain both **Allow** and **Deny**, the **Deny** permissions take precedence over the **Allow** permissions.

The following method can be used if you need to assign permissions of the **CSS Admin** policy to a user but you want to prevent the user from deleting

clusters. Create a custom policy for denying cluster deletion, and attach both policies to the group to which the user belongs. Then, the user can perform all operations on CSS except deleting clusters. The following is an example of a deny policy:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": [
        "css:cluster:delete"
      ]
    }
  ]
}
```

- Example 3: Defining permissions for multiple services in a policy

A custom policy can contain the actions of multiple services that are of the global or project-level type. Example of a policy containing multiple actions:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Action": [
        "ecs:cloudServers:resize",
        "ecs:cloudServers:delete",
        "ecs:cloudServers:delete",
        "css:cluster:restart",
        "css:*:get*",
        "css:*:list*"
      ],
      "Effect": "Allow"
    }
  ]
}
```

1.2 Example: Granting IAM Users the Permission to Create CSS Clusters

To implement fine-grained permission management for CSS, you can use Identity and Access Management (IAM) to create independent IAM users and assign policies or roles to IAM user groups. The policies and roles can be used to control access to CSS resources.

This section describes how to create an IAM user and add the IAM user to a user group, so that the IAM user has the permission to create CSS clusters.

1. **Creating a User Group:** The staff of company A is divided into an admin group and a developer group. The default group **admin** is generated after company A's administrator registers an account. The administrator needs to create another group in IAM for the developer team.
2. **Assigning Permissions to a User Group:** Developers in the company need to use cloud services CSS and OBS, so the administrator needs to assign the required permissions to the developer group to enable access to these services. After permissions are assigned, users in the two groups can access the corresponding resources.
3. **Creating IAM Users:** Create IAM users for your staff and add them to different groups. The users have their own credentials and can perform

operations on cloud resources based on the permissions inherited from their respective groups.

4. **Logging In Using an IAM User and Verifying Permissions:** After using the account of company A (**A-Company**) to create users **James, Alice, Charlie, Jackson, and Emily**, the administrator needs to send the account name, IAM user names, and IAM users' initial passwords to the corresponding employees, who can then use their own usernames and passwords to access the cloud platform. Then, they can create clusters on the CSS console and by doing so verify their permissions.

Creating a User Group

- Step 1** (Company A administrator) Log in to the Huawei Cloud management console.
 - Step 2** On the management console, hover the mouse pointer over the username in the upper right corner, and choose **Identity and Access Management** from the drop-down list.
 - Step 3** On the IAM console, choose **User Groups** and click **Create User Group**.
 - Step 4** Enter **Developers** for **Name**, and click **OK**.
- End

Assigning Permissions to a User Group


- Step 1** In the user group list, click **AuthorizeManage Permissions** in the row that contains the developer user group.
 - Step 2** In the **Select Policy/Role** step, search for **CSS FullAccess** and **OBS Administrator** in the search box, select them, and click **Next**.
 - Both the **CSS FullAccess** and **Elasticsearch Administrator** permissions enable users to create CSS clusters.
 - If users in the group need to view resource usage, attach the **BSS Administrator** role to the group for the same project.
 - Step 3** Select a scope. Set **Scope** to **Region-specific projects**, and select the region where your resources are.

Take the AP-Singapore region as an example. Set **Scope** to **Region-specific projects** and select **ap-southeast-3 [AP-Singapore]**.
 - Step 4** Click **OK**.
- End

Creating IAM Users


- Step 1** Log in to the IAM console as the administrator and choose **Users** from the left navigation pane.
- Step 2** On the **Users** page, click **Create User** in the upper right corner.

Figure 1-2 Creating a user



Step 3 Specify the user details and access type. To create more users, click **Add User**. A maximum of 10 users can be created at a time.

Figure 1-3 Configuring user information



NOTE

- Users can log in to Huawei Cloud using the username, email address, or mobile number.
- If users forget their passwords, they can reset them through email address or mobile number verification. If no email addresses or mobile numbers have been bound to users, users need to contact the administrator to reset their passwords.

Table 1-1 User information

Parameter	Description	Example Value
Username	(Mandatory) Username used by an IAM user to log in to Huawei Cloud.	James Alice
Email Address	Email address of the IAM user that can be used as a login credential. IAM users can bind an email address after they are created. This field is required if you have selected Set by user for Credential Type .	Skip it.
Mobile Number	Mobile phone number of the IAM user that can be used as a login credential. IAM users can bind a mobile number after they are created. This field is optional.	Skip it.
Description	Additional information about the IAM user. This field is optional.	Skip it.

Figure 1-4 Setting the access type

★ Access Type

☒ Programmatic access

Allows access to Huawei Cloud services only by using development tools, such as APIs, CLI, and SDKs, and requires an access key or password. ⓘ

☒ Management console access

Allows access to Huawei Cloud services only by using the management console and requires a password.

Credential Type

☒ Access key

You can download the access key after you create the user.

☒ Password

☒ Set now

👁

☒ Require password reset at first login

A password will be automatically generated. You can download the password file and provide it to the user.

☐ Automatically generated

A password will be automatically generated. You can download the password file and provide it to the user.

☐ Set by user

A one-time login URL will be emailed to the user. The user can then click on the link to set a password.

☐ USB Key

Give your account a security boost.

★ Login Protection

☐ Enable (Recommended)

☒ Disable

Table 1-2 Access types

Access Type	Description	Example Value
Programmatic access	An access key or password is generated for the IAM user. This type allows access to cloud services using development tools, such as APIs, CLI, and SDKs.	Select it.
Management console access	A password is generated for the IAM user. This type allows access to cloud services using the management console. If you select this parameter, Password must be selected for Credential Type .	Select it.

NOTE

- If the user accesses cloud services only by using the management console, select **Management console access** for **Access Type** and **Password** for **Credential Type**.
- If the user accesses cloud services only through programmatic calls, select **Programmatic access** for **Access Type** and **Access key** for **Credential Type**.
- If the user needs to use a password as the credential for programmatic access to certain APIs, select **Programmatic access** for **Access Type** and **Password** for **Credential Type**.
- If the user needs to perform access key verification when using certain services in the console, select **Programmatic access** and **Management console access** for **Access Type** and **Access key** and **Password** for **Credential Type**. For example, the user needs to perform access key verification when creating a data migration job in the Cloud Data Migration (CDM) console.

Table 1-3 Setting the credential type and login protection

Credential Type and Login Protection		Description	Example Value
Access key		An access key comprises an access key ID (AK) and secret access key (SK), and is used as a long-term identity credential to sign your requests for Huawei Cloud APIs . After creating the user, you can download the access key (AK/SK) access key (AK/SK) generated for the user.	Deselect it.
Password	Set now	Set a password for the user and determine whether to require the user to reset the password at first login. If you will use the IAM user by yourself, you are advised to select this option, set a password, and deselect Require password reset at first login .	Select it.
	Automatically generated	The system automatically generates a login password for the user. After the user is created, download the EXCEL password file and provide the password for the user. The user can then use this password for login. This option is available only when you create a single user.	-
	Set by user	A one-time login URL will be emailed to the user. The user can click the link to log in to the console and set a password. If you do not use the IAM user, select this option and enter the email address and mobile number of the IAM user. The user can then set a password by clicking the one-time login URL sent over email. The login URL is valid for two days .	-
Login Protection	Enabled	After login protection is enabled, IAM users need to enter verification codes in addition to their usernames and passwords during console login. It is a security practice to enable login protection. You can verify your identity by SMS, email, or virtual MFA device during login.	-

Credential Type and Login Protection		Description	Example Value
	Disabled	If login protection is disabled, you can enable it later by following the instructions provided in Viewing or Modifying IAM User Information .	Select it.


Step 4 Click **Next** and add the IAM users James and Alice to the created network O&M user group.

 NOTE

- If a user will be an administrator, add the user to the default group **admin**.
- Each user can be added to multiple user groups.

Step 5 Click **Next**. The created IAM users James and Alice are displayed in the user list. If you have selected **Access key** for **Credential Type** in step [Step 3](#), you can download the access key on the displayed page.

Figure 1-5 Users created successfully



-----End

Logging In Using an IAM User and Verifying Permissions

Step 1 Click **IAM User** **IAM User Login** **IAM User Login** on the login page, and then enter your **Tenant name or Huawei Cloud account name**, **IAM username or email address**, and **IAM user password** **Account name, IAM username or email**, and **Password** **Account name, IAM username or email**, and **Password**.

Figure 1-6 Logging in as an IAM user

The diagram illustrates the process of logging in as an IAM user. It shows two login forms side-by-side, connected by a red arrow pointing from left to right.

HUAWEI ID login (Left Form):

- Fields: Phone/Email/Login ID/HUAWEI CLOUD account name, Password.
- Buttons: LOG IN (red), Register, Forgot password?.
- Section: Use Another Account.
- Options: IAM User (highlighted with a red box), More (dropdown).
- Footer: Your account and network information will be used to help improve your login experience. [Learn more](#)

IAM User Login (Right Form):

- Fields: Company-A, Alice, Password (masked with asterisks).
- Buttons: Log In (red), Forgot Password, Remember me (checkbox).
- Footer: Use Another Account: HUAWEI ID | Federated User

Table 1-4 Login parameters

Parameter	Example	Description
Tenant name or Huawei Cloud account name	Company-A	Account used to create the IAM user, for example, Company-A.
IAM username or email address	Alice	IAM username or email address entered during the user creation. You can obtain the IAM username and IAM user's initial password from the administrator.
IAM user password	*****	Password of the IAM user, rather than the account. Enter the downloaded password.

Step 2 Click **Log In**.

Step 3 Choose **Service List**, search for CSS, and click the service name to go to the CSS console.

Step 4 In the upper right corner of the **Dashboard** page, click **Create Cluster**. Create a cluster by following the steps provided in [Creating a Cluster](#). If the cluster can be created, the permissions have taken effect.

----End

2 Using Elasticsearch for Data Search

2.1 Procedure for Using Elasticsearch

Table 2-1 Procedure for using an Elasticsearch cluster

Cate gory	Operation		Details
Use	Plann ing a cluste r	Before creating an Elasticsearch cluster, develop a plan for it, such as whether to deploy the cluster across multiple AZs to improve availability; the node quantity and specifications; the cluster version and security mode; and index sharding, in order to ensure the desired performance and reliability.	Elasticsearch Cluster Planning Suggestions
	Creati ng a cluste r	Create an Elasticsearch cluster based on the plan.	Creating an Elasticsearch Cluster (New Version)
	Acces sing a cluste r	There are many ways to access an Elasticsearch cluster, such as Kibana, Cerebro, open-source APIs, Java, Python, and Go clients, as well as multiple network configurations over an intranet and the public network. You can select the most appropriate access method based on the programming language you prefer as well as your network environment.	Elasticsearch Cluster Access Methods

Category	Operation	Details
	Importing data	There are many ways to import data to an Elasticsearch cluster, including Logstash, open-source Elasticsearch APIs, Cloud Data Migration (CDM), and Data Replication Service (DRS), with support for different data sources and formats, as well as real-time synchronization for relational databases. You can select the best way for yourself based on your use case and the characteristics of your data.
	Searching for data	With CSS, you are advised to use DSL for data search in Elasticsearch clusters. You may also use SQL.
	Enhancing the cluster's search capability	On top of the open-source version, CSS's Elasticsearch clusters offer a range of enhanced features, including vector search, storage-compute decoupling, flow control, large query isolation, aggregation enhancement, read/write splitting, switchover between hot and cold data storage classes, and index recycle bin. These features help you meet performance and cost optimization requirements for different use cases, while enhancing the service's cluster stability and search capability.
O&M	Backup and restoration	Creating Snapshots to Back Up the Data of an Elasticsearch Cluster Restoring the Data of an Elasticsearch Cluster Using a Snapshot

Category	Operation	Details
	Scaling a cluster	CSS provides flexible scale-out and scale-in options, using which you can add or reduce nodes (either randomly or with specified nodes), add node types, and increase or reduce node specifications. This allows you to dynamically adjust cluster resources to meet changing demand and optimize costs.
	Upgrade	Elasticsearch clusters support same-version upgrade, cross-version upgrade, and cross-engine upgrade. Same-version upgrade means to upgrade the kernel patches to fix problems or optimize performance. Cross-version upgrade means to upgrade the cluster version to enhance functionality or incorporate versions. Cross-engine upgrade means to upgrade an Elasticsearch cluster to an OpenSearch cluster.
	Scaling Out/Up an Elasticsearch Cluster Scaling In an Elasticsearch Cluster Upgrading the Version of an Elasticsearch Cluster	

Category	Operation	Details
	Managing clusters	CSS provides comprehensive cluster management functions. Users can view cluster information, authorize cluster access, change the cluster's security mode, manage tags, replace nodes, bind clusters with enterprise projects, switches AZs, and configure custom word dictionaries for Elasticsearch clusters. They help users efficiently manage Elasticsearch clusters and ensure cluster security, high availability, and optimized performance.
	Managing cluster index policies	The Index State Management (ISM) plug-in of Elasticsearch can be used to create and manage index lifecycle policies. These policies help automate index rollovers and deletions, helping optimize cluster performance and cut storage costs.

Category	Operation	Details
	Monitoring and log management	CSS provides comprehensive monitoring and log management functions. Users can configure and check monitoring metrics for clusters and nodes, configure alarm rules, and back up and view logs. Intelligent O&M tools help users efficiently monitor, analyze, and maintain Elasticsearch clusters and ensure cluster stability and performance.
	Audit logs	Cloud Trace Service (CTS) can be used to log mission-critical operations related to Elasticsearch clusters. Used for auditing and accountability purposes, these log records are retained for seven days on the management console.
		Elasticsearch Cluster Monitoring Metrics Configuring Elasticsearch Cluster Monitoring Setting Alarm Alerting via SMN for an Elasticsearch Cluster Intelligent Risk Detection for Elasticsearch Clusters Managing Elasticsearch Cluster Logs
		Viewing Elasticsearch Cluster Audit Logs

2.2 Elasticsearch Cluster Planning Suggestions

2.2.1 Planning Cluster AZs and HA

This topic describes how to improve cluster availability by deploying a cluster across multiple AZs, including the rules for distributing cluster nodes among AZs, the recommended number of replicas, and policies for handling AZ failures.

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

Multi-AZ deployment is a high availability feature provided by CSS. Deploying a cluster across two or three AZs located in the same region can help prevent data loss and lower the possibility of service outages.

Suggestions on Multi-AZ Deployment

If you select multi-AZ deployment when creating a cluster, CSS automatically enables cross-AZ HA to ensure that cluster nodes will be evenly distributed across the selected AZs (the difference between the number of nodes in each AZ cannot exceed 1).

You are advised to select three AZs, instead of two AZs, for multi-AZ deployment. If only two AZs are selected and one AZ becomes faulty, the cluster may not be able to elect a master node. As a result, the cluster may become unavailable.

Node Distribution Rules

When a multi-AZ cluster is created, nodes of all types are evenly distributed across different AZs. A maximum of three AZs can be used. [Table 2-2](#) shows the node distribution when the number of AZs varies.



WARNING

- When creating a multi-AZ cluster, ensure that the number of selected nodes of any type is greater than or equal to the number of AZs. Otherwise, multi-AZ cluster deployment will fail.
- If the number of data nodes or cold data nodes in a cluster is not divisible by the number of AZs, data in the cluster may be unevenly distributed, affecting data query or write performance.

Table 2-2 Node distribution when the numbers of nodes and AZs vary

Node Quantity	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
...	...	...	...	...	...	...

Replica Configuration Suggestions

For a multi-AZ cluster, configure the number of index replicas in a manner that can better capitalize on the high availability that comes with such as deployment.

- With a dual-AZ deployment, if one AZ becomes unavailable, the other AZ continues to provide services. In this case, configure at least one replica. If higher query performance is desired, you can increase the number of replicas.
- In the case of a three-AZ deployment, if one AZ becomes unavailable, the other AZs can continue to provide services. In this case, also configure at least one replica. To enhance the cluster's query performance, increase the number of replicas.

For Elasticsearch and OpenSearch clusters, the default number of index replicas is 1. To enable more replicas, modify settings. The following are two examples.

- Adjust the number of replicas for an existing index:

```
curl -XPUT http://ip:9200/{index_name}/_settings -d '{"number_of_replicas":2}'
```
- Set the number of replicas for new indexes using a template:

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

where, **ip** indicates the private IP address of the cluster; **index_name** indicates the index name; **templatename** indicates the template name; **template** indicates the index name matching rule (meaning the template will automatically apply to indexes that match this rule. The asterisk (*) indicates that the template will apply to all new indexes); and **number_of_replicas** indicates the number of index replicas to change to. In this example, the number of index replicas is changed to 2.

Service Outage Pattern Analysis

Master nodes manage cluster-wide operations, including metadata, indexes, and shard allocation. In a cluster with master nodes, the master nodes perform these tasks. In a cluster without them, data nodes and cold data nodes share the responsibilities of the master nodes.

In the case of a single AZ failure, if fewer than half (including half) of the nodes assuming master node responsibilities are still available, service availability will be affected. In this case, you need to restore services by referring to [Table 2-3](#).

For example, a cluster has three master nodes distributed in two AZs. If the AZ that contains one master node becomes unavailable while the other AZ remains available, services will still be available. However, if it is the other AZ, which contains two of the three master nodes, that fails, services will be interrupted, because fewer than half of the master nodes (or nodes that assume master node responsibilities) are available. In this case, restore services by referring to [Table 2-3](#).

Table 2-3 Possible service outage patterns in the face of the failure of a single AZ

Number of AZs	Number of Master Nodes	Service Outage Patterns and Handling Suggestions
2	0 (In this case, data nodes and cold data nodes share the responsibilities of master nodes.)	• If the number of data nodes plus cold data nodes is an even number: If half of the nodes become unavailable due to the failure of a single AZ, you must add one node to the still available AZ to ensure a master node can be elected for the cluster. To add nodes to your cluster, contact technical support. • If the number of data nodes plus cold data nodes is an odd number: – If the failed AZ is the one that has one more node, you must add two nodes to the surviving AZ to ensure a master node can be elected for the cluster. To add nodes to your cluster, contact technical support. – If the failed AZ is the other one—the one that has one less node, services will not be interrupted and a master node can be elected.

Number of AZs	Number of Master Nodes	Service Outage Patterns and Handling Suggestions
2	3	There is a 50% chance of service interruption when one AZ fails. In this situation, one AZ has two master nodes, and the other has one: • If the AZ that has one master node fails while the other AZ remains available, two master nodes are still available, so a master node can be elected from among them. In this case, services remain available. • If the failed AZ is the other one, only one master node survives, in which case, no master node can be elected, and services will be interrupted. In this case, contact technical support.
3	0 (In this case, data nodes and cold data nodes share the responsibilities of master nodes.)	There is a small chance of service interruption when a single node fails. However, there is one high-risk scenario: If the cluster spans three AZs and the total number of data nodes plus cold data nodes is 4, so the node distribution across the three AZs is 2-1-1, the failure of the AZ containing two nodes will cause a service interruption. To avoid this, avoid selecting four nodes for a three-AZ cluster.
3	3	A single AZ failure will not result in service interruptions.

2.2.2 Planning the Cluster Version and Security Mode

This topic introduces different Elasticsearch cluster versions, including features supported by each version and its application scenarios. It also describes the characteristics and application scenarios of different cluster security modes (including different application-layer protocols).

Introduction to Different Cluster Versions

When selecting an Elasticsearch cluster version, consider factors such as service requirements, available features, performance, security updates, and long-term support, ensuring that the selected version can meet both current and future needs and provide a stable, secure environment for your data.

- If you are deploying CSS Elasticsearch clusters for the first time, you are advised to use the latest version.
- If you are migrating an in-house built or third-party Elasticsearch cluster to CSS without altering the cluster, keep the version of the source cluster.
- If you are migrating an in-house built or third-party Elasticsearch cluster to CSS while recoding it, choose Elasticsearch 7.10.2.

Table 2-4 Features supported by different versions

Feature	Elasticsearch 7.6.2	Elasticsearch 7.10.2	Related Documents
Open-source vector search	x	x	-
CSS vector search	√	√	Configuring Vector Search for Elasticsearch Clusters
Decoupled storage and compute	√	√	Configuring Decoupled Storage and Compute for an Elasticsearch Cluster
Flow Control 2.0	√	√	Configuring Flow Control 2.0 for an Elasticsearch Cluster
Flow Control 1.0	√	√	Configuring Flow Control 1.0 for an Elasticsearch Cluster
Large query isolation	√	√	Configuring Large Query Isolation for an Elasticsearch Cluster
Enhanced aggregation	x	√	Configuring Enhanced Aggregation for an Elasticsearch Cluster
Read/write splitting	√	√	Configuring Read/Write Splitting Between Two Elasticsearch Clusters
Switching between hot and cold storage	√	√	Switching Between Hot and Cold Storage for an Elasticsearch Cluster
Index recycle bin	x	√	Configuring an Index Recycle Bin for an Elasticsearch Cluster
Enhanced import performance	x	√	Enhancing the Data Ingestion Performance of Elasticsearch Clusters
Enhanced kernel monitoring	√	√	Configuring Kernel Monitoring for an Elasticsearch Cluster
Index monitoring	√	√	Configuring Index Monitoring for an Elasticsearch Cluster

Cluster Security Modes

Configure the security mode based on the security needs of your cluster.

Table 2-5 Cluster security modes

Cluster Type	Key Settings	Description	Applicable Scenario
Non-security mode cluster	Security Mode: disabled	Access to such a cluster requires no user authentication, and data will be transmitted in plaintext using HTTP.	Use when creating a cluster for internal testing or workloads that have a low security standard. • Advantage: easy to access the cluster. • Disadvantage: poor security, as anyone can access the cluster. When the security mode is disabled, public network access and Kibana public network access cannot be enabled. Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.

Cluster Type	Key Settings	Description	Applicable Scenario
Security-mode cluster	Security-mode cluster + HTTP: • Security Mode: enabled • HTTPS Access: disabled	Such a cluster requires user authentication. It supports access control and data encryption, and it uses HTTP to transmit data in plaintext.	Use to balance security and performance. • Advantage: User authentication improves cluster security. HTTP-based access ensures high performance of the cluster. User permissions can be configured to ensure proper isolation. • Disadvantage: Public network access is not supported. Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.
	Security-mode cluster + HTTPS: • Security Mode: enabled • HTTPS Access: enabled	Such a cluster requires user authentication. It supports access control and data encryption, and it uses HTTPS to encrypt communication and enhance data security.	Use when security takes precedence over performance and public network access is required. • Advantage: User authentication improves cluster security. HTTPS enhances cluster security by encrypting all communication over the public network. User permissions can be configured to ensure proper isolation. • Disadvantage: When HTTPS is used, data encryption and decryption introduce computational overhead and impact the cluster's read and write performance.

To access a security-mode cluster, a username and password need to be provided. CSS supports authentication for the following two types of users:

- Administrator: The default administrator username is **admin**, and the password is the one specified during cluster creation.
- Cluster user: created by the cluster administrator on Kibana. User permissions can be configured to ensure proper isolation. For details, see [Creating Users for an Elasticsearch Cluster and Granting Cluster Access](#).

2.2.3 Planning Node Types and Specifications

This topic describes the application scenarios and configuration suggestions for each type of cluster nodes, including data nodes, master nodes, client nodes, and cold data nodes. The aim is to help you properly select and configure nodes for your cluster based on service requirements. It also provides suggestions on selecting node specifications and configuring the node storage type, storage capacity, and node quantity, helping you properly plan the capacities of your cluster.

Planning Node Types

Before creating a cluster, determine the types of nodes to use based on service requirements, query load, data growth patterns, and performance goals.

For each cluster, data nodes are mandatory, while master, client, and cold data nodes are optional. Choose whether to enable the latter three based on your service needs and performance requirements.

Table 2-6 Introduction to different node types

Node Type	Suggestions on Node Selection
Data node (ESS)	Data nodes are primarily used to store the cluster's data. The functions of the data nodes vary depending on the node types enabled for the cluster. • If Master node and Client node are both selected, data nodes will be used for data storage only. • If Master node and Client node are both unselected, data nodes will be used for all of the following purposes: cluster management, data storage, cluster access, and data analysis. To ensure reliability, a cluster should have a least three nodes. • If Master node is selected but Client node is not, data nodes will be used for data storage, cluster access, and data analysis. • If Master node is unselected but Client node is selected, data nodes will be used for data storage and cluster management.

Node Type	Suggestions on Node Selection
Master node (ess-master)	Master nodes manage cluster-wide operations, including metadata, indexes, and shard allocation. For large-scale deployments, using dedicated master nodes enhances cluster stability, service availability, and centralized control. • Large-scale cluster: For a cluster that has more than 10 nodes, you are advised to add dedicated master nodes to effectively manage the cluster status and metadata. • Large quantities of indexes and index shards: If the number of indexes or shards exceeds 10,000, master nodes will have better performance handling complex cluster management tasks, avoiding impact on the performance of data nodes. • Better management of cluster nodes: Master nodes maintain the cluster metadata, including index mapping, settings, and aliases. For a complex cluster structure, dedicated master nodes offer better management, including node joining, exiting, and fault detection. Master nodes play a critical role in cluster node management. • Improved cluster stability and reliability: Dedicated master nodes improve cluster stability and reliability by taking over cluster management responsibilities from data storage and query nodes. • Optimized performance for data nodes: By offloading cluster management tasks from data nodes to master nodes, you can allow data nodes to focus on data processing, which leads to improved performance.

Node Type	Suggestions on Node Selection
Client node (ess-client)	Client nodes route and coordinate search and index requests, offloading processing from data nodes for enhanced query performance and cluster scalability when there are heavy loads. • High queries per second (QPS): In the face of a high QPS, a dedicated client node can evenly distribute query requests, reducing the load of data nodes and improving the overall query performance. • Complex aggregation queries: For complex, compute-intensive aggregation queries, a client node can dedicate to the handling of aggregation results, thus improving the efficiency and response speed of such queries. • Large number of shards: In a cluster with a large number of shards, a client node can effectively coordinate and manage query requests to each shard, improving efficiency in request forwarding and processing. • Reducing the load of data nodes: A client node parses search requests, determines the locations of index shards, and coordinates different nodes to execute searches. This reduces the load of data nodes by allowing them to focus on data storage and indexing. • Improved cluster scalability: The use of client nodes allows for better cluster scalability and flexibility, enabling support for large datasets and more complex query requirements.

Node Type	Suggestions on Node Selection
Cold data node (ess-cold)	Cold data nodes are used to store and query latency-insensitive data in large quantities. They offer an effective way to manage large datasets while cutting storage costs. • Storage of historical data in large quantities: Cold data nodes offer a more cost-effective solution for storing large quantities of historical data that are infrequently accessed but useful for analytical purposes. • Optimizing hot data performance: By migrating cold data to cold data nodes, you reduce the storage load of hot data nodes, thereby optimizing their query and write performance. • Insensitivity to query latency: Cold data nodes are a better option for storing data that is insensitive to a high query latency. • Cost-effectiveness: Cold data nodes usually use large-capacity disks that offer inexpensive storage. CAUTION If no cold data nodes were enabled during cluster creation, they cannot be added later, so you have to determine whether to use cold data nodes while creating a cluster.

 NOTE

- If no master or client nodes were enabled when a cluster was created, you can add them if data nodes become overloaded later at some point. For details, see [Adding New Node Types](#).
- When cold data nodes are enabled, you can switch between cold and hot data storage. For details, see [Switching Between Hot and Cold Storage for an Elasticsearch Cluster](#). If there are no cold data nodes, we recommend that you use decoupled storage and compute, which can also cut storage costs. For details, see [Configuring Decoupled Storage and Compute for an Elasticsearch Cluster](#).

Suggestions on Data Node Configuration

Table 2-7 Data node configuration

Parameter	Configuration Suggestions
Node Specifications	In the node flavor list, vCPUs Memory indicate the number of vCPUs and memory capacity available for each flavor, and Recommended Storage indicates the supported storage capacity range. We recommend that you select node specifications based on service needs, such as the data volumes, performance requirements, and your spending budget. Node Specifications describes the application scenarios and core features of different node specifications. It can help you properly plan your cluster. For more information about different node specifications, see ECS Types.
Node Storage Type and Capacity	• If the selected node flavor uses EVS disks, you need to further select Node Storage Type and Capacity based on service requirements. – For more on EVS disk performance, see Disk Types and Performance. – The value range of node storage capacity is determined by the node flavor you select. The value must be divisible by 20. For how to calculate the required node storage capacity, see Recommended Node Storage Capacity. – Node storage capacity cannot be reduced once the cluster is created. Evaluate your long-term data needs and select an appropriate size. • If the selected node flavor uses local disks, there is no need to select the node storage type, and the node storage capacity is a fixed value. Both of them are determined by the local disk specifications.
Nodes	• If master nodes are configured, the number of data nodes ranges from 1 to 200. • If no master nodes are configured, the number of data nodes ranges from 1 to 32. • To ensure cluster availability, you should configure at least three data nodes. For how to calculate the number of nodes needed, see Recommended Node Quantity. If the number of data nodes in your cluster is not evenly divisible by the number of AZs, data distribution may become unbalanced across nodes. This will negatively impact both query and write performance.

Node Specifications

The choice of node specifications affects disk deployment. There are two types of disks: local disks and EVS disks. For example, high-performance nodes may use local disks, while standard nodes may use cost-effective EVS disks.

- Local disks refer to physical storage directly attached to the host machines running your ECSs. They deliver high I/O performance, low latency, and high throughput, making them ideal for performance-sensitive workloads that do not require long-term storage.
- EVS disks are a virtual block storage service that is independent of ECSs. They provide high reliability and fast elasticity, making them ideal for workloads that require high data reliability and highly scalable storage capacity.



WARNING

- Local disks are less reliable than EVS disks—if the underlying host fails, data stored on local disks may be permanently lost. For this reason, avoid using local disks to store mission-critical data. Instead, use EVS disks.
- If you do use local disks, you are advised to enable automatic snapshot creation when creating the cluster, and also enable replicas when creating indexes. This helps lower the risk of data loss.

The table below describes the application scenarios and core features of different node specifications. It can help you properly plan your cluster.

Table 2-8 Comparing different node specifications

CPU Architecture	Node Flavor	Description
x86	Compute-intensive	Core advantages • High-performance CPU: designed for high computational load and low latency, ideal for real-time search and complex queries. • High reliability: Local NVMe disks enhance cluster stability. Application scenarios • E-commerce search and recommendations: real-time responses to search requests, complex filtering and ranking supported. • App search service: quick search on mobile apps, such as social apps and content recommendations. • Database acceleration: low latency for complex OLAP queries. Precautions • This flavor is suitable for performance-demanding workloads, but it comes at a higher cost. Be sure to balance your performance needs with budget considerations. • Use ultra-high I/O disks with this node flavor, making sure the data read/write performance matches the CPU performance.
	General computing-plus - AC	Core advantages • Dedicated CPU: no resource contention between different instances, stable performance at a relatively low cost. • Low latency: ideal for latency-sensitive workloads. Application scenarios • High-concurrency search: e-commerce search, social platform search, etc., where fast responses are required. • Real-time recommendation system: recommendations are generated based on real-time user behavior. • OLTP scenarios: transactional data processing and real-time queries.

CPU Architecture	Node Flavor	Description
	General computing	Core advantages • Balanced configuration: default specifications, suitable for medium-scale workloads. • Cost-effective: able to meet the needs of most general search and analytical applications. Application scenarios This flavor meets the needs of general use cases through a standard deployment, without the need for special tuning. • Content search: medium-scale search and analytical workloads. (The data volume per node ranges from 100 GB to 1,000 GB.) • Log analytics: real-time query and analysis of medium-scale log data.
	Memory-optimized	Core advantages • Large memory capacity: ideal for memory-intensive tasks, such as complex aggregation and caching. • High throughput: fast processing of large data volumes. Application scenarios Use this flavor for search and analytical workloads where the per-node data volume ranges from 100 GB to 2,000 GB. • Large-scale data analytics, such as user behavior analysis and statistics. • OLAP scenarios: complex queries and multi-dimensional aggregation. • Cache-intensive workloads, such as real-time reporting, where data needs to be frequently loaded to the memory. Precautions Avoid using this flavor for latency-sensitive workloads. This flavor comes at a relatively high cost, primary due to high memory usage. Be sure to balance your performance needs with budget considerations.

CPU Architecture	Node Flavor	Description
	Ultra-high I/O	Core advantages - Local NVMe SSDs: impressively high disk I/O performance, ideal for workloads with high-concurrency reads/writes. - Low latency: Compared with cloud-based SSDs, local NVMe SSDs have lower latency and better performance. Application scenarios - Real-time public sentiment analysis: frequent writes and fast queries. - Patent search: quick search and matching of large volumes of text data. - Database acceleration: read/write splitting architecture for MySQL/PostgreSQL. Precautions - In terms of data loss prevention, local disks are less reliable than cloud-based risks. When using local disks, be sure to enable data replicas. - NVMe SSDs are expensive. Be sure to balance performance and costs.
Kunpeng	Kunpeng general computing	Core advantages Cost-effective: In general, the Arm architecture is more cost-effective than x86. Application scenarios - Cost-sensitive workloads, such as search services for small and medium enterprises and testing environments. - Arm ecosystem compatibility: when your search workloads must be compatible with Arm. Precautions Ecosystem compatibility: Make sure your applications (including Java virtual machine and third-party plugins) are compatible with the Arm architecture.

Recommended Node Storage Capacity

If you select EVS disks, you can select a node storage capacity within the range supported by the selected node flavor.

When planning the storage capacity of a cluster, consider the following factors: the original data size, number of data replicas, data bloat rate, and disk usage.

The following is a recommended formula for determining the needed cluster storage capacity:

Storage capacity = Original data size x (1 + Number of replicas) x (1 + Data bloat rate) x (1 + Ratio of reserved space)

where,

- Original data size: Determine the size of original data that needs to be stored.
- Number of replicas: The default value is 1.
- Data bloat rate: Extra data may be generated due to data indexing. Generally, you are advised to use a 25% data bloat rate.
- Disk usage: Considering the space occupied by the operating system and file system and the space reserved for optimized disk performance and redundancy, you are advised to keep the disk usage under 70%. That is, you need to reserve 30% of the total disk capacity.

Using the recommended values above, the formula should be: Storage capacity = Original data size x 2 x 1.25 x 1.3 = Original data size x 3.25

This formula is for quick reference only. You still need to adjust it based on the actual applications and projected data growth rate. Moreover, the storage capacity must be divisible by 20.

Recommended Node Quantity

Data nodes and cold data nodes are used to store data. Their quantities largely determine the cluster performance baseline. The following formulas can be used to estimate the number of data nodes and cold data nodes needed:

Number of data nodes + Number of cold data nodes = Number of write nodes + Number of query nodes

Number of write nodes = Traffic during peak hours (MB/s)/Number of vCPUs per node/Per-vCPU write performance baseline (MB/s) x Number of index replicas

Number of query nodes = Query QPS/(Number of vCPUs per node x 3/2/Average query latency (s)) x Number of index shards

where,

- The number of vCPUs per node is determined by the selected node flavor.
- The per-vCPU write performance baseline is determined by the selected disk type. For a node that uses EVS disks, the per-vCPU write performance baseline is 1 MB/s. For a node that uses local disks, this baseline is 1.5 MB/s.
- QPS indicates the number of queries that can be processed per second.
- The average query latency is used as the query performance baseline.
- Where a high write performance is desired, we recommend setting the number of index shards to equal the total number of data nodes plus cold data nodes, thus evenly distributing the write load across all nodes.
- Where a high QPS is desired, we recommend setting the number of index shards to equal the total number of data nodes plus cold data nodes **minus 1**, thus evenly distributing the query load across all nodes.

Example:

If the peak traffic is 100 MB/s; the node specifications are 16 vCPUs and 64 GB RAM, with EVS disks; the query throughput is 1,000 QPS; the average query latency is 100 ms; and each index has three shards and two replicas, then:

Number of write nodes = $100/16/1 \times 2 = 12$

Number of query nodes = $1000/(16 \times 3/2/0.1) \times 3 = 12$

Number of data nodes + Number of cold data nodes = Number of write nodes + Number of query nodes = $12 + 12 = 24$

Suggestions on Master Node Configuration

Table 2-9 Master node configuration

Parameter	Configuration Suggestions
Node Specifications	In the node flavor list, vCPUs Memory indicate the number of vCPUs and memory capacity available for each flavor, and Recommended Storage indicates the supported storage capacity range. For a large-scale cluster with a large number of index shards, use large-capacity nodes as master nodes. Master nodes support EVS disks only. For more information about different node specifications, see ECS Types.
Node Storage Type and Capacity	Select an appropriate storage type and capacity for master nodes. - For more on EVS disk performance, see Disk Types and Performance. - The node storage capacity is fixed at 40 GB.
Nodes	The number of master nodes can be 3, 5, 7, or 9. For split-brain protection, the number of master nodes must be an odd number. For the recommended master node quantity, see Recommended Master Node Quantity.

Recommended Master Node Quantity

Increasing the number of master nodes can improve the cluster's fault tolerance performance, but it can also lead to issues like increased communication overhead, longer election time, and low resource utilization. You are advised to select an appropriate number of master nodes based on service requirements.

The following table provides the recommended number of master nodes based on the cluster size.

Table 2-10 Recommended number of master nodes

Cluster Size	Recommended Number of Master Nodes	Performance Expectation
Small (< 10 nodes)	3	Basic high availability
Medium (50 > number of nodes ≥ 10)	3 or 5	Balanced performance and high availability
Large (≤ 50 nodes)	5 or 7	Enhanced fault tolerance and stability
Extra-large (≥ 100 nodes)	5 or 7	Pay attention to excessive communication overhead.

Suggestions on Client Node Configuration

Table 2-11 Client node configuration

Parameter	Configuration Suggestions
Node Specifications	In the node flavor list, vCPUs Memory indicate the number of vCPUs and memory capacity available for each flavor, and Recommended Storage indicates the supported storage capacity range. For a cluster with heavy read and write traffic, use large-capacity nodes as client nodes. Client nodes support EVS disks only. For more information about different node specifications, see ECS Types.
Node Storage Type and Capacity	Select an appropriate storage type and capacity for client nodes. - For more on EVS disk performance, see Disk Types and Performance. - The node storage capacity is fixed at 40 GB.
Nodes	The number of client nodes ranges from 1 to 32. For the recommended client node quantity, see Recommended Client Node Quantity.

Recommended Client Node Quantity

Client nodes route and coordinate search and index requests. Its quantity should be determined based on the cluster traffic. For your reference, the following table provides the recommended number of client nodes based on the number of data nodes.

Table 2-12 Recommended number of client nodes

Number of Data Nodes	Recommended Number of Client Nodes
< 5	Do not use client nodes.
≥ 5 and < 10	2
≥ 10	Number of client nodes/Number of data nodes = 1:5

Suggestions on Cold Data Node Configuration

Table 2-13 Cold data node configuration

Parameter	Configuration Suggestions
Node Specifications	In the node flavor list, vCPUs Memory indicate the number of vCPUs and memory capacity available for each flavor, and Recommended Storage indicates the supported storage capacity range. We recommend that you select node specifications based on service needs, such as the data volumes, performance requirements, and your spending budget. Node Specifications describes the application scenarios and core features of different node specifications. It can help you properly plan your cluster. For more information about different node specifications, see ECS Types.
Node Storage Type and Capacity	• If the selected node flavor uses EVS disks, you need to further select Node Storage Type and Capacity based on service requirements. – For more on EVS disk performance, see Disk Types and Performance. – The value range of node storage capacity is determined by the node flavor you select. The value must be divisible by 20. For how to calculate the required node storage capacity, see Recommended Node Storage Capacity. – Node storage capacity cannot be reduced once the cluster is created. Evaluate your long-term data needs and select an appropriate size. • If the selected node flavor uses local disks, there is no need to select the node storage type, and the node storage capacity is a fixed value. Both of them are determined by the local disk specifications.

Parameter	Configuration Suggestions
Nodes	The value range for cold data nodes is 1 to 32. For how to calculate the number of nodes needed, see Recommended Node Quantity. If the number of cold data nodes in your cluster is not evenly divisible by the number of AZs, data distribution may become unbalanced across nodes. This will negatively impact both query and write performance.

2.3 Creating an Elasticsearch Cluster (New Version)

You can use Elasticsearch clusters created in CSS for purposes like log analytics, data search, real-time data analytics, and more.

You can create an Elasticsearch cluster in either of the following ways:

- Method 1: Create an Elasticsearch cluster on the CSS management console. This topic provides the operation guide for the default UI (that is, the new UI). For the operation guide for the old UI, see [Creating an Elasticsearch Cluster \(Old Version\)](#).

NOTE

To improve efficiency in creating clusters, CSS has made several enhancements on the cluster creation page in terms of ease-of-use. Being more intuitive, the new UI streamlines cluster creation procedures.

- Method 2: Create an Elasticsearch cluster using CSS APIs. For details, see [Creating a Cluster](#).

Prerequisites

- You have planned the Elasticsearch cluster configuration by following the guidelines in [Elasticsearch Cluster Planning Suggestions](#), including the cluster size, node configuration, and storage capacity requirements.
- You have created a VPC, subnet, and security group beforehand, or you can create them when creating the cluster. For more information, see [Virtual Private Cloud User Guide](#).
- If you are going to enable automatic snapshot creation for the cluster, make sure you have created an OBS bucket for storing snapshots. The OBS bucket must be in the same region as the Elasticsearch clusters, and the storage class must be **Standard** or **Warm**.


Impact on Billing

- If you enable automatic snapshot creation, automatically created snapshots will be stored in standard OBS buckets, and you will need to pay extra for the storage. For more information, see [Billing Overview](#).
- VPC endpoints, if created for the cluster, will incur extra fees, depending on the resource usage. For details, see [VPC Endpoint Billing](#).

Opening the Create Cluster Page

- 1. Log in to the [CSS management console](#).
- 2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
- 3. In the upper right corner, click **Create Cluster**. The new-version UI for creating a cluster is displayed by default.

Figure 2-1 Create Cluster (new version)



Cluster Configuration

Select the cluster type and version.

Table 2-14 Cluster configuration

Parameter	Description
Cluster Type	Select Elasticsearch .
Cluster Version	Select a cluster version from the drop-down list.

Basic Settings

Select the region, billing mode, and AZ of the cluster.

Table 2-15 Basic settings

Parameter	Description
Region	Select the region where the cluster is located. A region is the location of a physical data center. Regions are defined based on their geographical location and network latency. For lower network latency and quicker resource access, select the nearest region.

Parameter	Description
AZ	Select one or more AZs associated with the cluster region. An AZ is a physical region where resources use independent power supplies and networks. AZs are physically isolated but interconnected through an internal network. A maximum of three AZs can be configured. For details about the use of multiple AZs, see Suggestions on Multi-AZ Deployment.
Billing Mode	Billing mode for the cluster, which can be Yearly/Monthly or Pay-per-use. • Yearly/Monthly: You prepay for a yearly or monthly subscription. This mode is cost-effective when the duration of your service usage is predictable. It is recommended for long-term users. A yearly/monthly subscription cannot be deleted. If you no longer need it, unsubscribe from it. • Pay-per-use (postpaid): You will be billed hourly by actual duration of use. Any partial hour of usage will be rounded up to one hour. This mode is ideal for short-term or unpredictable workloads. Pay-per-use resources can be provisioned and deleted at any time.

Data Nodes

Configure the cluster's data nodes.

Data nodes are primarily used to store the cluster's data. The functions of the data nodes vary depending on the node types enabled for the cluster.

- If the cluster has no client nodes, the data nodes will need to handle the additional task of query request distribution and data analysis.
- If the cluster has no master nodes, the data nodes will also need to handle cluster management.

Table 2-16 Configuring data nodes

Parameter	Description
CPU Architecture	Select the CPU architecture of the data nodes. x86 and Kunpeng nodes are supported. The architectures actually supported may vary depending on the regional environment.

Parameter	Description
Node Specifications	Select the specifications of the data nodes. Click Available. On the displayed page, select a flavor that suits your needs. The nodes may use either local or EVS disks. For more information about different node specifications, see ECS Types. In the node flavor list, vCPUs Memory indicate the number of vCPUs and memory capacity available for each flavor, and Recommended Storage indicates the supported storage capacity range. The node flavors available may vary depending on the region you select. WARNING • Local disks are less reliable than EVS disks—if the underlying host fails, data stored on local disks may be permanently lost. For this reason, avoid using local disks to store mission-critical data. Instead, use EVS disks. • If you do use local disks, you are advised to enable automatic snapshot creation when creating the cluster, and also enable replicas when creating indexes. This helps lower the risk of data loss.
Node Storage Type and Capacity	Select the storage type and capacity of the data nodes. • If the selected node flavor uses EVS disks, you need to further select Node Storage Type and Capacity based on service requirements. – Available EVS disk types vary depending on your region. For more on EVS disk performance, see Disk Types and Performance. – The value range of node storage capacity is determined by the node flavor you select. The value must be divisible by 20. – Node storage capacity cannot be reduced once the cluster is created. Evaluate your long-term data needs and select an appropriate size. • If the selected node flavor uses local disks, there is no need to select the node storage type, and the node storage capacity is a fixed value. Both of them are determined by the local disk specifications.

Parameter	Description
Nodes	Set the number of data nodes for the cluster. • If master nodes are configured, the number of data nodes ranges from 1 to 200. • If no master nodes are configured, the number of data nodes ranges from 1 to 32. • To ensure cluster availability, you should configure at least three data nodes. If the number of data nodes in your cluster is not evenly divisible by the number of AZs, data distribution may become unbalanced across nodes. This will negatively impact both query and write performance.

Master Nodes

Configure master nodes for the cluster. By default, master nodes are not enabled. To enable master nodes, select **Master Node**, and configure the node settings.

Master nodes manage cluster-wide operations, including metadata, indexes, and shard allocation. For large-scale deployments, using dedicated master nodes enhances cluster stability, service availability, and centralized control.

Table 2-17 Configuring master nodes

Parameter	Description
CPU Architecture	Select the CPU architecture of the master nodes. x86 and Kunpeng nodes are supported. The architectures actually supported may vary depending on the regional environment.
Node Specifications	Select the specifications of the master nodes. Click Available. On the displayed page, select a flavor that suits your needs. Master nodes support EVS disks only. For more information about different node specifications, see ECS Types. In the node flavor list, vCPUs Memory indicate the number of vCPUs and memory capacity available for each flavor. The node flavors available may vary depending on the region you select.

Parameter	Description
Node Storage Type and Capacity	Select the storage type and capacity of the master nodes. - Available EVS disk types vary depending on your region. For more on EVS disk performance, see Disk Types and Performance. - The node storage capacity is fixed at 40 GB.
Nodes	Set the number of master nodes for the cluster. The value can be 3 , 5 , 7 , or 9 .

Client Nodes

Configure client nodes for the cluster. By default, client nodes are not enabled. To enable client nodes, select **Client Node**, and configure the node settings.

Client nodes route and coordinate search and index requests, offloading processing from data nodes for enhanced query performance and cluster scalability when there are heavy loads.

Table 2-18 Configuring client nodes

Parameter	Description
CPU Architecture	Select the CPU architecture of the client nodes. x86 and Kunpeng nodes are supported. The architectures actually supported may vary depending on the regional environment.
Node Specifications	Select the specifications of the client nodes. Click Available . On the displayed page, select a flavor that suits your needs. Client nodes support EVS disks only. For more information about different node specifications, see ECS Types . In the node flavor list, vCPUs Memory indicate the number of vCPUs and memory capacity available for each flavor. The node flavors available may vary depending on the region you select.
Node Storage Type and Capacity	Select the storage type and capacity of the client nodes. - Available EVS disk types vary depending on your region. For more on EVS disk performance, see Disk Types and Performance. - The node storage capacity is fixed at 40 GB.

Parameter	Description
Nodes	Set the number of client nodes for the cluster. Value range: 1 to 32

Cold Data Nodes

Configure cold data nodes for the cluster. By default, cold data nodes are not enabled. To enable cold data nodes, select **Cold Data Node**, and configure the node settings. Cold data nodes must be configured at cluster creation and cannot be modified later.

Cold data nodes are used to store and query large quantities of latency-insensitive data. They offer an effective way to manage large datasets while cutting storage costs. When cold data nodes are enabled, you can switch between cold and hot data storage. For details, see [Switching Between Hot and Cold Storage for an Elasticsearch Cluster](#). If there are no cold data nodes, we recommend that you use decoupled storage and compute, which can also cut storage costs. For details, see [Configuring Decoupled Storage and Compute for an Elasticsearch Cluster](#).

Table 2-19 Configuring cold data nodes


Parameter	Description
CPU Architecture	Select the CPU architecture of the cold data nodes. x86 and Kunpeng nodes are supported. The architectures actually supported may vary depending on the regional environment.
Node Specifications	Select the specifications of the cold data nodes. Click Available. On the displayed page, select a flavor that suits your needs. The nodes may use either local or EVS disks. For more information about different node specifications, see ECS Types. In the node flavor list, vCPUs Memory indicate the number of vCPUs and memory capacity available for each flavor, and Recommended Storage indicates the supported storage capacity range. The node flavors available may vary depending on the region you select. WARNING - Local disks are less reliable than EVS disks—if the underlying host fails, data stored on local disks may be permanently lost. For this reason, avoid using local disks to store mission-critical data. Instead, use EVS disks. - If you do use local disks, you are advised to enable automatic snapshot creation when creating the cluster, and also enable replicas when creating indexes. This helps lower the risk of data loss.

Parameter	Description
Node Storage Type and Capacity	Select the storage type and capacity of the cold data nodes. • If the selected node flavor uses EVS disks, you need to further select Node Storage Type and Capacity based on service requirements. – Available EVS disk types vary depending on your region. For more on EVS disk performance, see Disk Types and Performance. – The value range of node storage capacity is determined by the node flavor you select. The value must be divisible by 20. – Node storage capacity cannot be reduced once the cluster is created. Evaluate your long-term data needs and select an appropriate size. • If the selected node flavor uses local disks, there is no need to select the node storage type, and the node storage capacity is a fixed value. Both of them are determined by the local disk specifications.
Nodes	Set the number of cold data nodes for the cluster. Value range: 1 to 32 If the number of cold data nodes in your cluster is not evenly divisible by the number of AZs, data distribution may become unbalanced across nodes. This will negatively impact both query and write performance.

Network Settings

Configure the VPC, subnet, IP addresses, and security group of the cluster.

Table 2-20 Network settings

Parameter	Description
VPC	Select a VPC for the cluster for proper network isolation. You can select a shared VPC, which is indicated by Shared VPC in the VPC name. To create a CSS cluster in a shared VPC, the default vpc subnet statement permission is required. VPC  VPC sharing allows you to centrally manage resources across multiple accounts, helping to improve resource management efficiency and reduce O&M costs. For more information about shared VPCs, see VPC Sharing. If none of the existing VPCs meets your requirements, click Create VPC to create one. (Make sure you have the permission to create VPCs.) For details, see Creating a VPC and Subnet.
Subnet	Select a subnet for the cluster. A subnet improves network security by providing exclusive network resources that are isolated from other networks. Select a subnet in the current VPC. Or you can select a subnet under a shared VPC. The owner of a VPC can share the VPC to specified users.

Parameter	Description
IPv4 Address	Assign IPv4 addresses to cluster nodes. • Assign automatically: Private IPv4 addresses will be automatically assigned to cluster nodes. • Assign manually: Manually assign private IPv4 addresses to cluster nodes. Before assigning an IP address, click View In-Use IP Address to check IP addresses that are already in use. – If you manually specify fewer IP addresses than the total number of cluster nodes, the system will automatically assign random private IPv4 addresses to the remaining nodes. – Manually specified IP addresses will be assigned in the following priority order: client nodes > data nodes > cold data nodes > master nodes. Default value: Assign automatically
Security Group	Select a security group for the cluster. A security group serves as a virtual firewall that provides access control policies for clusters. To select a security group that meets your requirements, click View Security Group to go to the security group list, where you can check the details of each security group. The selected security group must allow port 9200 in the inbound direction. Otherwise, the cluster may be inaccessible to external services. Furthermore, to enable Kibana private network access, port 5601 must also be allowed. If none of the existing security groups meets your requirements, click Create Security Group to create one. (Make sure you have the permission to create security groups.) For details, see Creating a Security Group.

Security Mode

Configure the cluster's security mode and access methods. The security mode is enabled by default when you create a cluster. To disable it, deselect **Security Mode**.

Configure the security mode based on the security needs of your cluster.

- **Disabling the security mode:** Use when creating a cluster for internal testing or workloads that have a low security standard. When the security mode is disabled, HTTP is used, data is transmitted in plaintext, and no user authentication is required for cluster access. Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.

- **Security mode + HTTP:** Use to balance security and performance. In this setting, user authentication is required for cluster access, yet data is transmitted via HTTP.
- **Security mode + HTTPS:** Use when security takes precedence over performance. In this setting, user authentication is required for cluster access, and data is transmitted via HTTPS in ciphertext.

 CAUTION

HTTPS, public network access, and Kibana public network access can be enabled only when the security mode is enabled.

Table 2-21 Configuring the security mode

Parameter	Description
Administrator account and password	When the security mode is enabled, cluster access requires user authentication. In this case, the administrator password must be configured. • Administrator: The default value is admin and cannot be changed. • Administrator Password: Set and confirm the administrator password. This password will be required when you access this cluster.
HTTPS Access	Whether to enable HTTPS access for the cluster. Enabled by default, HTTPS improves security by encrypting all communication with the cluster. To use HTTP instead, deselect HTTPS Access. Data security cannot be guaranteed when HTTPS is disabled. Exercise caution. When HTTPS is used, data encryption and decryption introduce computational overhead, which may impact the cluster's read and write performance.

Parameter	Description
Public Network Access	Whether to enable public network access. By default, it is disabled. To enable it, select Public Network Access, and configure the necessary settings. Public network access can be enabled only if the security mode and HTTPS are both enabled for the cluster. When public network access is enabled, the cluster is automatically assigned a public IP address with dedicated dynamic BGP bandwidth. When Public Network Access is selected, further configure the following parameters: • Bandwidth: Set the bandwidth for public network access. The value ranges from 1 Mbit/s to 200 Mbit/s. • Configure Whitelist Control public network access to the cluster using a whitelist. – If a whitelist is configured, only IP addresses that are on this whitelist can access the cluster over the public network. Click +Add. In the displayed text box, enter IP addresses or CIDR blocks that are allowed to access the cluster from the public network. Separate them using commas (,). Each value must be unique. An example of valid values: 192.168.1.1,10.0.0.0/24. Examples of invalid values: 0.0.0.0, xx.xx.xx.x0, 172.16.0.0-172.16.255.255, non-standard formats (e.g., 192.168.1), and duplicate values. – If no whitelist is configured, all public IP addresses can access the cluster. However, this can be a security risk and should be avoided.

Cluster Management

Configure the cluster name, enterprise project, and more.

Table 2-22 Cluster management information

Parameter	Description
Cluster Name	User-defined cluster name. The cluster name must start with a letter and can contain 4 to 32 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.

Parameter	Description
Add Description	Add a description for the cluster for easy recognition. Click + Add Description and then enter a description of the cluster. The value can contain 0 to 128 characters.
Enterprise Project	Associate the cluster with an enterprise project. An enterprise project groups cloud resources, so you can manage resources and members by project. The default project is default . If enterprise projects are enabled, you can select an enterprise project from the drop-down list. You can click View Enterprise Projects to go to the Enterprise Project Management Service console and check enterprise project details. For more information, see Accessing the Enterprise Center .
Tags	Adding tags to clusters helps you identify and manage your cluster resources. Click + Add . You can select existing tag keys and values from the drop-down lists or enter new ones. If you want to use the same tag to identify multiple cloud resources for better resource grouping, we recommend that you predefine tags in Tag Management Service (TMS). For details, see Predefined Tags . Each cluster can have a maximum of 20 tags. If your organization has configured tag policies for CSS, add tags to clusters based on these policies. If a tag does not comply with the tag policies, cluster creation may fail. Contact the administrator to learn more about tag policies.

Automatic Snapshot Creation

Click  to expand **More Settings** and configure automatic snapshot creation. Automatic snapshot creation is disabled by default. To enable it, select **Automatic Snapshot Creation**, and configure the necessary settings. Automatic snapshot creation provides continuous data protection by periodically backing up cluster data.

CAUTION

If you enable automatic snapshot creation, automatically created snapshots will be stored in standard OBS buckets, and you will need to pay extra for the storage. For more information, see [Billing Overview](#).

Table 2-23 Configuring automatic snapshot creation

Parameter	Description
OBS Bucket	From the drop-down list, select an OBS bucket for storing snapshots. If no OBS buckets meet your requirements, click Create Bucket to go to the OBS console and create one. For details, see Creating a Bucket. To grant an IAM user access to an OBS bucket, you need to grant the GetBucketStoragePolicy, GetBucketLocation, ListBucket, and ListAllMyBuckets permissions to that user.
Backup Path	Snapshot storage path in the OBS bucket. The backup path cannot: • Contain the following characters: \:*?"<> '{}. • Start with a slash (/). • Start or end with a period (.). • Contain more than two consecutive slashes (/) or periods (.). • Exceed 512 characters.
Maximum Backup Rate (per Second)	The parameter sets the maximum backup rate per node. When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual backup rate may not reach the configured value, as it depends on factors such as OBS performance and disk I/O. Value format: number + unit • Number range: 0-9999 • Unit: KB, MB, GB, TB, PB, or B Default value: 40 MB The value 0MB means there is no limit on how fast data is backed up to snapshots. An overly high backup rate may lead to excessive resource usage, which may impact cluster stability. Therefore, set this parameter carefully.

Parameter	Description
Maximum Recovery Rate (per Second)	The parameter sets the maximum recovery rate per node. When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual recovery rate may not reach the configured value, as it depends on factors such as OBS performance and disk I/O. Value format: number + unit • Number range: 0-9999 • Unit: KB, MB, GB, TB, PB, or B Default value: • For Elasticsearch 7.6.2 clusters or earlier, the default value is 40 MB. • For Elasticsearch clusters later than 7.6.2, the default value is 0 MB, indicating no limit. An overly high recovery rate may lead to excessive resource usage, which may impact cluster stability. Therefore, set this parameter carefully. For Elasticsearch clusters later than 7.6.2, the recovery rate is also limited by the indices.recovery.max_bytes_per_sec parameter. • If Maximum Recovery Rate (per Second) is less than indices.recovery.max_bytes_per_sec, the former takes effect. • If Maximum Recovery Rate (per Second) is greater than indices.recovery.max_bytes_per_sec, the latter takes effect. NOTE • To check the value of indices.recovery.max_bytes_per_sec, run the following command: GET _cluster/settings • To modify indices.recovery.max_bytes_per_sec, run the following command: PUT _cluster/settings { "transient": { "indices.recovery.max_bytes_per_sec": "100mb" } }

Parameter	Description
IAM Agency	Select an IAM agency to grant the current account the permission to access and use OBS. To store snapshots to an OBS bucket, you must have the required OBS access permissions. - If you are configuring an agency for the first time, click Automatically Create IAM Agency to create css-obs-agency. - If there is an IAM agency automatically created earlier, you can click One-click authorization to have the OBS Administrator permissions deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control. <pre> "obs:bucket:GetBucketLocation", "obs:object:GetObjectVersion", "obs:object:GetObject", "obs:object:DeleteObject", "obs:bucket:HeadBucket", "obs:bucket:GetBucketStoragePolicy", "obs:object:DeleteObjectVersion", "obs:bucket:ListBucketVersions", "obs:bucket:ListBucket", "obs:object:PutObject" </pre> - When OBS buckets use SSE-KMS encryption, the IAM agency must be granted KMS permissions. You can click Automatically Create IAM Agency and One-click authorization to have the following custom policies created automatically. <pre> "kms:cmk:create", "kms:dek:create", "kms:cmk:get", "kms:dek:decrypt", "kms:cmk:list" </pre> - To use Automatically Create IAM Agency and One-click authorization, the following minimum permissions are required: <pre> "iam:agencies:listAgencies", "iam:roles:listRoles", "iam:agencies:getAgency", "iam:agencies:createAgency", "iam:permissions:listRolesForAgency", "iam:permissions:grantRoleToAgency", "iam:permissions:listRolesForAgencyOnProject", "iam:permissions:revokeRoleFromAgency", "iam:roles:createRole" </pre> - To use an IAM agency, the following minimum permissions are required: <pre> "iam:agencies:listAgencies", "iam:agencies:getAgency", "iam:permissions:listRolesForAgencyOnProject", "iam:permissions:listRolesForAgency" </pre> WARNING The agency name can contain only letters (case-sensitive), digits, underscores (_), and hyphens (-). Otherwise, the backup will fail.

Parameter	Description
Snapshot Name Prefix	A snapshot name consists of the snapshot name prefix and timestamp, for example, <i>snapshot-1566921603720</i> . The snapshot name prefix can contain 1 to 32 characters. It must start with a lowercase letter, and it can only contain lowercase letters, digits, hyphens (-), and underscores (_). The default value is snapshot .
Time Zone	Select a time zone for the backup start time.
Backup Start Time	Specify the start time of auto backup. Select a value from the drop-down list. The interval can be Daily , Hourly , or weekly (by selecting a specific day of the week), and the backup time can be set to any hour from 00:00 to 23:00 (full hours only).
Retained Snapshots	Number of automatically created snapshots to be retained. The value ranges from 1 to 90. The system automatically deletes excess snapshots every half hour. (The expiration deletion policy applies only to the snapshots that were automatically created at the same frequency as the current automatic snapshot creation policy.) NOTE If the snapshot creation interval is short or if the data size of indexes is large, the number of automatic snapshots retained may not reach the value set using this parameter.

VPC Endpoint

Expand **More Settings** and configure a VPC endpoint for the cluster. VPC endpoint is disabled by default. To enable it, select **VPC Endpoint**, and configure the necessary settings.

VPC Endpoint enables you to access resources across VPCs using a dedicated gateway without exposing the network information of servers. The address of a VPC endpoint consists of an IPv4 address and internal domain name.

CAUTION

- If a shared VPC and a subnet within this shared VPC were selected earlier for the cluster, **VPC Endpoint** cannot be enabled.
 - VPC endpoints, if created for the cluster, will incur extra fees, depending on the resource usage. For details, see [VPC Endpoint Billing](#).
 - You need specific permissions to create VPC endpoints. For details, see [VPCEP Permissions](#).
-

Table 2-24 Configuring VPC Endpoint

Parameter	Description
Create Private Domain Name	Whether to create a private domain name for the VPC endpoint. • Enable: The system automatically assigns a private domain name to the VPC endpoint. After cluster creation, you can check this private domain name on the VPC Endpoint tab of the cluster details page. • Disable: No private domain name will be configured for the VPC endpoint. The cluster can only be accessed through an IP address assigned to the VPC endpoint.
Add Account	Add accounts that are allowed to access the cluster through the VPC endpoint service. To add an account, click Add Account and enter an account ID. • To obtain your authorized account ID, point to your username in the upper right corner, and choose My Credentials. Copy the value of Account ID. • If no account is added or the account ID is set to *, all users are allowed to access the cluster through the VPC endpoint.

Kibana Public Network Access

Expand **More Settings** and configure Kibana public network access for the cluster. Kibana public network access is disabled by default. To enable it, select **Kibana Public Network Access**, and configure the necessary settings.

When Kibana public network access is enabled, a public IP address will be automatically assigned for Kibana access. Users will be able to access this Elasticsearch cluster's Kibana console from the public network through this address.

CAUTION

This parameter is available only when the security mode is enabled for the cluster.

Table 2-25 Configuring public network access for Kibana

Parameter	Description
Bandwidth	Bandwidth for accessing Kibana from the public network Value range: 1 Mbit/s to 200 Mbit/s

Parameter	Description
Configure Whitelist	Control Kibana public network access using a whitelist. • If a whitelist is configured, only IP addresses that are on this whitelist can access the cluster's Kibana console over the public network. Click + Add. In the displayed text box, enter IP addresses or CIDR blocks that are allowed to access the cluster's Kibana console from the public network. Separate them using commas (,). Each value must be unique. An example of valid values: 192.168.1.1,10.0.0.0/24. Examples of invalid values: 0.0.0.0, xx.xx.xx.x0, 172.16.0.0-172.16.255.255, non-standard formats (e.g., 192.168.1), and duplicate values. • If no whitelist is configured, all public IP addresses can access the cluster's Kibana console. However, this can be a security risk and should be avoided.

Quantity

Configure the cluster's subscription duration and auto-renewal option. This parameter is available only when **Billing Mode** is set to **Yearly/Monthly**.

Table 2-26 Configuring the subscription duration

Parameter	Description
Required Duration	Set the duration of your subscription in the range of 1 month to 3 years. If you plan to use the cluster for more than nine months, you should choose a yearly subscription for a better price.
Auto-renew	Whether to automatically renew the subscription. By default, auto-renewal is disabled. To enable it, select Auto-renew. When Auto-renew is selected, a yearly/monthly subscription is automatically renewed upon expiry. • Monthly subscription: auto-renews for 1 month each time. • Yearly subscription: auto-renews for 1 year each time. For more information, see Auto-Renewal Rules.

Submit

1. In the **Configuration Summary** panel on the right, review the cluster configuration details.

Any mandatory fields that are missing are displayed in red, and you have to set them in the configuration area.

2. Click **Create Now**.

Cluster creation time depends on the number of nodes. Typically, this process takes less than 60 minutes, though clusters with a large number of nodes may require additional time.

3. Return to the cluster list and check the newly created cluster.

If the cluster is created successfully, **Cluster Status** changes to **Available**.

Related Documents

- For details about how to check basic cluster information, such as the cluster status, private IP address, and node information, see [Viewing Elasticsearch Cluster Information](#).
- For details about the methods you can use to access an Elasticsearch cluster, see [Elasticsearch Cluster Access Methods](#).
- You can enhance cluster performance by properly planning the number of index shards. For details, see [How Do I Plan the Quantity of Index Shards for a Cluster?](#)
- For how to optimize the write and query performance of a cluster and improve cluster utilization, see [Optimizing the Write Performance of a Cluster](#) and [Optimizing the Query Performance of a Cluster](#).
- For details about how to migrate data from Elasticsearch clusters elsewhere to Elasticsearch clusters in CSS, see [About Elasticsearch Cluster Migration Solutions](#).
- For details about how to synchronize data from other products to CSS's Elasticsearch clusters, see [Using Logstash to Synchronize Data to Elasticsearch](#).
- For a security-mode cluster, you need to specify the administrator account and password for cluster access. You are advised to create roles and users with the corresponding permissions on the Kibana console. For details, see [Creating Users for an Elasticsearch Cluster and Granting Cluster Access](#).
- If cluster creation fails, follow the guidelines in [Why Does Cluster Creation Fail?](#) for troubleshooting.

2.4 Creating an Elasticsearch Cluster (Old Version)

This topic describes how to create an Elasticsearch cluster. You can create an Elasticsearch cluster in either of the following ways:

- Method 1: Create an Elasticsearch cluster on the CSS management console. This topic provides the operation guide for the old UI. For the operation guide for the new UI, see [Creating an Elasticsearch Cluster \(New Version\)](#).

NOTE

To improve efficiency in creating clusters, CSS has made several enhancements on the cluster creation page in terms of ease-of-use. Being more intuitive, the new UI streamlines cluster creation procedures. You are advised to use the new UI to create Elasticsearch clusters.

- Method 2: Create an Elasticsearch cluster using CSS APIs. For details, see [Creating a Cluster](#).

Scenarios

[Table 2-27](#) lists key parameters that differentiate between different types of clusters.

Table 2-27 Parameters that differentiate between different types of clusters

Cluster Type	Security Mode	HTTPS Access	Internet Access	Kibana Public Network Access
Cluster in non-security mode	Disabled	N/A	Cannot be enabled	Cannot be enabled
Cluster in security mode + HTTP	Enabled	Disabled	Cannot be enabled	Can be enabled
Cluster in security mode + HTTPS	Enabled	Enabled	Can be enabled	Can be enabled


Prerequisites

You have planned the Elasticsearch cluster configuration by following the instructions in [Elasticsearch Cluster Planning Suggestions](#).

Creating a Cluster

1. Log in to the [CSS management console](#).
2. On the **Dashboard** page, click **Create Cluster** in the upper right corner. The **Create Cluster** page is displayed.
Alternatively, choose **Clusters** > **Elasticsearch** in the navigation tree on the left. Click **Create Cluster** in the upper right corner. The **Create Cluster** page is displayed.
3. Click **Back to Old Version** on the right to go to the old UI for cluster creation.

Figure 2-2 Create Cluster (old version)



4. On the **Basic Configuration** page, configure basic information and resources for the Elasticsearch cluster.

Table 2-28 Basic configuration of the Elasticsearch cluster

Parameter	Description
Billing Mode	Select Yearly/Monthly or Pay-per-use . • Yearly/monthly: You pay for the cluster by year or month, in advance. The service duration ranges from one month to three years. If you plan to use the cluster for more than nine months, you should choose a yearly subscription for a better price. • Pay-per-use: You will be billed hourly by actual duration of use. Any partial hour of usage will be rounded up to one hour.
Required Duration	The duration for which the purchased EIP will be used. The duration must be specified if the Billing Mode is set to Yearly/Monthly . Configure automatic renewal if necessary.
Region	Select the region where the cluster is located. ECSs in different regions cannot communicate with each other over an intranet. For lower network latency and quicker resource access, select the nearest region.
AZ	Select one or more AZs associated with the cluster region. A maximum of three AZs can be configured. For details about the use of multiple AZs, see Suggestions on Multi-AZ Deployment .
Type	Select Elasticsearch .
Version	Select a cluster version from the drop-down list.
Name	User-defined cluster name. The cluster name must start with a letter and can contain 4 to 32 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.
Cluster Description	Add a description for the cluster for easy recognition. The value can contain 0 to 128 characters.

Parameter	Description
Nodes	Number of nodes in the cluster. Select a number from 1 to 32. You are advised to configure three or more nodes to ensure high availability of the cluster. • If the cluster has no client nodes, the data nodes will need to handle the additional task of query request distribution and data analysis. • If the cluster has no master nodes, the data nodes will also need to handle cluster management. NOTE If the number of data nodes in a cluster is not evenly divisible by the number of AZs, data in the cluster may be unevenly distributed, affecting data query or write performance.
CPU Architecture	x86 and Kunpeng are supported. The supported types depend on the actual regional environment.
Node Specifications	Data node flavor. You can select a specified specification based on your needs. Each cluster supports only one specification. For details, see ECS Types .
Node Storage Type	If you select EVS for node storage, you need to further select the EVS disk type for data nodes of the cluster. Options include Common I/O, High I/O, Ultra-high I/O, and Extreme SSD. NOTE If the type of storage in use is not supported, the storage type is not displayed.
Node Storage Capacity	Data node storage capacity. Its value range varies with node specifications. The node storage capacity must be a multiple of 20. The node storage capacity cannot be reduced once the cluster is created. Choose an appropriate capacity based on service needs.
Master node	Master nodes manage cluster-wide operations, including metadata, indexes, and shard allocation. For large-scale deployments, using dedicated master nodes enhances cluster stability, service availability, and centralized control. After enabling master nodes, specify Node Specifications, Nodes, and Node Storage Type. The value of Nodes must be an odd number greater than or equal to 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.

Parameter	Description
Client node	Client nodes route and coordinate search and index requests, offloading processing from data nodes for enhanced query performance and cluster scalability when there are heavy loads. After enabling client nodes, 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.
Cold data node	Cold data nodes are used to store and query large quantities of latency-insensitive data. They offer an effective way to manage large datasets while cutting storage costs. After enabling cold data nodes, 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 required. When cold data nodes are enabled, you can switch between cold and hot data nodes. For details, see Switching Between Hot and Cold Storage for an Elasticsearch Cluster. NOTE If the number of cold data nodes in a cluster is not evenly divisible by the number of AZs, data in the cluster may be unevenly distributed, affecting data query or write performance.
Enterprise Project	When creating a CSS cluster, you can bind an enterprise project to the cluster if you have enabled the enterprise project function. Select an enterprise project from the Enterprise Project drop-down list, or click View Enterprise Project to go to the Enterprise Project Management Service page and check existing enterprise projects.

- Click **Next: Network**.
- On the **Network** page, configure the network settings and security mode for the Elasticsearch cluster.

Table 2-29 Network settings for the Elasticsearch cluster

Parameter	Description
VPC	Specify a VPC to isolate the cluster's network. Click View VPC to go to the VPC management console and check the created VPCs or VPCs shared with the current account. You can select a shared VPC when creating a cluster, in which case, the default vpc subnet statement permission is required. VPC sharing allows you to centrally manage resources across multiple accounts, helping to improve resource management efficiency and reduce O&M costs. For more information about shared VPCs, see VPC Sharing. If no VPC is available, contact the CSS administrator to create a new VPC. For details, see Creating a VPC and Subnet. NOTE The VPC must contain CIDRs. Otherwise, cluster creation will fail. By default, a created VPC contains CIDRs.
Subnet	A subnet provides dedicated network resources that are isolated from other networks, improving network security. Select a subnet needed by the cluster in the current VPC. Or you can select a subnet under a shared VPC. The owner of a VPC can share the VPC to specified users.
Security Group	A security group serves as a virtual firewall that provides access control policies for clusters. Select a security group for the cluster. Click View Security Group to go to the security group list, where you can view details about security groups. NOTE • Ensure that Port Range/ICMP Type is set to a port range that includes port 9200 for the selected security group. • 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. • After the port 9300 is enabled, if the cluster disk usage is high, delete expired data to release the disk storage space. For details, see How Do I Clear Expired Data to Release Storage Space?

Parameter	Description
Security Mode	Whether to enable the security mode for the cluster. • The security mode is enabled by default. In security mode, a cluster's communication is encrypted and access to the cluster requires user authentication. This is why Administrator Username and Administrator Password must be configured for the cluster. – The default administrator username is admin. – Set and confirm the Administrator Password. This password will be required when you access this cluster. • If Security Mode is disabled, a cluster in non-security mode will be created. With such a cluster, access to the cluster will require no user authentication, and data will be transmitted in plaintext using HTTP. Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.
HTTPS Access	HTTPS access can be enabled only when security mode is enabled for the cluster. With HTTPS access enabled, communication will be encrypted when you access the cluster. NOTE Compared with a non-security mode cluster that uses HTTP, a security-mode cluster that uses HTTPS has lower read performance. The performance loss is estimated at around 20% under high concurrency. If you want fast read performance as well as the isolation and permission control (such as indexes, documents, and fields) enabled by the security mode, you can disable HTTPS Access. 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.
Public IP Address	This parameter is available only when Security Mode and HTTPS Access are enabled. When Public IP Address is enabled, a public IP address is automatically assigned, which will enable access to the security cluster from the Internet. For details, see Configuring Public Network Access for an Elasticsearch Cluster.

7. Click **Next: Advanced Settings**.
8. On the **Advanced Settings** page, configure a snapshot policy and other advanced settings for the Elasticsearch cluster.
 - a. Set a cluster snapshot policy.

Cluster snapshots are disabled by default. To enable them, toggle on **Cluster Snapshots**. To store snapshots automatically created in OBS, an

agency will need to be created in order to access OBS. Fees will be incurred for using standard OBS storage.

Table 2-30 Basic configuration for a cluster snapshot policy

Parameter	Description
OBS Bucket	From the drop-down list, select an OBS bucket for storing snapshots. You can also click Create Bucket on the right to create an OBS bucket. For details, see Creating a Bucket. The created or existing OBS bucket must meet the following requirements: • Storage Class: Standard. • Region must be the same as that of the created cluster.
Backup Path	Storage path of the snapshot in the OBS bucket. The backup path cannot: • Contain the following characters: \:*?"<> '{}. • Start with a slash (/). • Start or end with a period (.). • Contain more than two consecutive slashes (/) or periods (.). • Exceed 512 characters.
Maximum Backup Rate (per Second)	The parameter sets the maximum backup rate per node. When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual backup rate may not reach the configured value, as it depends on factors such as OBS performance and disk I/O. Value format: number + unit • Number range: 0–9999 • Unit: KB, MB, GB, TB, PB, or B Default value: 40 MB The value 0MB means there is no limit on how fast data is backed up to snapshots. An overly high backup rate may lead to excessive resource usage, which may impact cluster stability. Therefore, set this parameter carefully.

Parameter	Description
Maximum Recovery Rate (per Second)	The parameter sets the maximum recovery rate per node. When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual recovery rate may not reach the configured value, as it depends on factors such as OBS performance and disk I/O. Value format: number + unit - Number range: 0–9999 - Unit: KB, MB, GB, TB, PB, or B Default value: - For Elasticsearch 7.6.2 clusters or earlier, the default value is 40 MB. - For Elasticsearch clusters later than 7.6.2, the default value is 0 MB, indicating no limit. An overly high recovery rate may lead to excessive resource usage, which may impact cluster stability. Therefore, set this parameter carefully. For Elasticsearch clusters later than 7.6.2, the recovery rate is also limited by the indices.recovery.max_bytes_per_sec parameter. - If Maximum Recovery Rate (per Second) is less than indices.recovery.max_bytes_per_sec, the former takes effect. - If Maximum Recovery Rate (per Second) is greater than indices.recovery.max_bytes_per_sec, the latter takes effect. NOTE - To check the value of indices.recovery.max_bytes_per_sec, run the following command: GET _cluster/settings - To modify indices.recovery.max_bytes_per_sec, run the following command: PUT _cluster/settings { "transient": { "indices.recovery.max_bytes_per_sec": "100mb" } }

Parameter	Description
IAM Agency	To store snapshots to an OBS bucket, you must have the required OBS access permissions. Select an IAM agency to grant the current account the permission to access and use OBS. - If you are configuring an agency for the first time, click Automatically Create IAM Agency to create css-obs-agency. - If there is an IAM agency automatically created earlier, you can click One-click authorization to have the OBS Administrator permissions deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control. "obs:bucket:GetBucketLocation", "obs:object:GetObjectVersion", "obs:object:GetObject", "obs:object:DeleteObject", "obs:bucket:HeadBucket", "obs:bucket:GetBucketStoragePolicy", "obs:object:DeleteObjectVersion", "obs:bucket:ListBucketVersions", "obs:bucket:ListBucket", "obs:object:PutObject" - When OBS buckets use SSE-KMS encryption, the IAM agency must be granted KMS permissions. You can click Automatically Create IAM Agency and One-click authorization to have the following custom policies created automatically. "kms:cmk:create", "kms:dek:create", "kms:cmk:get", "kms:dek:decrypt", "kms:cmk:list" - To use Automatically Create IAM Agency and One-click authorization, the following minimum permissions are required: "iam:agencies:listAgencies", "iam:roles:listRoles", "iam:agencies:getAgency", "iam:agencies:createAgency", "iam:permissions:listRolesForAgency", "iam:permissions:grantRoleToAgency", "iam:permissions:listRolesForAgencyOnProject", "iam:permissions:revokeRoleFromAgency", "iam:roles:createRole" - To use an IAM agency, the following minimum permissions are required: "iam:agencies:listAgencies", "iam:agencies:getAgency", "iam:permissions:listRolesForAgencyOnProject", "iam:permissions:listRolesForAgency" WARNING The agency name can contain only letters (case-sensitive), digits, underscores (_), and hyphens (-). Otherwise, the backup will fail.

Table 2-31 Setting Automatic Snapshot Creation

Parameter	Description
Snapshot Name Prefix	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 .
Time Zone	Time zone for the backup time. Specify Backup Started Time based on the time zone.
Backup Start Time	Specify the start time of auto backup. Select a time from the drop-down list. The interval can be Daily , Hourly , or weekly (by selecting a specific day of the week), and the backup time can be set to any hour from 00:00 to 23:00 (full hours only).
Retained Snapshots	Number of automatically created snapshots to be retained. The value ranges from 1 to 90. The system automatically deletes excess snapshots every half hour. (The expiration deletion policy applies only to the snapshots that were automatically created at the same frequency as the current automatic snapshot creation policy.) NOTE If the snapshot creation interval is short or if the data size of indexes is large, the number of automatic snapshots retained may not reach the value set using this parameter.

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

VPC Endpoint

VPC Endpoint enables you to access resources across Virtual Private Clouds (VPCs) using a dedicated gateway, without exposing the network information of servers. When VPC Endpoint is enabled, a VPC endpoint will be created by default. You can select **Create Private Domain Name** if necessary. Users will be able to access this cluster across VPCs through node IP addresses or a private domain name.

NOTE

- If a shared VPC and a subnet within this shared VPC were selected earlier for the cluster on the **Network** page, **VPC Endpoint** cannot be enabled for the cluster.
- After the VPC endpoint service is enabled for a cluster, you will be billed per use for the service. For more information, see [Billing Modes](#).

Table 2-32 Configuring VPC Endpoint

Parameter	Description
Create Private Domain Name	If Create Private Domain Name is selected, the system generates a node IP address and also automatically creates a private domain name, which enables users to access this cluster from within the same VPC. If it is not selected, only a node IP address is generated.
VPC Endpoint Whitelist	In VPC Endpoint Whitelist, you can add accounts that are allowed to access the cluster using a node IP address or private domain name. - Click Add to add accounts in Account ID. If no account is added or the account ID is set to *, all users are allowed to access the cluster through the VPC endpoint service. - Click Delete in the Operation column to delete accounts. NOTE To obtain your authorized account ID, point to your username in the upper right corner, and choose My Credentials. Copy the value of Account ID.

Kibana Public Network Access

This parameter is available only when security mode is enabled for the cluster. By enabling this option, you can obtain a public IP address for accessing Kibana.

Table 2-33 Configuring public network access for Kibana

Parameter	Description
Bandwidth	Bandwidth for accessing Kibana through a 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 ranges that are on the whitelist can access Kibana through from the public network.

Parameter	Description
Whitelist	IP addresses or IP address ranges that are allowed to access the cluster from the public network. Use commas (,) to separate different values. This parameter can be configured only when Access Control is enabled. You are advised to enable the whitelist. NOTE The whitelist that controls Kibana public network access depends on whitelist support by the ELB service. After you update 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 in approximately 1 minute after these connections are disconnected.

Tags

Adding tags to clusters helps you identify and manage your cluster resources. You can configure your own tags or use tags predefined by the Tag Management Service (TMS). If you want to use the same tag to identify multiple cloud resources for better resource grouping, we recommend that you predefine tags in Tag Management Service (TMS). For details, see [Predefined Tags Overview](#).

NOTE

If your organization has configured tag policies for CSS, add cluster tags based on these policies. Tags that do not comply with predefined tag policies will cause a cluster creation failure. Contact the administrator to learn about the tag policies.

- Click **Next: Confirm Configuration**. Check the configuration and click **Next** to create a cluster.
- Return to the cluster list and check the newly created cluster. If the cluster is created successfully, **Cluster Status** changes to **Available**. Cluster creation time depends on the number of nodes. Typically, this process takes less than 60 minutes, though clusters with a large number of nodes may require additional time.

If cluster creation fails, try creating the cluster again by rectifying the errors returned.

Follow-up Operations

After an Elasticsearch cluster is created, you are advised to optimize the query performance of the cluster to improve efficiency by referring to [Cluster Performance Tuning](#).

2.5 Accessing Elasticsearch Clusters

2.5.1 Elasticsearch Cluster Access Methods

There are many ways to connect to an Elasticsearch cluster. Select the most appropriate access method based on the programming language you prefer. The various access methods support different types of network configurations. The cluster access addresses vary depending on the network configuration you choose. Select an appropriate network configuration method based on your network environment.

Access Methods

Table 2-34 Cluster access methods

Access Method	Scenario	Details
Kibana (recommended)	• A graphical user interface • Instance monitoring • Data management Language-irrelevant	Logging In to an Elasticsearch Cluster Using Kibana
Cerebro		Logging In to an Elasticsearch Cluster Through Cerebro
Open-source Elasticsearch API	Run Curl commands to access and manage Elasticsearch clusters.	Accessing an Elasticsearch Cluster Using cURL Commands
Java	Use Java to access and manage Elasticsearch clusters. To avoid compatibility issues, use the Java client that matches your Elasticsearch cluster version.	Accessing an Elasticsearch Cluster Through the High Level REST Client Accessing an Elasticsearch Cluster Through the Low Level REST Client Accessing an Elasticsearch Cluster Through the Transport Client Accessing an Elasticsearch Cluster Using Spring Boot
Python	Use Python to access and manage Elasticsearch clusters.	Accessing an Elasticsearch Cluster Using Python
MRS Hive	Use the Elasticsearch-Hadoop (ES-Hadoop) connector of MRS to access and manage Elasticsearch clusters using Hive.	Accessing an Elasticsearch Cluster Through the MRS Hive Client

Access Method	Scenario	Details
Go	Use Go to access and manage Elasticsearch clusters.	Accessing an Elasticsearch Cluster Using Go
LDAP Authentication	Enable Lightweight Directory Access Protocol (LDAP) authentication for an Elasticsearch cluster.	Accessing an Elasticsearch Cluster Using LDAP

Network Configuration

Configure the network of the Elasticsearch cluster, and ensure that you have a server that can connect to the cluster's network. Obtain the cluster access address based on the selected network. Use this address to access the cluster.

Table 2-35 Network configuration for the cluster

Network Environment	Access Method	Scenario	Details
Private network	Private IP address	Obtain the private IP address of each node in the cluster and access the cluster by connecting to one of these IP addresses. How to obtain: 1. Log in to the CSS management console. 2. In the navigation pane on the left, choose Clusters > Elasticsearch. 3. In the cluster list, obtain the target cluster's private IP address from the Private IP Address column. Generally, the IP address format is <code><host>:<port></code> or <code><host>:<port>,<host>:<port></code>. If the cluster has only one node, the IP address and port number of this 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.	This is the default method.

Network Environment	Access Method	Scenario	Details
Private network	IP address or private domain name of a VPC endpoint	Configure a VPC endpoint for the cluster, enabling cross-VPC access over the internal network. The cluster can be accessed via the VPC endpoint IP address or a private domain name. How to obtain: 1. Log in to the CSS management console. 2. In the navigation pane on the left, choose Clusters > Elasticsearch. 3. In the cluster list, click the name of the target cluster. The cluster information page is displayed. 4. Choose Cluster Access > VPC Endpoint. In the VPC endpoint list below, record Service Address or Private Domain Name of the target VPC endpoint.	Configuring VPC Endpoints for an Elasticsearch Cluster
Private network	Private IP address of a dedicated load balancer	Associate a dedicated load balancer with the cluster to enable access through the private IP address of this dedicated load balancer. How to obtain: 1. Log in to the CSS management console. 2. In the navigation pane on the left, choose Clusters > Elasticsearch. 3. In the cluster list, click the name of the target cluster. The cluster information page is displayed. 4. Choose Cluster Access > Load Balancing. On the Elasticsearch tab, find and record the private or public IP address of the target load balancer.	Configuring a Dedicated Load Balancer for an Elasticsearch Cluster

Network Environment	Access Method	Scenario	Details
Public network	Public IP address	Enable public network access for a security-mode cluster that uses HTTPS, and configure a whitelist for access control. The cluster can be accessed from the public network via its public IP address. How to obtain: 1. Log in to the CSS management console. 2. In the navigation pane on the left, choose Clusters > Elasticsearch. 3. In the cluster list, click the name of the target cluster. The cluster information page is displayed. 4. Click the Overview tab. In the Configuration area, find and record the IP address used for public network access.	Configuring Public Network Access for an Elasticsearch Cluster
Public network	Public IP address of a dedicated load balancer	Associate with the cluster a dedicated load balancer configured with an EIP to enable access through the public IP address of this dedicated load balancer. How to obtain: 1. Log in to the CSS management console. 2. In the navigation pane on the left, choose Clusters > Elasticsearch. 3. In the cluster list, click the name of the target cluster. The cluster information page is displayed. 4. Choose Cluster Access > Load Balancing. On the Elasticsearch tab, find and record the private or public IP address of the target load balancer.	Configuring a Dedicated Load Balancer for an Elasticsearch Cluster

2.5.2 Logging In to an Elasticsearch Cluster Using Kibana

Kibana is the official data visualization and exploration platform for Elasticsearch. It enables in-depth data analysis and interactive visualization. In CSS, Kibana is

pre-built for each Elasticsearch cluster. You can start Kibana in one click, without installing anything. Kibana offers a comprehensive suite of dashboard features and visualization tools. It seamlessly integrates Elasticsearch's analytical capabilities, supporting the full analytics process from data exploration to actionable business insights.

CSS's Kibana supports multiple access methods. Steps needed to log in to an Elasticsearch cluster vary depending on the access method you choose. See [Table 2-36](#).

Table 2-36 Methods for logging in to an Elasticsearch cluster through Kibana

Kibana Access Method	Constraints	Details
One-click access to Kibana from the service console	N/A	Logging In to an Elasticsearch Cluster by Accessing Kibana Through the Console
Accessing Kibana using a public IP address	- Only clusters in security mode support Kibana access through a public IP address. - Kibana public network access cannot be enabled for Elasticsearch security-mode clusters created before June 2020, that is, when the feature first became available.	Logging In to an Elasticsearch Cluster by Accessing Kibana Using a Public IP Address
Accessing Kibana using a private network address	The client that accesses Kibana must be able to reach Kibana's private IP address.	Logging In to an Elasticsearch Cluster by Accessing Kibana Using a Private Network Address

Constraints on Kibana Usage

- You can customize the username, role name, and tenant name in Kibana.
- With an Elasticsearch 7.10.2 cluster (the image version is no earlier than 7.10.2_24.3.3_x.x.x), you can switch between Chinese and English on Kibana.

In the upper right corner of the Kibana console, click , select a language, and confirm the language change.

Logging In to an Elasticsearch Cluster by Accessing Kibana Through the Console

- Log in to the [CSS management console](#).
- In the navigation pane on the left, choose **Clusters > Elasticsearch**.

3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
 - Non-security cluster: The Kibana console is displayed without asking for a username and password.
 - Security cluster: Enter the username and password on the login page and click **Log In** to log in to the Kibana console. The default username is **admin** and the password is the one specified during cluster creation.
4. After the login is successful, you can access the cluster through Kibana.

Logging In to an Elasticsearch Cluster by Accessing Kibana Using a Public IP Address

WARNING

- Only clusters in security mode support Kibana access through a public IP address.
 - Kibana public network access cannot be enabled for Elasticsearch security-mode clusters created before June 2020, that is, when the feature first became available.
 - The whitelist that controls Kibana public network access depends on whitelist support by the ELB service. After you update 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 in approximately 1 minute after these connections are disconnected.
 - If you disable Kibana public network access and then re-enable it, the public IP address for accessing Kibana may change. Exercise caution.
-

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Cluster Access > Kibana Public Network Access** to check whether Kibana public network access is enabled for the cluster.
 - Yes: Go to [6](#).
 - No: Go to the next step.
5. Enable Kibana public network access for the Elasticsearch cluster.
 - a. On the **Kibana Public Network Access** page, toggle on the **Kibana Public Network Access** button.
 - b. In the displayed **Kibana Public Network Access** dialog box, set the parameters.


Table 2-37 Configuring public network access for Kibana

Parameter	Description
Bandwidth	Bandwidth for accessing Kibana from the public network Value range: 1 Mbit/s to 200 Mbit/s
Configure Whitelist	Control Kibana public network access using a whitelist. - If a whitelist is configured, only IP addresses that are on this whitelist can access the cluster's Kibana console over the public network. Click + Add. In the displayed text box, enter IP addresses or CIDR blocks that are allowed to access the cluster's Kibana console from the public network. Separate them using commas (,). Each value must be unique. An example of valid values: 192.168.1.1,10.0.0.0/24. Examples of invalid values: 0.0.0.0, xx.xx.xx.x0, 172.16.0.0-172.16.255.255, non-standard formats (e.g., 192.168.1), and duplicate values. - If no whitelist is configured, all public IP addresses can access the cluster's Kibana console. However, this can be a security risk and should be avoided.

- c. Click **OK** to confirm the settings.

If Kibana public network access is already enabled, you can modify relevant settings.

6. After Kibana public network access is enabled, obtain the Kibana public IP address on the **Kibana Public Network Access** page.

Figure 2-3 Obtaining the Kibana public IP address

7. Enter the public IP address for Kibana in the browser address box to go to the Kibana login page.
Enter the username and password on the login page and click **Log In** to log in to the Kibana console. The default username is **admin** and the password is the one specified during cluster creation.
8. After the login is successful, you can access the Elasticsearch cluster through Kibana.

Logging In to an Elasticsearch Cluster by Accessing Kibana Using a Private Network Address

CAUTION

The client that accesses Kibana must be able to reach Kibana's private IP address.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Overview** tab. In the **Configuration** area, record **Private IPv4 Address**.

Figure 2-4 Obtaining the cluster's private IP address

Enterprise Project	default
HTTPS Access	Disabled
Private IPv4 Address	10  00

5. Obtain Kibana's private IP address.
By changing the port number in the cluster's private IP address from **9200** to **5601**, you get Kibana's private IP address. For example, if the cluster's private IPv4 address is **192.168.0.xx:9200**, the private IP address of Kibana is **192.168.0.xx:5601**.
6. (Optional) Configure a return route. To enable a client to access the cluster's Kibana console across different networks, configure a route for the Elasticsearch cluster.
 - a. Connect the client and Kibana through a Direct Connect or VPC peering connection.
 - b. Configure the route connecting the Elasticsearch cluster and the client.
On the **Cluster Information** page, find **Cluster Route**, and click **Add Route** on the right. In the displayed dialog box, set **IP Address** and **Subnet Mask**.
For details, see [Configuring Routes for an Elasticsearch Cluster](#).
7. On the client, enter the private network address of Kibana (for example, **https://192.168.0.xx:5601**) to go to the Kibana login page.
 - Non-security cluster: The Kibana console is displayed without asking for a username and password.
 - Security cluster: Enter the username and password on the login page and click **Log In** to log in to the Kibana console. The default username is **admin** and the password is the one specified during cluster creation.
8. After the login is successful, you can access the Elasticsearch cluster through Kibana.

More: Configuring a Custom Kibana Base Path for Public Network Access

With Elasticsearch 7.10.2, you can configure a custom Base Path for Kibana access after Kibana public network access is enabled. Or you can use the default Kibana Base Path instead.




WARNING

- This feature is available for Elasticsearch 7.10.2 clusters (the image version is no earlier than 7.10.2_24.3.3_x.x.x) for which Kibana public network access is enabled.
- Only the cluster administrator **admin** under Global Tenant has the permission to configure a custom Kibana Base Path, which applies globally.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console using the administrator account.
 - Username: **admin** (default administrator account)
 - Password: Enter the administrator password you set when creating the cluster in security mode.
4. After a successful login, choose **Stack Management** in the left navigation pane of the Kibana console.
5. Choose **Advanced Settings** on the left of the **Stack Management** page.
6. On the **Settings** page, set **Base Path Alias**. The value must start with a slash (/) and must not end with one. A multi-layer path is allowed, but its length cannot exceed 255 characters.

Figure 2-5 Custom Base Path



7. Click **Save changes**. The saved changes take effect in approximately 10 seconds.
8. Access Kibana using the Kibana public network address plus *Base Path Alias*. For example, if the Kibana public network address of the Elasticsearch cluster is **https://xx.xx.xx.xx:5601** and the configured **Base Path Alias** is **/test**, you can access Kibana via **https://xx.xx.xx.xx:5601/test** from the public network.

Related Documents

- For details about how to use in-house developed Kibana to access Elasticsearch clusters in CSS, see [How Do I Connect In-house Developed Kibana to an Elasticsearch Cluster in CSS?](#)
- If Kibana cannot be launched, see [Failed to Open Kibana](#).
- For routine O&M tasks (such as shard adjustment, index management, and performance monitoring), you are advised to use Cerebro for cluster login. For details, see [Logging In to an Elasticsearch Cluster Through Cerebro](#).

2.5.3 Logging In to an Elasticsearch Cluster Through Cerebro

Cerebro is an open-source Elasticsearch web admin tool, designed for real-time monitoring and efficient cluster operations. In CSS, Cerebro is pre-built for each Elasticsearch cluster. You can start Cerebro with one click, without installing anything. Cerebro allows you to quickly check a cluster's health status, node distribution, and index details. For routine O&M tasks, including shard adjustment, index management, and performance monitoring, we recommend using Cerebro. Its intuitive cluster topology views and comprehensive management features streamline O&M workflows, reducing complexity while improving efficiency. Cerebro is suitable for different roles, including developers, O&M teams, and data analysts.

Accessing a Cluster Through Cerebro on the CSS Console

The Elasticsearch clusters created in CSS provide a built-in Cerebro component, using which you can quickly access these clusters.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and choose **More > Cerebro** in the **Operation** column. The Cerebro login page is displayed.
 - 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.
4. After the login is successful, you can access clusters through Cerebro.

Accessing a Cluster Using an In-House Built Cerebro

CAUTION

- When using in-house built Cerebro to access a cluster, make sure the network between them is connected.
- Make sure the version of the in-house built Cerebro is compatible with the Elasticsearch cluster version.

-
1. Obtain the cluster access address. For details about how to obtain the cluster access address for different network configurations, see [Network Configuration](#).
 2. Start Cerebro, and enter the cluster access address.
 - For a cluster in security mode, enter `https://access address:9200`.
To log in to a cluster in security mode, you need to provide a username and password.
 - For a cluster in non-security mode, enter `http://access address:9200`.

2.5.4 Accessing an Elasticsearch Cluster Using cURL Commands

Open-source Elasticsearch provides a series of RESTful APIs. You can run cURL commands to access these APIs using tools such as Kibana and Postman. This topic describes how to use cURL commands to access an Elasticsearch cluster.

Prerequisites

- The target Elasticsearch cluster is available.
- An ECS has been created, and it is in the same VPC and security group as the CSS cluster.
 - If they are not in the same security group, modify the ECS security group, or configure the inbound and outbound rules of the group to allow all access from the cluster. For details, see [Configuring Security Group Rules](#).
 - For details about how to use an ECS, see [Purchasing a Custom ECS](#).

Accessing a Cluster

1. Obtain the private network address of the cluster. It is used to access the cluster.

This topic uses a private IP address as an example to describe how to access a cluster. The cluster access address varies with the network configurations used. For details, see [Network Configuration](#).

- a. Log in to the [CSS management console](#).
- b. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
- c. In the cluster list, obtain the target cluster's private IP address from the **Private IP Address** column. Generally, the IP address format is *<host>:<port>* or *<host>:<port>,<host>:<port>*.

If the cluster has only one node, the IP address and port number of this 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**.

2. Run one of the following commands on the ECS to access the cluster. The access command varies according to the security mode of the cluster.
 - Cluster in non-security mode

```
curl "http://<host>:<port>"
```
 - Cluster in security mode + HTTP

```
curl -u <user>:<password> "http://<host>:<port>"
```
 - Cluster in security mode + HTTPS

```
curl -u <user>:<password> -k "https://<host>:<port>"
```

Table 2-38 Variables

Variable	Description
<host>	IP address of each node in the cluster. If the cluster contains multiple nodes, there will be multiple IP addresses. You can use any of them.
<port>	Port number for accessing a cluster node. Generally, the port number is 9200.
<user>	Username for accessing the cluster.
<password>	Password of the user. If the password contains special characters, enclose the username and password in single quotation marks, for example, <code>curl -u "user:password!" "http://<host>:<port>"</code> .

An access example is as follows:

```
curl "http://10.62.176.32:9200"
```

Information similar to the following is displayed:

```
HTTP/1.1 200 OK
content-type: application/json; charset=UTF-8
content-length: 513

{
  "name" : "xxx-1",
  "cluster_name" : "xxx",
  "cluster_uuid" : "xxx_uuid",
  "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"
}
```

NOTE

For more commands, see the [Elasticsearch documentation](#).

2.5.5 Accessing an Elasticsearch Cluster Using Java

2.5.5.1 Accessing an Elasticsearch Cluster Through the High Level REST Client

Elasticsearch provides SDK (High Level REST Client) for connecting to a cluster. This client encapsulates Elasticsearch APIs. You only need to construct required structures to access the Elasticsearch cluster. For details about how to use the Rest

Client, see the official document at <https://www.elastic.co/guide/en/elasticsearch/client/java-rest/current/java-rest-high.html>.

This section describes how to use the High Level REST Client to access a CSS cluster. The High Level REST Client can be connected to the cluster in one of the following ways:

- **Connecting to a Non-Security Mode Cluster Through the High Level REST Client:** suitable for clusters with the security mode disabled
- **Connecting to a Security-Mode Cluster Through High Level REST Client (Without a Security Certificate):** suitable for clusters in security mode+HTTP, and for clusters in security mode+HTTPS (without using a certificate)
- **Connecting to a Security-Mode Cluster Through High Level REST Client (With a Security Certificate):** suitable for clusters in security mode+HTTPS

Constraints

The High Level REST Client version should match the Elasticsearch version. For example, use High Level REST Client 7.6.2 to access an Elasticsearch 7.6.2 cluster. If your Java High Level REST Client version is later than the Elasticsearch cluster and there are incompatibility issues with a few requests, you can use **RestHighLevelClient.getLowLevelClient()** to obtain Low Level Client and customize the Elasticsearch request content.

Prerequisites

- There are available CSS clusters.
- You have a server running Java that can communicate with the CSS cluster.
- Depending on the network configuration method used, obtain the cluster access address. For details, see [Network Configuration](#).
- Install JDK 1.8 on the server. You can download JDK 1.8 from: <https://www.oracle.com/technetwork/java/javase/downloads/jdk8-downloads-2133151.html>
- Java dependencies have been installed on the server.

7.6.2 indicates the version of the Elasticsearch Java client.

– Maven mode:

```
<dependency>
  <groupId>org.elasticsearch.client</groupId>
  <artifactId>elasticsearch-rest-high-level-client</artifactId>
  <version>7.6.2</version>
</dependency>
<dependency>
  <groupId>org.elasticsearch</groupId>
  <artifactId>elasticsearch</artifactId>
  <version>7.6.2</version>
</dependency>
```

– Gradle mode:

```
compile group: 'org.elasticsearch.client', name: 'elasticsearch-rest-high-level-client', version: '7.6.2'
```

Connecting to a Non-Security Mode Cluster Through the High Level REST Client

Use the High Level REST Client to connect to a non-security mode cluster and check whether the **test** index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
import org.elasticsearch.client.RequestOptions;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;
import org.elasticsearch.client.indices.GetIndexRequest;

import java.io.IOException;
import java.util.Arrays;
import java.util.List;

/**
 * Use Rest Hive Level to connect to a non-security cluster.
 */
public class Main {
    public static void main(String[] args) throws IOException {
        List<String> host = Arrays.asList("xx.xx.xx.xx", "xx.xx.xx.xx");
        RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, 9200, "http"));
        final RestHighLevelClient client = new RestHighLevelClient(builder);
        GetIndexRequest indexRequest = new GetIndexRequest("test");
        boolean exists = client.indices().exists(indexRequest, RequestOptions.DEFAULT);
        System.out.println(exists);
        client.close();
    }

    /**
     * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
     */
    public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
        return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
    }
}
```

host indicates the IP address of the cluster. If there are multiple IP addresses, separate them using commas (,). *test* indicates the index name to be queried.

Connecting to a Security-Mode Cluster Through High Level REST Client (Without a Security Certificate)

Connect to a security-mode cluster that uses either HTTP or HTTPS without using a security certificate.

The sample code is as follows:

```
import org.apache.http.HttpHost;
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.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.elasticsearch.action.admin.cluster.health.ClusterHealthRequest;
import org.elasticsearch.action.admin.cluster.health.ClusterHealthResponse;
import org.elasticsearch.client.RequestOptions;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;
import org.elasticsearch.client.indices.GetIndexRequest;
import org.elasticsearch.common.Nullable;

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 java.util.Arrays;
import java.util.List;
import java.util.Objects;

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;

/**
 * Connect to a security cluster through Rest High Level (without using certificates).
 */
public class Main {
    /**
     * Create a class for the client. Define the create function.
     */
    public static RestHighLevelClient create(List<String> host, int port, String protocol, int connectTimeout,
int connectionRequestTimeout, int socketTimeout, String username, String password) throws IOException{
        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(constructHttpHosts(host, port, protocol))
            .setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout))
            .setConnectionRequestTimeout(connectionRequestTimeout)
            .setSocketTimeout(socketTimeout)
            .setHttpClientConfigCallback(httpClientConfigCallback);
        final RestHighLevelClient client = new RestHighLevelClient(builder);
        logger.info("es rest client build success {} ", client);

        ClusterHealthRequest request = new ClusterHealthRequest();
        ClusterHealthResponse response = client.cluster().health(request, RequestOptions.DEFAULT);
        logger.info("es rest client health response {} ", response);
        return client;
    }

    /**
     * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
     */
    public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
        return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
    }

    /**
     * Configure trustAllCerts to ignore the certificate configuration.
     */
    public 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;
    }
}

};

private static final Logger logger = LogManager.getLogger(Main.class);

static 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;
    }
}

public static class NullHostNameVerifier implements HostnameVerifier {
    @Override
```

```
public boolean verify(String arg0, SSLSession arg1) {  
    return true;  
}  
}  
  
/**  
 * The following is an example of the main function. Call the create function to create a client and check  
 * whether the test index exists.  
 */  
public static void main(String[] args) throws IOException {  
    RestHighLevelClient client = create(Arrays.asList("xx.xx.xx.xx", "xx.xx.xx.xx"), 9200, "https", 1000, 1000,  
1000, "username", "password");  
    GetIndexRequest indexRequest = new GetIndexRequest("test");  
    boolean exists = client.indices().exists(indexRequest, RequestOptions.DEFAULT);  
    System.out.println(exists);  
    client.close();  
}  
}
```

Table 2-39 Variables

Parameter	Description
host	IP address for accessing the Elasticsearch cluster. If there are multiple IP addresses, separate them with commas (,).
port	Access port of the Elasticsearch cluster. The default value is 9200 .
protocol	Connection protocol, which can be http or https .
connectTimeout	Socket connection timeout (in ms).
connectionRequestTimeout	Socket connection request timeout (in ms).
socketTimeout	Socket request timeout (in ms).
username	Username for accessing the cluster.
password	Password of the user.

Connecting to a Security-Mode Cluster Through High Level REST Client (With a Security Certificate)

Connect to a security-mode Elasticsearch cluster that uses HTTPS using a security certificate. The sample code is as follows:

CAUTION

For how to obtain and upload a security certificate, see [Obtaining and Uploading a Security Certificate](#).

```
import org.apache.http.HttpHost;  
import org.apache.http.auth.AuthScope;
```

```
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.client.CredentialsProvider;
import org.apache.http.conn.ssl.NoopHostnameVerifier;
import org.apache.http.impl.client.BasicCredentialsProvider;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.elasticsearch.action.admin.cluster.health.ClusterHealthRequest;
import org.elasticsearch.action.admin.cluster.health.ClusterHealthResponse;
import org.elasticsearch.client.RequestOptions;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;
import org.elasticsearch.client.indices.GetIndexRequest;
import org.elasticsearch.common.Nullable;

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 java.util.Arrays;
import java.util.List;
import java.util.Objects;

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

/**
 * Use Rest Hive Level to connect to a security cluster (using an HTTPS certificate).
 */
public class Main {
    public static RestHighLevelClient create(List<String> host, int port, String protocol, int connectTimeout,
int connectionRequestTimeout, int socketTimeout, String username, String password, String certFilePath,
String certPassword) throws IOException {

        final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
        credentialsProvider.setCredentials(AuthScope.ANY, new UsernamePasswordCredentials(username,
password));
        SSLContext sc = null;
        try {
            TrustManager[] tm = {new MyX509TrustManager(certFilePath, certPassword)};
            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();
        }

        SSLIOStrategy sessionStrategy = new SSLIOStrategy(sc, new NoopHostnameVerifier());
        SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,
credentialsProvider);

        RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, port, protocol))
            .setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout))
            .setConnectionRequestTimeout(connectionRequestTimeout)
            .setSocketTimeout(socketTimeout)
            .setHttpClientConfigCallback(httpClientConfigCallback);
        final RestHighLevelClient client = new RestHighLevelClient(builder);
        logger.info("es rest client build success {} ", client);

        ClusterHealthRequest request = new ClusterHealthRequest();
```



```
ClusterHealthResponse response = client.cluster().health(request, RequestOptions.DEFAULT);
logger.info("es rest client health response {} ", response);
return client;
}

/**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
    return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
}

/**
 * SecuredHttpClientConfigCallback class definition
 */
static class SecuredHttpClientConfigCallback implements RestClientBuilder.HttpClientConfigCallback {
    @Nullable
    private final CredentialsProvider credentialsProvider;

    private final SSLIOSessionStrategy sslStrategy;

    SecuredHttpClientConfigCallback(final SSLIOSessionStrategy sslStrategy,
        @Nullable final CredentialsProvider credentialsProvider) {
        this.sslStrategy = Objects.requireNonNull(sslStrategy);
        this.credentialsProvider = credentialsProvider;
    }

    @Nullable
    CredentialsProvider getCredentialsProvider() {
        return credentialsProvider;
    }

    SSLIOSessionStrategy getSSLStrategy() {
        return sslStrategy;
    }

    @Override
    public HttpAsyncClientBuilder customizeHttpClient(final HttpAsyncClientBuilder httpClientBuilder) {
        httpClientBuilder.setSSLStrategy(sslStrategy);
        if (credentialsProvider != null) {
            httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
        }
        return httpClientBuilder;
    }
}

private static final Logger logger = LogManager.getLogger(Main.class);

public static class MyX509TrustManager implements X509TrustManager {
    X509TrustManager sunJSSEX509TrustManager;

    MyX509TrustManager(String certFilePath, String certPassword) throws Exception {
        File file = new File(certFilePath);
        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, certPassword.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];
    }
}

/**
 * The following is an example of the main function. Call the create function to create a client and check
 * whether the test index exists.
 */
public static void main(String[] args) throws IOException {
    RestHighLevelClient client = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200, "https",
    1000, 1000, 1000, "username", "password", "certFilePath", "certPassword");
    GetIndexRequest indexRequest = new GetIndexRequest("test");
    boolean exists = client.indices().exists(indexRequest, RequestOptions.DEFAULT);
    System.out.println(exists);
    client.close();
}
```

Table 2-40 Function parameters

Parameter	Description
host	IP address for accessing the Elasticsearch cluster. If there are multiple IP addresses, separate them with commas (,).
port	Access port of the Elasticsearch cluster. The default value is 9200 .
protocol	Connection protocol. Set this parameter to https .
connectTimeout	Socket connection timeout (in ms).
connectionRequestTimeout	Socket connection request timeout (in ms).
socketTimeout	Socket request timeout (in ms).
username	Username for accessing the cluster.
password	Password of the user.
certFilePath	Certificate path.
certPassword	Certificate password.

Obtaining and Uploading a Security Certificate

To access a security-mode Elasticsearch cluster that uses HTTPS, a security certificate must be loaded. Perform the following steps to obtain the security certificate and upload it to the client:

1. Obtain the security certificate **CloudSearchService.cer**.
 - a. Log in to the [CSS management console](#).
 - b. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
 - c. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
 - d. Click the **Overview** tab. In the **Configuration** area, click **Download Certificate** next to **HTTPS Access**.

Figure 2-6 Downloading a security certificate



2. Convert the security certificate **CloudSearchService.cer**. Upload the downloaded security certificate to the client and use keytool to convert the **.cer** certificate into a **.jks** certificate that can be read by Java.

- In Linux, run the following command to convert the certificate:

```
keytool -import -alias newname -keystore ./truststore.jks -file ./CloudSearchService.cer
```
- In Windows, run the following command to convert the certificate:

```
keytool -import -alias newname -keystore .\truststore.jks -file .\CloudSearchService.cer
```

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

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

2.5.5.2 Accessing an Elasticsearch Cluster Through the Low Level REST Client

The high-level client is encapsulated based on the low-level client. If the method calls (such as **.search** and **.bulk**) in the high-level client cannot meet the requirements or has compatibility issues, you can use the low-level client. You can even use **HighLevelClient.getLowLevelClient()** to directly obtain the low-level client. A low-level client allows you to customize the request structure, which is more flexible and supports all the request formats of Elasticsearch, such as GET, POST, DELETE, and HEAD.

This section describes how to use the Low Level REST Client to access the CSS cluster. The methods are as follows. For each method, you can directly create a REST low-level client, or create a high-level client and then invoke **getLowLevelClient()** to obtain a low-level client.

- **Connecting to a Non-Security Cluster Through the Low Level REST Client:** suitable for clusters using non-security mode

- **Connecting to a Security-Mode Cluster Through Low Level REST Client (Without a Security Certificate)**: suitable for clusters in security mode+HTTP, and for clusters in security mode+HTTPS (without using certificates)
- **Connecting to a Security-Mode Cluster Through Low Level REST Client (With a Security Certificate)**: suitable for clusters in security mode+HTTPS

Precautions

The Low Level REST Client version should match the Elasticsearch version. For example, use Low Level REST Client 7.6.2 to access the Elasticsearch 7.6.2 cluster.

Prerequisites

- The CSS cluster is available.
- Ensure that the server running Java can communicate with the CSS cluster.
- Depending on the network configuration method used, obtain the cluster access address. For details, see [Network Configuration](#).
- Install JDK 1.8 on the server. You can download JDK 1.8 from: <https://www.oracle.com/technetwork/java/javase/downloads/jdk8-downloads-2133151.html>
- Declare the Apache version in Maven mode. The following code uses version 7.6.2 as an example.

7.6.2 indicates the version of the Elasticsearch Java client.

```
<dependency>
  <groupId>org.elasticsearch.client</groupId>
  <artifactId>elasticsearch-rest-client</artifactId>
  <version>7.6.2</version>
</dependency>
<dependency>
  <groupId>org.elasticsearch</groupId>
  <artifactId>elasticsearch</artifactId>
  <version>7.6.2</version>
</dependency>
```

Connecting to a Non-Security Cluster Through the Low Level REST Client

- **Method 1: Directly create a Low Level REST Client.**

```
import org.apache.http.HttpHost;
import org.elasticsearch.client.Request;
import org.elasticsearch.client.Response;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;

import java.io.IOException;
import java.util.Arrays;
import java.util.List;

public class Main {

    public static void main(String[] args) throws IOException {
        List<String> host = Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx");
        RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, 9200, "http"));
        /**
         * Create a Low Level REST Client.
         */
        RestClient lowLevelClient = builder.build();
        /**
         * Check whether the test index exists. If the index exists, 200 is returned. If the index does not
         exist, 404 is returned.
```

```
    */
    Request request = new Request("HEAD", "/" + test);
    Response response = lowLevelClient.performRequest(request);
    System.out.println(response.getStatusLine().getStatusCode());
    lowLevelClient.close();
}

/**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
    return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
}
}
```

- **Method 2: Create a high-level client and then call `getLowLevelClient()` to obtain a low-level client.**

```
import org.apache.http.HttpHost;
import org.elasticsearch.client.Request;
import org.elasticsearch.client.Response;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;

import java.io.IOException;
import java.util.Arrays;
import java.util.List;

public class Main {

    public static void main(String[] args) throws IOException {
        List<String> host = Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx");
        RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, 9200, "http"));
        final RestHighLevelClient restHighLevelClient = new RestHighLevelClient(builder);
        /**
         * Create a high-level client and then call getLowLevelClient() to obtain a low-level client.
         * The code differs from the client creation code only in the following line:
         */
        final RestClient lowLevelClient = restHighLevelClient.getLowLevelClient();
        /**
         * Check whether the test index exists. If the index exists, 200 is returned. If the index does not
         * exist, 404 is returned.
         */
        Request request = new Request("HEAD", "/" + test);
        Response response = lowLevelClient.performRequest(request);
        System.out.println(response.getStatusLine().getStatusCode());
        lowLevelClient.close();
    }

    /**
     * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
     */
    public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
        return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
    }
}
```

host indicates the IP address of the cluster. If there are multiple IP addresses, separate them using commas (,). *test* indicates the index name to be queried.

Connecting to a Security-Mode Cluster Through Low Level REST Client (Without a Security Certificate)

- **Method 1: Directly create a Low Level REST Client.**

```
import org.apache.http.HttpHost;
import org.apache.http.HttpResponse;
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.impl.client.DefaultConnectionKeepAliveStrategy;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOSessionStrategy;
import org.apache.http.protocol.HttpContext;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.elasticsearch.client.Request;
import org.elasticsearch.client.Response;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.common.Nullable;

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 java.util.Arrays;
import java.util.List;
import java.util.Objects;
import java.util.concurrent.TimeUnit;

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 Main {

    /**
     * Create a class for the client. Define the create function.
     */
    public static RestClient create(List<String> host, int port, String protocol, int connectTimeout, int
connectionRequestTimeout, int socketTimeout, String username, String password) throws
IOException {
        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(constructHttpHosts(host, port, protocol))
            .setRequestConfigCallback(requestConfig ->
requestConfig.setConnectTimeout(connectTimeout)
                .setConnectionRequestTimeout(connectionRequestTimeout)
                .setSocketTimeout(socketTimeout))
            .setHttpClientConfigCallback(httpClientConfigCallback);
        final RestClient client = builder.build();
        logger.info("es rest client build success {} ", client);
        return client;
    }

    /**
     * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
     */
    public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
```

```
        return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
    }

    /**
     * Configure trustAllCerts to ignore the certificate configuration.
     */
    public 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;
            }
        }
    };

    /**
     * The CustomConnectionKeepAliveStrategy function is used to set the connection keepalive time when
     there are a large number of short connections or when the number of data requests is small.
     */
    public static class CustomConnectionKeepAliveStrategy extends
DefaultConnectionKeepAliveStrategy {
        public static final CustomConnectionKeepAliveStrategy INSTANCE = new
CustomConnectionKeepAliveStrategy();

        private CustomConnectionKeepAliveStrategy() {
            super();
        }

        /**
         * Maximum keepalive time (in minutes)
         * The default value is 10 minutes. You can set it based on the number of TCP connections in
TIME_WAIT state. If there are too many TCP connections, you can increase the value.
         */
        private final long MAX_KEEP_ALIVE_MINUTES = 10;

        @Override
        public long getKeepAliveDuration(HttpResponse response, HttpContext context) {
            long keepAliveDuration = super.getKeepAliveDuration(response, context);
            // <0 indicates an unlimited keepalive period.
            // Change the period from unlimited to a default period.
            if (keepAliveDuration < 0) {
                return TimeUnit.MINUTES.toMillis(MAX_KEEP_ALIVE_MINUTES);
            }
            return keepAliveDuration;
        }
    }

    private static final Logger logger = LogManager.getLogger(Main.class);

    static 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;
}
}

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

/**
* The following is an example of the main function. Call the create function to create a Low Level
REST Client and check whether the test index exists.
*/
public static void main(String[] args) throws IOException {
    RestClient lowLevelClient = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200, "http",
1000, 1000, 1000, "username", "password");
    Request request = new Request("HEAD", "/test");
    Response response = lowLevelClient.performRequest(request);
    System.out.println(response.getStatusLine().getStatusCode());
    lowLevelClient.close();
}
}
```


- **Method 2: Create a high-level client and then call `getLowLevelClient()` to obtain a low-level client.**

```
import org.apache.http.HttpHost;
import org.apache.http.HttpResponse;
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.impl.client.DefaultConnectionKeepAliveStrategy;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOSessionStrategy;
import org.apache.http.protocol.HttpContext;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.elasticsearch.client.Request;
import org.elasticsearch.client.Response;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.common.Nullable;

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 java.util.Arrays;
import java.util.List;
import java.util.Objects;
import java.util.concurrent.TimeUnit;

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;

import org.elasticsearch.client.RestHighLevelClient;

public class Main13 {

    /**
     * Create a class for the client. Define the create function.
     */
    public static RestHighLevelClient create(List<String> host, int port, String protocol, int
connectTimeout, int connectionRequestTimeout, int socketTimeout, String username, String
password) throws IOException {

        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(constructHttpHosts(host, port, protocol))
            .setRequestConfigCallback(requestConfig ->
requestConfig.setConnectTimeout(connectTimeout)
                .setConnectionRequestTimeout(connectionRequestTimeout)
                .setSocketTimeout(socketTimeout))
            .setHttpClientConfigCallback(httpClientConfigCallback);
```

```
        final RestHighLevelClient client = new RestHighLevelClient(builder);
        logger.info("es rest client build success {} ", client);
        return client;
    }

    /**
     * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
     */
    public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
        return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
    }

    /**
     * Configure trustAllCerts to ignore the certificate configuration.
     */
    public 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;
            }
        }
    };

    /**
     * The CustomConnectionKeepAliveStrategy function is used to set the connection keepalive time when
     there are a large number of short connections or when the number of data requests is small.
     */
    public static class CustomConnectionKeepAliveStrategy extends
DefaultConnectionKeepAliveStrategy {
        public static final CustomConnectionKeepAliveStrategy INSTANCE = new
CustomConnectionKeepAliveStrategy();

        private CustomConnectionKeepAliveStrategy() {
            super();
        }

        /**
         * Maximum keepalive time (in minutes)
         * The default value is 10 minutes. You can set it based on the number of TCP connections in
TIME_WAIT state. If there are too many TCP connections, you can increase the value.
         */
        private final long MAX_KEEP_ALIVE_MINUTES = 10;

        @Override
        public long getKeepAliveDuration(HttpResponse response, HttpContext context) {
            long keepAliveDuration = super.getKeepAliveDuration(response, context);
            // <0 indicates an unlimited keepalive period.
            // Change the period from unlimited to a default period.
            if (keepAliveDuration < 0) {
                return TimeUnit.MINUTES.toMillis(MAX_KEEP_ALIVE_MINUTES);
            }
            return keepAliveDuration;
        }
    }

    private static final Logger logger = LogManager.getLogger(Main.class);

    static 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
    HttpClientBuilder#setDefaultCredentialsProvider(CredentialsProvider) credential provider},
     *
     * @param httpClientBuilder The client to configure.
     * @return Always {@code httpClientBuilder}.
     */
    @Override
    public HttpClientBuilder customizeHttpClient(final HttpClientBuilder
    httpClientBuilder) {
        // enable SSL / TLS
        httpClientBuilder.setSSLStrategy(sslStrategy);
        // enable user authentication
        if (credentialsProvider != null) {
            httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
        }
        return httpClientBuilder;
    }
}

public static class NullHostNameVerifier implements HostnameVerifier {
    @Override
    public boolean verify(String arg0, SSLSession arg1) {
        return true;
    }
}
/**
 * The following is an example of the main function. Call the create function to create a high-level
    client, call the getLowLevelClient() function to obtain a low-level client, and check whether the test
    index exists.
 */
```

```
public static void main(String[] args) throws IOException {  
    RestHighLevelClient client = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200,  
    "http", 1000, 1000, 1000, "username", "password");  
    RestClient lowLevelClient = client.getLowLevelClient();  
    Request request = new Request("HEAD", "test");  
    Response response = lowLevelClient.performRequest(request);  
    System.out.println(response.getStatusLine().getStatusCode());  
    lowLevelClient.close();  
}  
}
```

Table 2-41 Variables

Parameter	Description
host	IP address for accessing the Elasticsearch cluster. If there are multiple IP addresses, separate them with commas (,).
port	Access port of the Elasticsearch cluster. The default value is 9200 .
protocol	Connection protocol, which can be http or https .
connectTimeout	Socket connection timeout (in ms).
connectionRequestTimeout	Socket connection request timeout (in ms).
socketTimeout	Socket request timeout (in ms).
username	Username for accessing the cluster.
password	Password of the user.

Connecting to a Security-Mode Cluster Through Low Level REST Client (With a Security Certificate)

CAUTION

For how to obtain and upload a security certificate, see [Obtaining and Uploading a Security Certificate](#).

- **Method 1: Directly create a Low Level REST Client.**

```
import org.apache.http.HttpHost;  
import org.apache.http.auth.AuthScope;  
import org.apache.http.auth.UsernamePasswordCredentials;  
import org.apache.http.client.CredentialsProvider;  
import org.apache.http.conn.ssl.NoopHostnameVerifier;  
import org.apache.http.impl.client.BasicCredentialsProvider;  
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;  
import org.apache.http.nio.conn.ssl.SSLIOStrategy;  
import org.apache.logging.log4j.LogManager;  
import org.apache.logging.log4j.Logger;  
import org.elasticsearch.client.Request;
```

```
import org.elasticsearch.client.Response;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.common.Nullable;

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 java.util.Arrays;
import java.util.List;
import java.util.Objects;

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

public class Main13 {

    private static final Logger logger = LogManager.getLogger(Main.class);

    /**
     * Create a class for the client. Define the create function.
     */
    public static RestClient create(List<String> host, int port, String protocol, int connectTimeout, int
connectionRequestTimeout, int socketTimeout, String username, String password, String certFilePath,
String certPassword) throws IOException {
        final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
        credentialsProvider.setCredentials(AuthScope.ANY, new
UsernamePasswordCredentials(username, password));
        SSLContext sc = null;
        try {
            TrustManager[] tm = {new MyX509TrustManager(certFilePath, certPassword)};
            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
NoopHostnameVerifier());
        SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,
credentialsProvider);

        RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, port, protocol))
            .setRequestConfigCallback(requestConfig ->
requestConfig.setConnectTimeout(connectTimeout)
            .setConnectionRequestTimeout(connectionRequestTimeout)
            .setSocketTimeout(socketTimeout))
            .setHttpClientConfigCallback(httpClientConfigCallback);
        final RestClient client = builder.build();
        logger.info("es rest client build success {} ", client);
        return client;
    }

    /**
     * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
     */
    public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
        return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);}

    static class SecuredHttpClientConfigCallback implements
RestClientBuilder.HttpClientConfigCallback {
```

```
@Nullable
private final CredentialsProvider credentialsProvider;

private final SSLIOSessionStrategy sslStrategy;

SecuredHttpClientConfigCallback(final SSLIOSessionStrategy sslStrategy,
    @Nullable final CredentialsProvider credentialsProvider) {
    this.sslStrategy = Objects.requireNonNull(sslStrategy);
    this.credentialsProvider = credentialsProvider;
}

@Nullable
CredentialsProvider getCredentialsProvider() {
    return credentialsProvider;
}

SSLIOSessionStrategy getSSLStrategy() {
    return sslStrategy;
}

@Override
public HttpAsyncClientBuilder customizeHttpClient(final HttpAsyncClientBuilder
httpClientBuilder) {
    httpClientBuilder.setSSLStrategy(sslStrategy);
    if (credentialsProvider != null) {
        httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
    }
    return httpClientBuilder;
}

public static class MyX509TrustManager implements X509TrustManager {
    X509TrustManager sunJSSEX509TrustManager;

    MyX509TrustManager(String certFilePath, String certPassword) throws Exception {
        File file = new File(certFilePath);
        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, certPassword.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];
    }
}
```

```
    }  
  }  
  
  /**  
   * The following is an example of the main function. Call the create function to create a Low Level  
   * REST Client and check whether the test index exists.  
   */  
  public static void main(String[] args) throws IOException {  
    RestClient lowLevelClient = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200,  
"https", 1000, 1000, 1000, "username", "password", "certFilePath", "certPassword");  
    Request request = new Request("HEAD", "test");  
    Response response = lowLevelClient.performRequest(request);  
    System.out.println(response.getStatusLine().getStatusCode());  
    lowLevelClient.close();  
  }  
}
```

- **Method 2: Create a high-level client and then call `getLowLevelClient()` to obtain a low-level client.**

```
import org.apache.http.HttpHost;  
import org.apache.http.auth.AuthScope;  
import org.apache.http.auth.UsernamePasswordCredentials;  
import org.apache.http.client.CredentialsProvider;  
import org.apache.http.conn.ssl.NoopHostnameVerifier;  
import org.apache.http.impl.client.BasicCredentialsProvider;  
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;  
import org.apache.http.nio.conn.ssl.SSLIOStrategy;  
import org.apache.logging.log4j.LogManager;  
import org.apache.logging.log4j.Logger;  
import org.elasticsearch.action.admin.cluster.health.ClusterHealthRequest;  
import org.elasticsearch.action.admin.cluster.health.ClusterHealthResponse;  
import org.elasticsearch.client.Request;  
import org.elasticsearch.client.RequestOptions;  
import org.elasticsearch.client.Response;  
import org.elasticsearch.client.RestClient;  
import org.elasticsearch.client.RestClientBuilder;  
import org.elasticsearch.client.RestHighLevelClient;  
import org.elasticsearch.common.Nullable;  
  
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 java.util.Arrays;  
import java.util.List;  
import java.util.Objects;  
  
import javax.net.ssl.SSLContext;  
import javax.net.ssl.TrustManager;  
import javax.net.ssl.TrustManagerFactory;  
import javax.net.ssl.X509TrustManager;  
  
public class Main {  
  
    private static final Logger logger = LogManager.getLogger(Main.class);  
  
    /**  
     * Create a class for the client. Define the create function.  
     */  
    public static RestClient create(List<String> host, int port, String protocol, int  
connectTimeout, int connectionRequestTimeout, int socketTimeout, String username, String password,  
String certFilePath, String certPassword) throws IOException {  
        final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();  
        credentialsProvider.setCredentials(AuthScope.ANY, new  
UsernamePasswordCredentials(username, password));  
        SSLContext sc = null;  
        try {
```

```
TrustManager[] tm = {new MyX509TrustManager(certFilePath, certPassword)};
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
NoopHostnameVerifier());
SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,
    credentialsProvider);

RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, port, protocol))
    .setRequestConfigCallback(requestConfig ->
requestConfig.setConnectTimeout(connectTimeout)
    .setConnectionRequestTimeout(connectionRequestTimeout)
    .setSocketTimeout(socketTimeout))
    .setHttpClientConfigCallback(httpClientConfigCallback);
final RestHighLevelClient client = new RestHighLevelClient(builder);
logger.info("es rest client build success {}", client);

ClusterHealthRequest request = new ClusterHealthRequest();
ClusterHealthResponse response = client.cluster().health(request, RequestOptions.DEFAULT);
logger.info("es rest client health response {}", response);
return client;
}

/**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
    return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
}

static class SecuredHttpClientConfigCallback implements
RestClientBuilder.HttpClientConfigCallback {
    @Nullable
    private final CredentialsProvider credentialsProvider;

    private final SSLIOSessionStrategy sslStrategy;

    SecuredHttpClientConfigCallback(final SSLIOSessionStrategy sslStrategy,
        @Nullable final CredentialsProvider credentialsProvider) {
        this.sslStrategy = Objects.requireNonNull(sslStrategy);
        this.credentialsProvider = credentialsProvider;
    }

    @Nullable
    CredentialsProvider getCredentialsProvider() {
        return credentialsProvider;
    }

    SSLIOSessionStrategy getSSLStrategy() {
        return sslStrategy;
    }

    @Override
    public HttpAsyncClientBuilder customizeHttpClient(final HttpAsyncClientBuilder
httpClientBuilder) {
        httpClientBuilder.setSSLStrategy(sslStrategy);
        if (credentialsProvider != null) {
            httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
        }
        return httpClientBuilder;
    }
}

public static class MyX509TrustManager implements X509TrustManager {
    X509TrustManager sunJSSEX509TrustManager;
```



```
MyX509TrustManager(String certFilePath, String certPassword) throws Exception {
    File file = new File(certFilePath);
    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, certPassword.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];
}
}

/**
 * The following is an example of the main function. Call the create function to create a high-level
 * client, call the getLowLevelClient() function to obtain a low-level client, and check whether the test
 * index exists.
 */
public static void main(String[] args) throws IOException {
    RestHighLevelClient client = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200,
"https", 1000, 1000, 1000, "username", "password", "certFilePath", "certPassword");
    RestClient lowLevelClient = client.getLowLevelClient();
    Request request = new Request("HEAD", "test");
    Response response = lowLevelClient.performRequest(request);
    System.out.println(response.getStatusLine().getStatusCode());
    lowLevelClient.close();
}
}
```

Table 2-42 Function parameters

Parameter	Description
host	IP address for accessing the Elasticsearch cluster. If there are multiple IP addresses, separate them with commas (,).

Parameter	Description
port	Access port of the Elasticsearch cluster. The default value is 9200 .
protocol	Connection protocol. Set this parameter to https .
connectTimeout	Socket connection timeout (in ms).
connectionRequestTimeout	Socket connection request timeout (in ms).
socketTimeout	Socket request timeout (in ms).
username	Username for accessing the cluster.
password	Password of the user.
certFilePath	Certificate path.
certPassword	Certificate password.

Obtaining and Uploading a Security Certificate

To access a security-mode Elasticsearch cluster that uses HTTPS, a security certificate must be loaded. Perform the following steps to obtain the security certificate and upload it to the client:

1. Obtain the security certificate **CloudSearchService.cer**.
 - a. Log in to the [CSS management console](#).
 - b. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
 - c. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
 - d. Click the **Overview** tab. In the **Configuration** area, click **Download Certificate** next to **HTTPS Access**.

Figure 2-7 Downloading a security certificate



2. Convert the security certificate **CloudSearchService.cer**. Upload the downloaded security certificate to the client and use keytool to convert the **.cer** certificate into a **.jks** certificate that can be read by Java.
 - In Linux, run the following command to convert the certificate:
`keytool -import -alias newname -keystore ./truststore.jks -file ./CloudSearchService.cer`
 - In Windows, run the following command to convert the certificate:
`keytool -import -alias newname -keystore .\truststore.jks -file .\CloudSearchService.cer`In the preceding command, *newname* indicates the user-defined certificate name.

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

2.5.5.3 Accessing an Elasticsearch Cluster Through the Transport Client

You can use Transport Client to access a CSS cluster with the security mode disabled. For a cluster in security mode, you are advised to use [Accessing an Elasticsearch Cluster Through the High Level REST Client](#).

Precautions

- You are advised to use a Transport Client version that matches the Elasticsearch version. For example, use Transport Client 7.6.2 to access an Elasticsearch 7.6.2 cluster.
- This solution is suitable for clusters that are not using the security mode. Such clusters can only be accessed through private IP addresses.
- Transport Client has been deprecated in Elasticsearch 7.x. You are advised to use High Level Rest Client to access clusters instead.

Prerequisites

- The CSS cluster is available.
- Ensure that the server running Java can communicate with the CSS cluster.
- Install JDK 1.8 on the server. You can download JDK 1.8 from: <https://www.oracle.com/technetwork/java/javase/downloads/jdk8-downloads-2133151.html>
- Declare Java dependencies.

7.6.2 indicates the version of the Elasticsearch Java client.

```
<dependency>
  <groupId>org.elasticsearch.client</groupId>
  <artifactId>transport</artifactId>
  <version>7.6.2</version>
</dependency>
<dependency>
  <groupId>org.elasticsearch</groupId>
  <artifactId>elasticsearch</artifactId>
  <version>7.6.2</version>
</dependency>
```

Accessing a Cluster

The following code is an example of using Transport Client to connect to the Elasticsearch cluster and check whether the **test** index exists.

```
import org.elasticsearch.action.ActionFuture;
import org.elasticsearch.action.admin.indices.exists.indices.IndicesExistsRequest;
import org.elasticsearch.action.admin.indices.exists.indices.IndicesExistsResponse;
import org.elasticsearch.client.transport.TransportClient;
import org.elasticsearch.common.settings.Settings;
import org.elasticsearch.common.transport.TransportAddress;
import org.elasticsearch.transport.client.PreBuiltTransportClient;

import java.net.InetAddress;
import java.net.UnknownHostException;
import java.util.concurrent.ExecutionException;
```

```
public class Main {
    public static void main(String[] args) throws ExecutionException, InterruptedException,
    UnknownHostException {
        String cluster_name = "xxx";
        String host1 = "xx.xx.xx.xx";
        String host2 = "y.y.y.y";
        Settings settings = Settings.builder()
            .put("client.transport.sniff", false)
            .put("cluster.name", cluster_name)
            .build();
        TransportClient client = new PreBuiltTransportClient(settings)
            .addTransportAddress(new TransportAddress(InetAddress.getByName(host1), 9300))
            .addTransportAddress(new TransportAddress(InetAddress.getByName(host2), 9300));
        IndicesExistsRequest indicesExistsRequest = new IndicesExistsRequest("test");
        ActionFuture<IndicesExistsResponse> exists = client.admin().indices().exists(indicesExistsRequest);
        System.out.println(exists.get().isExists());
    }
}
```

In the information above, *cluster_name* indicates the cluster name, and *host1* and *host2* indicate the IP addresses of the cluster nodes. You can run the **GET _cat/nodes** command to obtain the IP addresses of the nodes.

2.5.5.4 Accessing an Elasticsearch Cluster Using Spring Boot

You can access a CSS cluster using Spring Boot. Spring Boot can connect to a cluster in any of the following ways:

- **Accessing an HTTP Cluster Through Spring Boot:** applicable to clusters in non-security mode and clusters in security mode+HTTP
- **Using Spring Boot to Access an HTTPS Cluster (Without Using Any Security Certificate):** suitable for clusters in security mode+HTTPS
- **Using Spring Boot to Access an HTTPS Cluster (Using a Security Certificate):** suitable for clusters in security mode+HTTPS

NOTE

For details about how to use Spring Boot, see the official document: <https://docs.spring.io/spring-boot/docs/current/reference/htmlsingle/>

Precautions

- The Elasticsearch High Level REST Client version should match the Elasticsearch version. For example, use High Level REST Client 7.10.2 to access an Elasticsearch 7.10.2 cluster.
- This section uses Spring Boot 2.5.5 as an example to describe how to connect Spring Boot to a cluster. The corresponding Spring Data Elasticsearch version is 4.2.x.

Prerequisites

- The CSS cluster is available.
- Ensure that the server running Java can communicate with the CSS cluster.
- Depending on the network configuration method used, obtain the cluster access address. For details, see [Network Configuration](#).
- Install JDK 1.8 on the server. You can download JDK 1.8 from: <https://www.oracle.com/technetwork/java/javase/downloads/jdk8-downloads-2133151.html>

- Create a Spring Boot project.
 - Declare Java dependencies.
- 7.10.2 indicates the version of the Elasticsearch Java client.

– Maven mode:

```
<parent>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-starter-parent</artifactId>
  <version>2.5.5</version>
</parent>
<dependencies>
  <dependency>
    <groupId>org.springframework.boot</groupId>
    <artifactId>spring-boot-starter-web</artifactId>
  </dependency>
  <dependency>
    <groupId>org.springframework.boot</groupId>
    <artifactId>spring-boot-starter-data-elasticsearch</artifactId>
  </dependency>
  <dependency>
    <groupId>org.elasticsearch.client</groupId>
    <artifactId>elasticsearch-rest-high-level-client</artifactId>
    <version>7.10.2</version>
  </dependency>
</dependencies>
```

Accessing an HTTP Cluster Through Spring Boot

This scenario applies to clusters in non-security mode or clusters in security mode +HTTP.

Configuration file:

```
elasticsearch.url=host1:9200,host2:9200
// You do not need to configure the following two lines for a non-security cluster.
elasticsearch.username=username
elasticsearch.password=password
```

Table 2-43 Parameter description

Parameter	Description
host	IP address for accessing the Elasticsearch cluster.
username	Username for accessing the cluster.
password	Password of the user.

Code:

 NOTE

- **com.xxx** indicates the project directory, for example, **com.company.project**.
- **com.xxx.repository** is the repository directory, which is defined by **extends org.springframework.data.elasticsearch.repository.ElasticsearchRepository**.

```
package com.xxx.configuration;

import org.elasticsearch.client.RestHighLevelClient;
import org.springframework.beans.factory.annotation.Value;
```

```
import org.springframework.context.annotation.Bean;
import org.springframework.context.annotation.ComponentScan;
import org.springframework.context.annotation.Configuration;
import org.springframework.data.elasticsearch.client.ClientConfiguration;
import org.springframework.data.elasticsearch.client.RestClients;
import org.springframework.data.elasticsearch.config.AbstractElasticsearchConfiguration;
import org.springframework.data.elasticsearch.repository.config.EnableElasticsearchRepositories;

@Configuration
@EnableElasticsearchRepositories(basePackages = "com.xxx.repository")
@ComponentScan(basePackages = "com.xxx")
public class Config extends AbstractElasticsearchConfiguration {

    @Value("${elasticsearch.url}")
    public String elasticsearchUrl;

    // You do not need to set the following two parameters for a non-security cluster.
    @Value("${elasticsearch.username}")
    public String elasticsearchUsername;

    @Value("${elasticsearch.password}")
    public String elasticsearchPassword;

    @Override
    @Bean
    public RestHighLevelClient elasticsearchClient() {
        final ClientConfiguration clientConfiguration = ClientConfiguration.builder()
            .connectedTo(StringHostParse(elasticsearchUrl))
            // For a non-security cluster, there is no need to configure withBasicAuth.
            .withBasicAuth(elasticsearchUsername, elasticsearchPassword)
            .build();

        return RestClients.create(clientConfiguration).rest();
    }

    private String[] StringHostParse(String hostAndPorts) {
        return hostAndPorts.split(",");
    }
}
```

Using Spring Boot to Access an HTTPS Cluster (Without Using Any Security Certificate)

You can connect to a cluster in Security mode + HTTPS without using any security certificate.

Configuration file:

```
elasticsearch.url=host1:9200,host2:9200
elasticsearch.username=username
elasticsearch.password=password
```

Table 2-44 Parameter description

Parameter	Description
host	IP address for accessing the Elasticsearch cluster.
username	Username for accessing the cluster.
password	Password of the user.

Code:

 NOTE

- **com.xxx** indicates the project directory, for example, **com.company.project**.
- **com.xxx.repository** is the repository directory, which is defined by **extends org.springframework.data.elasticsearch.repository.ElasticsearchRepository**.

```
package com.xxx.configuration;
import org.elasticsearch.client.RestHighLevelClient;
import org.springframework.beans.factory.annotation.Value;
import org.springframework.context.annotation.Bean;
import org.springframework.context.annotation.ComponentScan;
import org.springframework.context.annotation.Configuration;
import org.springframework.data.elasticsearch.client.ClientConfiguration;
import org.springframework.data.elasticsearch.client.RestClients;
import org.springframework.data.elasticsearch.config.AbstractElasticsearchConfiguration;
import org.springframework.data.elasticsearch.repository.config.EnableElasticsearchRepositories;
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;
@Configuration
@EnableElasticsearchRepositories(basePackages = "com.xxx.repository")
@ComponentScan(basePackages = "com.xxx")
public class Config extends AbstractElasticsearchConfiguration {
    @Value("${elasticsearch.url}")
    public String elasticsearchUrl;
    @Value("${elasticsearch.username}")
    public String elasticsearchUsername;
    @Value("${elasticsearch.password}")
    public String elasticsearchPassword;
    @Override
    @Bean
    public RestHighLevelClient elasticsearchClient() {
        SSLContext sc = null;
        try {
            sc = SSLContext.getInstance("SSL");
            sc.init(null, trustAllCerts, new SecureRandom());
        } catch (KeyManagementException | NoSuchAlgorithmException e) {
            e.printStackTrace();
        }
        final ClientConfiguration clientConfiguration = ClientConfiguration.builder()
            .connectedTo(StringHostParse(elasticsearchUrl))
            .usingSsl(sc, new NullHostNameVerifier())
            .withBasicAuth(elasticsearchUsername, elasticsearchPassword)
            .build();
        return RestClients.create(clientConfiguration).rest();
    }
    private String[] StringHostParse(String hostAndPorts) {
        return hostAndPorts.split(",");
    }
    public 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;  
    }  
}
```

Using Spring Boot to Access an HTTPS Cluster (Using a Security Certificate)

You can use a security certificate to connect to a cluster in security mode + HTTPS.

1. **Obtaining and Uploading a Security Certificate.**

2. **application.properties** configuration file:

```
elasticsearch.url=host1:9200,host2:9200  
elasticsearch.username=username  
elasticsearch.password=password
```

Table 2-45 Parameter description

Parameter	Description
host	IP address for accessing the Elasticsearch cluster.
username	Username for accessing the cluster.
password	Password of the user.

3. Code:

 NOTE

- **com.xxx** indicates the project directory, for example, **com.company.project**.
- **com.xxx.repository** is the repository directory, which is defined by **extends org.springframework.data.elasticsearch.repository.ElasticsearchRepository**.

```
package com.xxx.configuration;  
import org.elasticsearch.client.RestHighLevelClient;  
import org.springframework.beans.factory.annotation.Value;  
import org.springframework.context.annotation.Bean;  
import org.springframework.context.annotation.ComponentScan;  
import org.springframework.context.annotation.Configuration;  
import org.springframework.data.elasticsearch.client.ClientConfiguration;  
import org.springframework.data.elasticsearch.client.RestClients;  
import org.springframework.data.elasticsearch.config.AbstractElasticsearchConfiguration;  
import org.springframework.data.elasticsearch.repository.config.EnableElasticsearchRepositories;  
import java.io.File;  
import java.io.FileInputStream;  
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;
@Configuration
@EnableElasticsearchRepositories(basePackages = "com.xxx.repository")
@ComponentScan(basePackages = "com.xxx")
public class Config extends AbstractElasticsearchConfiguration {
    @Value("${elasticsearch.url}")
    public String elasticsearchUrl;
    @Value("${elasticsearch.username}")
    public String elasticsearchUsername;
    @Value("${elasticsearch.password}")
    public String elasticsearchPassword;
    @Override
    @Bean
    public RestHighLevelClient elasticsearchClient() {
        SSLContext sc = null;
        try {
            TrustManager[] tm = {new MyX509TrustManager(certFilePath, certPassword)};
            sc = SSLContext.getInstance("SSL", "SunJSSE");
            sc.init(null, tm, new SecureRandom());
        } catch (Exception e) {
            e.printStackTrace();
        }
        final ClientConfiguration clientConfiguration = ClientConfiguration.builder()
            .connectedTo(StringHostParse(elasticsearchUrl))
            .usingSsl(sc, new NullHostNameVerifier())
            .withBasicAuth(elasticsearchUsername, elasticsearchPassword)
            .build();
        return RestClients.create(clientConfiguration).rest();
    }

    private String[] StringHostParse(String hostAndPorts) {
        return hostAndPorts.split(",");
    }

    public static class MyX509TrustManager implements X509TrustManager {
        X509TrustManager sunJSSEX509TrustManager;
        MyX509TrustManager(String certFilePath, String certPassword) throws Exception {
            File file = new File(certFilePath);
            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, certPassword.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;  
    }  
}  
}
```

In the preceding command, *certFilePath* and *certPassword* indicate the path and password of the .jks certificate, respectively.

Obtaining and Uploading a Security Certificate

To access a security-mode Elasticsearch cluster that uses HTTPS, a security certificate must be loaded. Perform the following steps to obtain the security certificate and upload it to the client:

1. Obtain the security certificate **CloudSearchService.cer**.
 - a. Log in to the [CSS management console](#).
 - b. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
 - c. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
 - d. Click the **Overview** tab. In the **Configuration** area, click **Download Certificate** next to **HTTPS Access**.

Figure 2-8 Downloading a security certificate



2. Convert the security certificate **CloudSearchService.cer**. Upload the downloaded security certificate to the client and use keytool to convert the .cer certificate into a .jks certificate that can be read by Java.
 - In Linux, run the following command to convert the certificate:
`keytool -import -alias newname -keystore ./truststore.jks -file ./CloudSearchService.cer`
 - In Windows, run the following command to convert the certificate:
`keytool -import -alias newname -keystore .\truststore.jks -file .\CloudSearchService.cer`

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

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

2.5.6 Accessing an Elasticsearch Cluster Using Python

You can access an Elasticsearch cluster created in CSS using Python.

Prerequisites

- The CSS cluster is available.
- Ensure that the server running Python can communicate with the CSS cluster.

Accessing a Cluster

1. Install the Elasticsearch Python client. You are advised to use the client version that matches the Elasticsearch version. For example, if the cluster version is 7.6.2, install the Elasticsearch Python client 7.6.

```
pip install Elasticsearch==7.6
```

2. Create an Elasticsearch client and check whether the **test** index exists. The examples for clusters in different security modes are as follows:

- Cluster in non-security mode

```
from elasticsearch import Elasticsearch

class ElasticFactory(object):

    def __init__(self, host: list, port: str, username: str, password: str):
        self.port = port
        self.host = host
        self.username = username
        self.password = password

    def create(self) -> Elasticsearch:
        addrs = []
        for host in self.host:
            addr = {'host': host, 'port': self.port}
            addrs.append(addr)

        if self.username and self.password:
            elasticsearch = Elasticsearch(addrs, http_auth=(self.username, self.password))
        else:
            elasticsearch = Elasticsearch(addrs)
        return elasticsearch

es = ElasticFactory(["xxx.xxx.xxx.xxx"], "9200", None, None).create()
print(es.indices.exists(index='test'))
```

- Cluster in security mode + HTTP

```
from elasticsearch import Elasticsearch

class ElasticFactory(object):

    def __init__(self, host: list, port: str, username: str, password: str):
        self.port = port
        self.host = host
        self.username = username
        self.password = password

    def create(self) -> Elasticsearch:
        addrs = []
        for host in self.host:
            addr = {'host': host, 'port': self.port}
            addrs.append(addr)

        if self.username and self.password:
            elasticsearch = Elasticsearch(addrs, http_auth=(self.username, self.password))
        else:
            elasticsearch = Elasticsearch(addrs)
        return elasticsearch

es = ElasticFactory(["xxx.xxx.xxx.xxx"], "9200", "username", "password").create()
print(es.indices.exists(index='test'))
```

- Cluster in security mode + HTTPS

```
from elasticsearch import Elasticsearch
import ssl

class ElasticFactory(object):

    def __init__(self, host: list, port: str, username: str, password: str):
        self.port = port
```

```
self.host = host
self.username = username
self.password = password

def create(self) -> Elasticsearch:
    context = ssl_create_unverified_context()

    addrs = []
    for host in self.host:
        addr = {'host': host, 'port': self.port}
        addrs.append(addr)

    if self.username and self.password:
        elasticsearch = Elasticsearch(addrs, http_auth=(self.username, self.password),
        scheme="https", ssl_context=context)
    else:
        elasticsearch = Elasticsearch(addrs)
    return elasticsearch

es = ElasticFactory(["xxx.xxx.xxx.xxx"], "9200", "username", "password").create()
print(es.indices.exists(index='test'))
```

Table 2-46 Variables

Parameter	Description
host	IP address for accessing the Elasticsearch cluster. If there are multiple IP addresses, separate them with commas (,).
port	Access port of the Elasticsearch cluster. Enter 9200 .
username	Username for accessing the cluster.
password	Password of the user.

3. Create a cluster index through the Elasticsearch client.

```
mappings = {
    "settings": {
        "index": {
            "number_of_shards": number_of_shards,
            "number_of_replicas": 1,
        },
    },
    "mappings": {
        properties
    }
}
result = es.indices.create(index=index, body=mappings)
```

4. Query the index created in the previous step through the Elasticsearch client.

```
body = {
    "query": {
        "match": {
            "Query field": "Query content"
        }
    }
}
result = es.search(index=index, body=body)
```

2.5.7 Accessing an Elasticsearch Cluster Through the MRS Hive Client

The Elasticsearch-Hadoop (ES-Hadoop) connector combines the massive data storage and in-depth processing capabilities of Hadoop with the real-time search and analysis capabilities of Elasticsearch. It allows you to quickly get to know big data and work better in the Hadoop ecosystem.

This topic uses the ES-Hadoop of MRS as an example to describe how to connect to a CSS cluster. You can configure any other applications that need to use the Elasticsearch cluster. Ensure the network connection between the client and the Elasticsearch cluster is normal.

Prerequisites

- The CSS cluster is available.
- The client can communicate with the CSS cluster.
- The CSS and MRS clusters are in the same region, AZ, VPC, and subnet.

Accessing a Cluster

1. Obtain the private network address of the cluster. It is used to access the cluster.

This topic uses a private IP address as an example to describe how to access a cluster. The cluster access address varies with the network configurations used. For details, see [Network Configuration](#).

- a. Log in to the [CSS management console](#).
- b. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
- c. In the cluster list, obtain the target cluster's private IP address from the **Private IP Address** column. Generally, the IP address format is `<host>:<port>` or `<host>:<port>,<host>:<port>`.

If the cluster has only one node, the IP address and port number of this 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**.

2. Log in to an MRS cluster node. For details, see [Logging In to an ECS](#).
3. Run the cURL command on an MRS cluster node to check the network connectivity. Ensure every node in the MRS cluster can connect to the CSS cluster.

- Cluster in non-security mode
`curl -X GET http://<host>:<port>`
- Cluster in security mode + HTTP
`curl -X GET http://<host>:<port> -u <user>:<password>`
- Cluster in security mode + HTTPS
`curl -X GET https://<host>:<port> -u <user>:<password> -ik`

Table 2-47 Variables

Variable	Description
<host>	IP address of each node in the cluster. If the cluster contains multiple nodes, there will be multiple IP addresses. You can use any of them.
<port>	Port number for accessing a cluster node. Generally, the port number is 9200.
<user>	Username for accessing the cluster.
<password>	Password of the user. If the password contains special characters, enclose the username and password in single quotation marks, for example, <code>curl -u "user:password!" "http://<host>:<port>"</code> .

- Download the [ES-Hadoop lib package](#) and decompress it to obtain the **elasticsearch-hadoop-x.x.x.jar** file. The version must be the same as the CSS cluster version. For example, if the CSS cluster version is 7.6.2, you are advised to download **elasticsearch-hadoop-7.6.2.zip**.
- Download the HttpClient dependency package [commons-httpclient:commons-httpclient-3.1.jar](#). In the package name, 3.1 indicates the version number. Select the package of the version you need.
- Install the MRS client. If the MRS client has been installed, skip this step. For details, see [Installing a Client \(MRS 3.x or Later\)](#).
- Log in to the MRS client. Upload the JAR dependency packages of ES-Hadoop and HttpClient to the MRS client.
- Create an HDFS directory on the MRS client. Upload the ES-Hadoop lib package and the HttpClient dependency package to the directory.

```
hadoop fs -mkdir /tmp/hadoop-es
hadoop fs -put elasticsearch-hadoop-x.x.x.jar /tmp/hadoop-es
hadoop fs -put commons-httpclient-3.1.jar /tmp/hadoop-es
```
- Log in to the Hive client from the MRS client. For details, see [Using a Hive Client](#).
- On the Hive client, add the ES-Hadoop lib package and the HttpClient dependency package. This command is valid only for the current session. Enter **beeline** or **hive** to go to the execution page and run the following commands:

```
add jar hdfs:///tmp/hadoop-es/commons-httpclient-3.1.jar;
add jar hdfs:///tmp/hadoop-es/elasticsearch-hadoop-x.x.x.jar;
```
- On the Hive client, create a Hive foreign table.

```
- Cluster in non-security mode
CREATE EXTERNAL table IF NOT EXISTS student(
  id BIGINT,
  name STRING,
  addr STRING
)
STORED BY 'org.elasticsearch.hadoop.hive.EsStorageHandler'
TBLPROPERTIES(
  'es.nodes' = 'xxx.xxx.xxx.xxx:9200',
  'es.port' = '9200',
```

- ```
'es.net.ssl' = 'false',
'es.nodes.wan.only' = 'false',
'es.nodes.discovery'='false',
'es.input.use.sliced.partitions'='false',
'es.resource' = 'student/_doc'
);
```
- Cluster in security mode + HTTP

```
CREATE EXTERNAL table IF NOT EXISTS student(
 id BIGINT,
 name STRING,
 addr STRING
)
STORED BY 'org.elasticsearch.hadoop.hive.EsStorageHandler'
TBLPROPERTIES(
 'es.nodes' = 'xxx.xxx.xxx.xxx:9200',
 'es.port' = '9200',
 'es.net.ssl' = 'false',
 'es.nodes.wan.only' = 'false',
 'es.nodes.discovery'='false',
 'es.input.use.sliced.partitions'='false',
 'es.nodes.client.only'='true',
 'es.resource' = 'student/_doc',
 'es.net.http.auth.user' = 'username',
 'es.net.http.auth.pass' = 'password'
);
```
  - Cluster in security mode + HTTPS
    - i. **Obtaining and Uploading a Security Certificate.**
    - ii. Put the .jks file to the same path of each node in the MRS cluster, for example, /tmp. You can run the **scp** command to transfer the file. Ensure user **omm** has the permission to read the file. You can run the following command to set the permission:

```
chown -R omm truststore.jks
```
    - iii. Create a Hive foreign table.

```
CREATE EXTERNAL table IF NOT EXISTS student(
 id BIGINT,
 name STRING,
 addr STRING
)
STORED BY 'org.elasticsearch.hadoop.hive.EsStorageHandler'
TBLPROPERTIES(
 'es.nodes' = 'https://xxx.xxx.xxx.xxx:9200',
 'es.port' = '9200',
 'es.net.ssl' = 'true',
 'es.net.ssl.truststore.location' = 'certFilePath',
 'es.net.ssl.truststore.pass' = 'certPassword',
 'es.nodes.wan.only' = 'false',
 'es.nodes.discovery'='false',
 'es.nodes.client.only'='true',
 'es.input.use.sliced.partitions'='false',
 'es.resource' = 'student/_doc',
 'es.net.http.auth.user' = 'username',
 'es.net.http.auth.pass' = 'password'
);
```

**Table 2-48** ES-Hadoop parameters

| Parameter | Default Value | Description                                                                                           |
|-----------|---------------|-------------------------------------------------------------------------------------------------------|
| es.nodes  | localhost     | Address for accessing the CSS cluster. You can check the private network address in the cluster list. |

| Parameter                      | Default Value | Description                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| es.port                        | 9200          | Port number for accessing a cluster. Generally, the port number is 9200.                                                                                                                                                                                                                                                                                                                |
| es.nodes.wan.only              | false         | Whether to perform node sniffing.                                                                                                                                                                                                                                                                                                                                                       |
| es.nodes.discovery             | true          | Whether to disable node discovery.                                                                                                                                                                                                                                                                                                                                                      |
| es.input.use.sliced.partitions | true          | Whether to use slices. Its value can be: <ul style="list-style-type: none"><li>• <b>true</b></li><li>• <b>false</b></li></ul> <b>NOTE</b><br>If this parameter is set to <b>true</b> , the index prefetch time may be significantly prolonged, and may even be much longer than the data query time. You are advised to set this parameter to <b>false</b> to improve query efficiency. |
| es.resource                    | NA            | Specifies the index and type to be read and written.                                                                                                                                                                                                                                                                                                                                    |
| es.net.http.auth.user          | NA            | Username for accessing the cluster. Set this parameter only if the security mode is enabled.                                                                                                                                                                                                                                                                                            |
| es.net.http.auth.pass          | NA            | Password of the user. Set this parameter only if the security mode is enabled.                                                                                                                                                                                                                                                                                                          |
| es.net.ssl                     | false         | Whether to enable SSL. If SSL is enabled, you need to configure the security certificate information.                                                                                                                                                                                                                                                                                   |
| es.net.ssl.truststore.location | NA            | Path of the .jks certificate file, for example, <b>file:///tmp/truststore.jks</b> .                                                                                                                                                                                                                                                                                                     |
| es.nodes.client.only           | false         | Check whether the IP address of an independent Client node is configured for <b>es.nodes</b> (that is, whether the Client node is enabled during Elasticsearch cluster creation). If yes, change the value to <b>true</b> , or an error will be reported, indicating that the data node cannot be found.                                                                                |



| Parameter                  | Default Value | Description                            |
|----------------------------|---------------|----------------------------------------|
| es.net.ssl.truststore.pass | NA            | Password of the .jks certificate file. |

For details about ES-Hadoop configuration items, see the [official configuration description](#).

12. On the Hive client, insert data.

```
INSERT INTO TABLE student VALUES (1, "Lucy", "address1"), (2, "Lily", "address2");
```

13. On the Hive client, execute a query.

```
select * from student;
```

The query result is as follows:


```
+-----+-----+-----+
| student.id | student.name | student.addr |
+-----+-----+-----+
| 1 | Lucy | address1 |
| 2 | Lily | address2 |
+-----+-----+-----+
2 rows selected (0.116 seconds)
```

14. Log in to Kibana.

- Log in to the [CSS management console](#).
- In the navigation pane on the left, choose **Clusters > Elasticsearch**.
- In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
- In the left navigation pane, choose **Dev Tools**.

15. On the **Dev Tools** page of Kibana, run a query and check the result.

```
GET /student/_search
```

**Figure 2-9** Kibana query result

## Obtaining and Uploading a Security Certificate

To access a security-mode Elasticsearch cluster that uses HTTPS, a security certificate must be loaded. Perform the following steps to obtain the security certificate and upload it to the client:

1. Obtain the security certificate **CloudSearchService.cer**.
  - a. Log in to the **CSS management console**.
  - b. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
  - c. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  - d. Click the **Overview** tab. In the **Configuration** area, click **Download Certificate** next to **HTTPS Access**.

**Figure 2-10** Downloading a security certificate

2. Convert the security certificate **CloudSearchService.cer**. Upload the downloaded security certificate to the client and use keytool to convert the **.cer** certificate into a **.jks** certificate that can be read by Java.

- In Linux, run the following command to convert the certificate:  
`keytool -import -alias newname -keystore ./truststore.jks -file ./CloudSearchService.cer`
- In Windows, run the following command to convert the certificate:  
`keytool -import -alias newname -keystore .\truststore.jks -file .\CloudSearchService.cer`

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

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

## 2.5.8 Accessing an Elasticsearch Cluster Using Go

This section describes how to access a CSS cluster using Go.

### Preparations

- The CSS cluster is available.
- Ensure that the server running Go can communicate with the CSS cluster.
- Ensure that Go has been installed on the server. You can download Go from the official website: <https://go.dev/dl/>.

### Connecting to a Non-Security Mode Cluster

Connect to a non-security mode cluster. The sample code is as follows:

```
package main

import (
 "github.com/elastic/go-elasticsearch/v7"
 "log"
)

func main() {
 cfg := elasticsearch.Config{
 Addresses: []string{
 "http://HOST:9200/",
 },
 }

 es, _ := elasticsearch.NewClient(cfg)
 log.Println(es.Info())
}
```

In the information above, **HOST** indicates the internal IP address of a cluster node.

### Connecting to a Security-Mode Cluster

- Connect to a security-mode cluster with HTTPS disabled. The sample code is as follows:

```
package main

import (
 "github.com/elastic/go-elasticsearch/v7"
 "log"
)
```

```
func main() {
 cfg := elasticsearch.Config{
 Addresses: []string{
 "http://HOST:9200/",
 },
 Username: "USERNAME",
 Password: "PASSWORD",
 }

 es, _ := elasticsearch.NewClient(cfg)
 log.Println(es.Info())
}
```

- Connect to a security cluster that has enabled HTTPS and does not use certificates. The sample code is as follows:

```
package main

import (
 "crypto/tls"
 "github.com/elastic/go-elasticsearch/v7"
 "log"
 "net/http"
)

func main() {
 cfg := elasticsearch.Config{
 Addresses: []string{
 "https://HOST:9200/",
 },
 Username: "USERNAME",
 Password: "PASSWORD",
 Transport: &http.Transport{
 TLSClientConfig: &tls.Config{
 InsecureSkipVerify: true,
 },
 },
 }

 es, _ := elasticsearch.NewClient(cfg)
 log.Println(es.Info())
}
```

- Connect to a security cluster that has enabled HTTPS and uses certificates. The sample code is as follows:

```
package main

import (
 "crypto/tls"
 "crypto/x509"
 "flag"
 "github.com/elastic/go-elasticsearch/v7"
 "io/ioutil"
 "log"
 "net"
 "net/http"
 "time"
)

func main() {
 insecure := flag.Bool("insecure-ssl", false, "Accept/Ignore all server SSL certificates")
 flag.Parse()

 // Get the SystemCertPool, continue with an empty pool on error
 rootCAs, _ := x509.SystemCertPool()
 if rootCAs == nil {
 rootCAs = x509.NewCertPool()
 }

 // Read in the cert file
 certs, err := ioutil.ReadFile("/tmp/CloudSearchService.cer")
```

```
if err != nil {
 log.Fatalf("Failed to append %q to RootCAs: %v", "xxx", err)
}

// Append our cert to the system pool
if ok := rootCAs.AppendCertsFromPEM(certs); !ok {
 log.Println("No certs appended, using system certs only")
}

config := elasticsearch.Config{
 Addresses: []string{
 "https://HOST:9200/",
 },
 Username: "USERNAME",
 Password: "PASSWORD",
 Transport: &http.Transport{
 MaxIdleConnsPerHost: 10,
 ResponseHeaderTimeout: time.Second,
 DialContext: (&net.Dialer{
 Timeout: 30 * time.Second,
 KeepAlive: 30 * time.Second,
 }).DialContext,
 TLSClientConfig: &tls.Config{
 InsecureSkipVerify: *insecure,
 RootCAs: rootCAs,
 },
 },
}
es, _ := elasticsearch.NewClient(config)
log.Println(elasticsearch.Version)
log.Println(es.Info())
}
```

**Table 2-49** Variables

| Parameter | Description                                                                                                            |
|-----------|------------------------------------------------------------------------------------------------------------------------|
| HOST      | IP address for accessing the Elasticsearch cluster. If there are multiple IP addresses, separate them with commas (,). |
| USERNAME  | Username for accessing the cluster.                                                                                    |
| PASSWORD  | Password of the user.                                                                                                  |

## Running Code

Write the code above to the **EsTest.gc** file based on the cluster type and save the file to an independent directory. Run the following command in that directory to run the code:

```
go env -w GO111MODULE=on
go env -w GOPROXY=https://repo.huaweicloud.com/repository/goproxy/
go env -w GONOSUMDB=*

go mod init test
go mod tidy
go run EsTest.go
```

## 2.5.9 Accessing an Elasticsearch Cluster Using LDAP

The Light Directory Access Protocol (LDAP) is a lightweight version of the directory access protocol based on the X.500 standard. An LDAP service provides user authentication and authorization. Using the Security plugin for Open Distro for Elasticsearch, CSS adds Active Directory as an authentication backend for clusters, connecting them seamlessly to the LDAP service. This topic describes the steps needed to connect a CSS cluster to an LDAP service.

It also describes how to enable LDAP authentication for a CSS cluster to allow access by LDAP users of specific roles.

### Constraints

Only Elasticsearch 7.10.2 security-mode clusters can be accessed through LDAP.

### Preparations

- A security-mode Elasticsearch cluster has been created in CSS and its status is available.
- The LDAP service that is in the same VPC as the Elasticsearch cluster and the necessary user data have been prepared. For details, see [the OpenLDAP document: A Quick-Start Guide](#).


### Accessing a Cluster

- Step 1** Install an LDAP service on an ECS. If the LDAP service and user data have already been prepared, skip this step.
1. Create an ECS. The ECS must run a Windows OS and must be in the same VPC and security group as the security-mode Elasticsearch cluster of CSS. The Windows Server running on the ECS provides the built-in Active Directory service that supports the LDAP protocol.  
For how to create an ECS, see [Purchasing an ECS in Custom Config Mode](#).
  2. Log in to the ECS, and enable the Active Directory service. Create a domain, administrator, users, and user groups.
- Step 2** Modify the parameter settings of the security-mode Elasticsearch cluster on CSS. Configure a static parameter in `elasticsearch.yml` to connect the cluster to the LDAP service.
1. Log in to the [CSS management console](#).
  2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
  3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  4. Choose **Cluster Settings > Parameter Settings**, and click **Edit**. Expand **Custom**, and add the following parameter and value.
    - **Parameter:**  
opendistro\_security.unsupported.restapi.allow\_securityconfig\_modification
    - **Value:** true
  5. Then click **Submit**. In the displayed dialog box, confirm the settings, select the box indicating "I understand that the modification will take effect after the cluster is restarted." and click **Yes**.

If **Status** is **Succeeded** in the parameter change list, the change has been saved. Up to 20 change records can be displayed.

6. After saving the change, click **Restart** in the upper-right corner of the page to restart the cluster, thus making the change take effect.

**Figure 2-11** Restarting a cluster



- You need to restart the cluster after the change, or **Configuration not updated** will be displayed in the **Task Status** column in the cluster list.
- If you restart the cluster after the change, and **Task Status** displays **Configuration error** in the cluster list, the parameter configuration file has failed to be modified.

**Step 3** Configure a route for an Elasticsearch cluster on the CSS console to connect the cluster to the LDAP service.

1. On the cluster information page, click the **Overview** tab.
2. In the **Configuration** area, click **Add Route** next to **Cluster Route**.
3. In the displayed dialog box, configure the route information.

**Table 2-50** Adding a route

| Parameter   | Description                                                                                                         |
|-------------|---------------------------------------------------------------------------------------------------------------------|
| IP Address  | Enter the IP address of the LDAP server. If the LDAP service on the ECS is used, enter the IP address of the ECS.   |
| Subnet Mask | Enter the subnet mask of the LDAP server. If the LDAP service on the ECS is used, enter the subnet mask of the ECS. |

4. Click OK.

**Step 4** Configure LDAP authentication for a security-mode Elasticsearch cluster.

1. On the cluster information page, click **Kibana** in the upper-right corner to log in to Kibana.
2. In the navigation pane on the left, choose **Dev Tools**.
3. Run the following commands to configure LDAP authentication.

 NOTE

- Concepts used in an X.500 directory access protocol (including LDAP):

- CN = Common Name
- OU = Organizational Unit
- DC = Domain Component
- DN = Distinguished Name

The CN, OU, and DC must be provided in the correct order. Otherwise, authentication will fail.

- The configuration file consists of two parts: **authc** and **authz**.
  - **authc** (authentication): verifies whether a user is truly who they claim they are (password verification).
  - **authz** (authorization): verifies what the current user has access to.

```
PUT _opendistro/_security/api/securityconfig/config
{
 "dynamic": {
 "authc": {
 "basic_internal_auth_domain": {
 "description": "Authenticate via HTTP Basic against internal users database",
 "http_enabled": true,
 "transport_enabled": true,
 "order": 1,
 "http_authenticator": {
 "type": "basic",
 "challenge": true
 },
 "authentication_backend": {
 "type": "intern"
 }
 },
 "ldap": {
 "description": "Authenticate via LDAP or Active Directory",
 "http_enabled": true,
 "transport_enabled": true,
 "order": 2,
 "http_authenticator": {
 "type": "basic",
 "challenge": false
 },
 "authentication_backend": {
 "type": "ldap",
 "config": {
 "enable_ssl": false,
 "enable_start_tls": false,
 "enable_ssl_client_auth": false,
 "verify_hostnames": true,
 "hosts": ["10.0.XXX.XXX:389"],
 "bind_dn": "CN=adminAD,DC=test,DC=ldap,DC=com",
 "password": "<password>",
 "userbase": "OU=ITDepartment,DC=test,DC=ldap,DC=com",
 "usersearch": "(sAMAccountName={0})",
 "username_attribute": "uid"
 }
 }
 }
 },
 "authz": {
 "roles_from_myldap": {
 "description": "Authorize via LDAP or Active Directory",
```



```
"http_enabled": true,
"transport_enabled": true,
"authorization_backend": {
 "type": "ldap",
 "config": {
 "enable_ssl": false,
 "enable_start_tls": false,
 "enable_ssl_client_auth": false,
 "verify_hostnames": true,
 "hosts": ["10.0.XXX.XXX:389"],
 "bind_dn": "CN=adminAD,DC=test,DC=ldap,DC=com",
 "password": "<password>",
 "rolebase": "OU=groups,DC=test,DC=ldap,DC=com",
 "rolesearch": "(member={0})",
 "userroleattribute": null,
 "userrolename": "disabled",
 "rolename": "CN",
 "resolve_nested_roles": true,
 "userbase": "OU=ITDepartment,DC=test,DC=ldap,DC=com",
 "usersearch": "(uid={0})"
 }
}
```


The parameters in [Table 2-51](#) need to be modified based on the actual environment.

**Table 2-51** Parameter description

| Parameter | Description                                                                                                                                                                   |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| hosts     | Address of the LDAP service. The port number is 389. If the LDAP service on the ECS is used, enter the IP address of the ECS.                                                 |
| bind_dn   | It is similar to the LDAP user name (CN - OU - DC) and is used to access the LDAP server. Select a user name from the user data of the LDAP service.                          |
| password  | Password of the LDAP user configured using <b>bind_dn</b> .                                                                                                                   |
| userbase  | After the LDAP service is connected, the DN that the user belongs to is obtained. In this example, all user information in the <b>ITDepartment</b> directory is synchronized. |
| rolebase  | The collection of permissions that can be configured for the <b>userbase</b> user group of the LDAP service.                                                                  |

**Step 5** Configure the mapping between LDAP user permissions and Elasticsearch permissions in the Elasticsearch security-mode cluster to enable fine-grained access control.

The rolebase permissions group of the LDAP server must be mapped to the roles in the Elasticsearch cluster. [Figure 2-12](#) illustrates the mapping. For details about the configuration, see [Creating Users for an Elasticsearch Cluster and Granting Cluster Access](#).

**Figure 2-12** Permissions mapping

1. On the Kibana console, expand the menu in the upper-left corner, and choose **Security**. The **Security** page is displayed.
2. Click **Roles** to go to the Open Distro Security Roles page. Click **Create Role**, set **Name**, **Cluster Permissions**, **Index Permissions**, and **Tenant Permissions**. Then click **Save Role Definition** to save the role settings. The parameters are as follows:
  - Name (name of the role)
  - Cluster permissions
  - Index permissions
  - Tenant permissions
3. Click the newly created role, select **Mapped users**, enter a permissions group of the LDAP service in **Backend roles**, and click **Map**.
4. Check the configuration result.

**Figure 2-13** Permissions mapping

|                                      |                          |                   |                |                 |                               |               |
|--------------------------------------|--------------------------|-------------------|----------------|-----------------|-------------------------------|---------------|
| Q reserved:false                     | Cluster permissions      | Index permissions | Internal users | Backend roles   | Tenants                       | Customization |
| <input type="checkbox"/> Role        | Cluster permissions      | Index permissions | Internal users | Backend roles   | Tenants                       | Customization |
| <input type="checkbox"/> role_google | cluster_composite_ops_ro | index*            | —              | groupGoogleInAD | global_tenant<br>admin_tenant | Custom        |

5. Repeat **Step 5.2** to **Step 5.4** to map other permissions groups.

**Step 6** Log in to Kibana using the LDAP user to verify the configuration.

1. Log in to the **CSS management console**.
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column. Use the LDAP user to log in to the Kibana console.  
If the login is successful, the configuration is successful, and users can access the Elasticsearch cluster through LDAP. The specific permissions authorized are controlled by role permissions configured in Elasticsearch.

-----End

## 2.6 Importing Data to an Elasticsearch Cluster

## 2.6.1 Different Ways to Ingest Data into an Elasticsearch Cluster

### Introduction

Elasticsearch clusters support multiple data ingestion methods, as listed in [Table 2-52](#). Select one that fits your needs the best. Before starting to ingest data, determine whether to enhance the data ingestion performance of Elasticsearch clusters first. For details, see [Enhancing the Data Ingestion Performance of Elasticsearch Clusters](#).

**Table 2-52** Different ways to ingest data into an Elasticsearch cluster

| Data Ingestion Method         | Scenario                                                                                                                                                                                                                        | Supported Data Formats/Sources                                                                      | Details                                                                              |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| CSS Logstash                  | Use CSS Logstash to ingest data from multiple sources, such as relational databases, a Kafka service, and OBS, into Elasticsearch clusters.                                                                                     | <ul style="list-style-type: none"><li>• RDS MySQL</li><li>• Kafka data</li><li>• OBS data</li></ul> | <a href="#">Using Logstash to Synchronize Data to Elasticsearch</a>                  |
| Open-source Logstash          | Open-source Logstash offers a server-side, real-time data processing pipeline, which supports data ingestion from multiple sources. It can be used to ingest various types of data, such as logs, monitoring data, and metrics. | JSON, CSV, and text                                                                                 | <a href="#">Using In-house Built Logstash to Import Data to Elasticsearch</a>        |
| Open-source Elasticsearch API | Open-source Elasticsearch APIs can be used to ingest data. This method is flexible, as you can write your own application code.                                                                                                 | JSON                                                                                                | <a href="#">Using Open Source Elasticsearch APIs to Import Data to Elasticsearch</a> |
| Cloud Data Migration (CDM)    | You can use CDM for batch data migration. For example, if data is stored in OBS or an Oracle database, CDM is recommended.                                                                                                      | JSON                                                                                                | <a href="#">Using CDM to Import Data to Elasticsearch</a>                            |

| Data Ingestion Method          | Scenario                                                                          | Supported Data Formats/Sources    | Details                                                                     |
|--------------------------------|-----------------------------------------------------------------------------------|-----------------------------------|-----------------------------------------------------------------------------|
| Data Replication Service (DRS) | DRS can be used for online database migration and real-time data synchronization. | Relational Database Service (RDS) | <a href="#">Using DRS to Ingest Data from a Database into Elasticsearch</a> |

## Using DRS to Ingest Data from a Database into Elasticsearch

DRS is an easy-to-use, stable, and efficient cloud service for online database migration and real-time database synchronization. Real-time data synchronization refers to the real-time replication of data from one source to another while ensuring data consistency.

DRS can be used to ingest data from multiple types of relational databases to Elasticsearch clusters. For details about the supported software versions for source databases and destination clusters, see [Table 2-53](#).

**Table 2-53** Supported versions for using DRS to ingest data from a database to Elasticsearch

| Scenario                                                                            | Source DB                                                             | Destination Elasticsearch Cluster                   | Details                                                                        |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------|-----------------------------------------------------|--------------------------------------------------------------------------------|
| Ingesting data from an RDS for MySQL database to a CSS Elasticsearch cluster        | RDS for MySQL 5.5, 5.6, 5.7, or 8.0                                   | Elasticsearch 5.5, 6.2, 6.5, 7.1, 7.6, 7.9, or 7.10 | <a href="#">(Real-Time Synchronization) From MySQL to CSS/ES</a>               |
| Ingesting data from a TaurusDB database to a CSS Elasticsearch cluster              | Primary/standby TaurusDB instances                                    | Elasticsearch 5.5, 6.2, 6.5, 7.1, 7.6, 7.9, or 7.10 | <a href="#">(Real-Time Synchronization) From TaurusDB to CSS/Elasticsearch</a> |
| Ingesting data from an in-house built MySQL database to a CSS Elasticsearch cluster | MySQL database 5.5, 5.6, 5.7, or 8.0 created on a local server or ECS | Elasticsearch 5.5, 6.2, 6.5, 7.1, 7.6, 7.9, or 7.10 | <a href="#">(Real-Time Synchronization) From MySQL to CSS/ES</a>               |

## 2.6.2 Using In-house Built Logstash to Import Data to Elasticsearch

With CSS, you can use in-house developed Logstash to ingest data into Elasticsearch for efficient search and exploration. Data files can be in the JSON or CSV format.

Logstash is an open-source, server-side data processing pipeline that ingests data from multiple sources simultaneously, processes and transforms the data, 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).
- Logstash must use an OSS version that is consistent with that of the CSS cluster. To download Logstash, visit the following website: <https://www.elastic.co/downloads/logstash-oss>
- 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.
- 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>
- 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 2-14** illustrates how data is imported when Logstash is deployed on an external network.

**Figure 2-14** 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 configure port mapping to forward requests sent to the opened port on the jump host to the destination 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>
```

  - 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
```

4. 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/** (create the access\_log folder if it does not already exist), 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     |

5. On 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
```

6. Enter the following content in **logstash-simple.conf**:

```
input {
 Location of data
}
filter {
 Related data processing
}
output {
 elasticsearch {
 hosts => "<EIP 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 *<EIP address of the jump host>* with the IP address (with **EIP**) 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** to try and rectify the error.

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

- a. **Obtaining the Security Certificate.**
- b. Store the downloaded 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" # Username for accessing the security-mode cluster
 password => "*****" # Password for accessing the security-mode cluster
 cacert => "/logstash/logstash6.8/config/CloudSearchService.cer"
 manager_template => false
 ilm_enabled => false
 ssl => true
 ssl_certificate_verification => false
 }
}
```

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

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

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

8. Log in to the **CSS management console**.
9. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
10. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
11. In the left navigation pane, choose **Dev Tools**.
12. On the Kibana console, search for the ingested data.

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 2-15** 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 2-15** 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.
  - If there are multiple servers in a VPC, you only need to associate an EIP with one of these servers. Switch to the node where Logstash is deployed from the node with which the EIP is associated.
  - If a private line or VPN is available, there is no need for an EIP.

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     |

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

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

Enter 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"
 }
}
```

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

- Obtaining the Security Certificate.**
- Store the downloaded 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 step 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 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" # Username for accessing the security-mode cluster
 password => "*****" # Password for accessing the security-mode cluster
 cacert => "/logstash/logstash6.8/config/CloudSearchService.cer"
 manager_template => false
 ilm_enabled => false
 ssl => true
 ssl_certificate_verification => false
 }
}
```

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**.
7. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
8. In the left navigation pane, choose **Dev Tools**.
9. On the Kibana console, search for the ingested data.

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
```

## Obtaining the Security Certificate

To access a security-mode Elasticsearch cluster that uses HTTPS, a security certificate must be loaded. To obtain this security certificate (CloudSearchService.cer), follow these steps:

1. Log in to the **CSS management console**.
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.

4. Click the **Overview** tab. In the **Configuration** area, click **Download Certificate** next to **HTTPS Access**.

**Figure 2-16** Downloading a security certificate



## 2.6.3 Using Open Source Elasticsearch APIs to Import Data to Elasticsearch

With CSS, you can use open-source Elasticsearch APIs on Kibana or an ECS server to import data to an Elasticsearch cluster. JSON files are supported.

- **Using an Open-Source Elasticsearch API to Import Data on Kibana:** Run POST commands to import data.
- **Using an Open-Source Elasticsearch API to Import Data on an ECS Server:** Run cURL commands to import data.

### Using an Open-Source Elasticsearch API to Import Data on Kibana

On Kibana, you can run POST commands to import single pieces of data using an open-source Elasticsearch API.

1. Log in to Kibana.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
  - c. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
  - d. In the left navigation pane, choose **Dev Tools**.
2. Run the following command on the console to check whether the cluster has indexes:

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

  - If indexes are available in the cluster to which you want to import data, you do not need to create an index. Go to [7](#).
  - If no indexes are available in the cluster, go to the next step to create an index.
3. Run the following command to create an index for storing imported data and create a custom mapping to define the data type.

For example, run the following command to create index **my\_store**:

Sample code for Elasticsearch 7.x or later:

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

```

 },
 "size": {
 "type": "keyword"
 }
 }
}

```

Sample code for Elasticsearch earlier than 7.x:

```

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

```

4. Run commands on Kibana to import data. In the example below, only a single piece of data is imported:

Sample code for Elasticsearch 7.x or later:

```

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

```

Sample code for Elasticsearch earlier than 7.x:

```

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

```

The command output is similar to that shown in [Figure 2-17](#). If the value of the **errors** field in the result is **false**, the data is successfully imported.

Figure 2-17 Response message

```
1 {
2 "took": 42,
3 "errors": false,
4 "items": [
5 {
6 "index": {
7 "_index": "my_store",
8 "_type": "products",
9 "_id": "AWTGBt7BwpN-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 }
```

## Using an Open-Source Elasticsearch API to Import Data on an ECS Server

On an ECS server, you can run cURL commands to use an open-source Elasticsearch API to import JSON files.

In the example below, a cluster in non-security mode is used to describe how to import data using cURL commands. For the commands for a security cluster, see [Accessing an Elasticsearch Cluster Using cURL Commands](#).

1. Obtain the private network address of the cluster.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
  - c. In the cluster list, obtain the target cluster's private IP address from the **Private IP Address** column. Generally, the IP address format is `<host>:<port>` or `<host>:<port>,<host>:<port>`.

If the cluster has only one node, the IP address and port number of this 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**.
2. Log in to the ECS that you are going to use to import data.

For how to buy and use an ECS, see [Purchasing and Logging In to a Linux ECS](#).
3. Upload a JSON file to the ECS.

For example, save the following data as a JSON file and upload the file to the ECS:

For Elasticsearch 7.x or later:

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

For Elasticsearch 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"}
```

4. Run the following command in the path where the JSON file is stored in the ECS to import the JSON file to an Elasticsearch cluster.

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

Replace **{Private network address and port number of the node}** with the private network address and port number of a node in the cluster. 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. **test.json** indicates the JSON file to be ingested.

#### NOTE

In the case of the failure of a cluster node, if the cluster contains multiple nodes, you can replace **{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 this node and execute the command again.

**Example:** In this example, assume that you need to import data in the **test.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.

- a. Run the following command to create an index named **my\_store**:

Sample code for Elasticsearch 7.x or later:

```
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"
 }
 }
 }
}'
```

Sample code for Elasticsearch 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"
}
}
}
}
```

- b. Run the following command to import the **test.json** file:

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

In this case, if the following information is displayed, the data is successfully imported:

```
{"took":204,"errors":false,"items":[{"index":
{"_index":"my_store","_type":"_doc","_id":"DJQkBlwBbJvUd2769Wi-","_version":1,"result":"created",
"_shards":{"total":2,"successful":1,"failed":0},"_seq_no":0,"_primary_term":1,"status":201}},
{"index":
{"_index":"my_store","_type":"_doc","_id":"DZQkBlwBbJvUd2769Wi-","_version":1,"result":"create
d","_shards":{"total":2,"successful":1,"failed":0},"_seq_no":1,"_primary_term":1,"status":201}}}]}
```

## 2.6.4 Using CDM to Import Data to Elasticsearch

With CSS, you can use the web user interface of the Cloud Data Migration (CDM) service to import data from an Oracle database or OBS to an Elasticsearch cluster. Only JSON files are supported.

**Table 2-54** Using CDM to import data to CSS

| Scenario                                      | Source Data                            | Target Cluster                                      |
|-----------------------------------------------|----------------------------------------|-----------------------------------------------------|
| Importing data from an Oracle database to CSS | A local or third-party Oracle database | Elasticsearch 5.5, 6.2, 6.5, 7.1, 7.6, 7.9, or 7.10 |
| Importing data from OBS to CSS                | JSON files in OBS buckets              | Elasticsearch 5.5, 6.2, 6.5, 7.1, 7.6, 7.9, or 7.10 |

### Preparations

1. Confirm the source data.

For example, the source data can be the following JSON files:

For Elasticsearch 7.x or later:

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

For Elasticsearch 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"}
```

2. Obtain data source information.

- If the data source is an Oracle database, you need to obtain the IP address, database name, username, and password of the database.
- If the data source is JSON files in an OBS bucket, you need to obtain its domain name, port, AK, and SK.



3. If the data source is an Oracle database, make sure the Oracle database can be accessed using a public IP address, or a VPN or Direct Connect connection has been established between the customer's data center and Huawei Cloud.

## Importing Data

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, obtain the target cluster's private IP address from the **Private IP Address** column. Generally, the IP address format is `<host>:<port>` or `<host>:<port>,<host>:<port>`.  
If the cluster has only one node, the IP address and port number of this 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**.
4. In the cluster list, locate the destination cluster, and click **Access Kibana** in the **Operation** column to log in to the Kibana console.
5. In the Kibana navigation pane on the left, choose **Dev Tools**.
6. Run the following command on the console to check whether the cluster has indexes:  
GET \_cat/indices?v
  - If indexes are available in the cluster to which you want to import data, you do not need to create an index. Go to [8](#).
  - If no indexes are available in the cluster, go to the next step to create an index.
7. Run the following command to create an index for storing imported data and create a custom mapping to define the data type.

For example, run the following command to create index **demo**:

For Elasticsearch 7.x or later:

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

For Elasticsearch earlier than 7.x:

```
PUT /demo
{
 "settings": {
 "number_of_shards": 1
 },
 "mappings": {
 "products": {
 "properties": {
 "productName": {
```

```
{
 "type": "text",
 "analyzer": "ik_smart"
},
"size": {
 "type": "keyword"
}
}
}
```

The command is successfully executed if the following information is displayed.

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

8. Use CDM to import data from Oracle or OBS to the Elasticsearch cluster.
  - If the data source is an Oracle database, see [Migrating Data from Oracle to CSS](#).
  - If the source data is JSON files in an OBS bucket, see [Migrating Data from OBS to CSS](#).
9. After the data migration is complete, go to the Kibana console of the Elasticsearch cluster again, and search for the imported data.

Run the following command to search for data. If the imported data is consistent with the source data, data importing is successful.

```
GET demo/_search
```

**demo** is the name of the created index. Replace it if another index is used.

The command is successfully executed if the following information is displayed.

```
{
 "took": 18,
 "timed_out": false,
 "_shards": {
 "total": 1,
 "successful": 1,
 "skipped": 0,
 "failed": 0
 },
 "hits": {
 "total": 2,
 "max_score": 1,
 "hits": [
 {
 "_index": "demo",
 "_type": "products",
 "_id": "g6UepnEBuvdFwWkRmn4V",
 "_score": 1,
 "_source": {
 "size": "L",
 "productName": "Latest art shirts for women in autumn 2019"
 }
 },
 {
 "_index": "demo",
 "_type": "products",
 "_id": "hKUepnEBuvdFwWkRmn4V",
 "_score": 1,
 "_source": {
 "size": "M",
 "productName": "Latest art shirts for women in autumn 2019"
 }
 }
]
 }
}
```

```
}
}
]
}
}
```

## 2.6.5 Enhancing the Data Ingestion Performance of Elasticsearch Clusters

### Overview

This unique CSS feature significantly enhances data ingestion performance and reduces write rejections through bulk route optimization, text index acceleration, and word segmentation acceleration. You are advised to use this feature for clusters that contain a large number of index shards and text indexes, or have a high inbound data throughput.

**Table 2-55** Ways to improve data ingestion performance

| Method                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Details                                       |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|
| Bulk route optimization       | <p>According to the default routing rule of Elasticsearch, each piece of data in a bulk request is routed to a different shard. When a large amount of data is written in and a large number of index shards exist, excessive internal request forwarding may trigger bulk rejection. Additionally, in a large-scale cluster, the long tail effect causes a high latency for bulk requests.</p> <p>You can specify the <b>index.bulk_routing</b> configuration item to enable bulk route optimization. You can use it to reduce the number of requests that need to be internally forwarded. For a cluster that has a large number of shards, this setting can improve write performance and reduce bulk rejection.</p> | <a href="#">Bulk Route Optimization</a>       |
| Bulk aggregation optimization | <p>You can specify the <b>index.aggr_perf_batch_size</b> configuration item to enable or disable bulk aggregation optimization. When bulk aggregation optimization is enabled, documents in a bulk request are written in batches, rather than individually. This helps to reduce the overheads of memory requests, lock requests, and other calls, enhancing data ingestion performance.</p>                                                                                                                                                                                                                                                                                                                           | <a href="#">Bulk Aggregation Optimization</a> |

| Method                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Details                                       |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|
| Text index acceleration       | You can configure <b>index.native_speed_up</b> to enable or disable text index acceleration. This setting optimizes the indexing process and memory access to accelerate index building for text fields (text and keyword). When text index acceleration is enabled, you can configure <b>index.native_analyzer</b> to enable word segmentation acceleration as well. For texts that need common word segmentation, you can use the analyzer to accelerate word segmentation. | <a href="#">Text Index Acceleration</a>       |
| Index merge task optimization | Enabling the three optimization options above will increase the number of index merge tasks in the cluster. You can configure <b>index.merge.scheduler.max_thread_count</b> to reduce the impact of index merge tasks on data ingestion performance. The optimization is about increasing the number of shard merging threads and thereby alleviate the throttling of data import.                                                                                            | <a href="#">Index Merge Task Optimization</a> |

## Constraints

Only Elasticsearch 7.10.2 clusters support data ingestion performance enhancement.

## Prerequisites

The target cluster is available.

## Logging In to Kibana

Log in to Kibana and go to the command execution page. Elasticsearch clusters support multiple access methods. This topic uses Kibana as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
4. In the left navigation pane, choose **Dev Tools**.

## Bulk Route Optimization

You can specify the **index.bulk\_routing** configuration item to enable bulk route optimization. You can use it to reduce the number of requests that need to be internally forwarded. For a cluster that has a large number of shards, this setting can improve write performance and reduce bulk rejection.

 **CAUTION**

After bulk route optimization is enabled (that is, **index.bulk\_routing** is set to **pack** or **local\_pack**), data writes are no longer routed based on **\_id**, and routing-related functionality may be affected. For example, ID-based GET requests may fail.

Run the following command to enable bulk route optimization:

```
PUT my_index
{
 "settings": {
 "index.bulk_routing": "local_pack"
 }
}
```

The options for the **index.bulk\_routing** configuration item are as follows:

- **default**: Use the default routing mechanism. Records in a bulk request are split and then routed separately.
- **pack**: All data of a single bulk request is randomly routed to the same shard.
- **local\_pack**: The data of a single bulk request is routed to a local shard of the data node that has received the bulk request. If the node does not contain the corresponding index shard, the data is randomly routed to another node that contains the index shard. This solution depends on the random distribution of client bulk requests and the balanced distribution of primary shards.

## Bulk Aggregation Optimization

You can specify the **index.aggr\_perf\_batch\_size** configuration item to enable or disable bulk aggregation optimization. When bulk aggregation optimization is enabled, documents in a bulk request are written in batches, rather than individually. This helps to reduce the overheads of memory requests, lock requests, and other calls, enhancing data ingestion performance.

Run the following command to enable bulk aggregation optimization:

```
PUT my_index
{
 "settings": {
 "index.aggr_perf_batch_size": "128"
 }
}
```

The value range of the **index.aggr\_perf\_batch\_size** configuration item is [1, Integer.MAX\_VALUE]. The default value is 1, indicating that bulk aggregation optimization is disabled. If the value is greater than 1, bulk aggregation optimization is enabled, the value of **MIN(bulk\_doc\_size, aggr\_perf\_batch\_size)** indicates the bulk size.

## Text Index Acceleration

You can configure **index.native\_speed\_up** to enable or disable text index acceleration. This setting optimizes the indexing process and memory access to accelerate index building for text fields (text and keyword). When text index acceleration is enabled, you can configure **index.native\_analyzer** to enable word segmentation acceleration as well. For texts that need common word segmentation, you can use the analyzer to accelerate word segmentation.

 **CAUTION**

- Word segmentation acceleration can be enabled (set **index.native\_analyzer** to **true**) only when text index acceleration is enabled (set **index.native\_speed\_up** to **true**). Otherwise, word segmentation acceleration will not take effect.
- Text index acceleration cannot be enabled for indexes that contain nested fields.

Run the following command to enable text index acceleration:

```
PUT my_index
{
 "settings": {
 "index.native_speed_up": true,
 "index.native_analyzer": true
 }
}
```

The values of **index.native\_speed\_up** and **index.native\_analyzer** are **true** or **false**. The default value is **false**.

## Index Merge Task Optimization

Enabling the three optimization options above will increase the number of index merge tasks in the cluster. You can configure **index.merge.scheduler.max\_thread\_count** to reduce the impact of index merge tasks on data ingestion performance. The optimization is about increasing the number of shard merging threads and thereby alleviate the throttling of data import.

Run the following command to start index merge task optimization:

```
PUT my_index
{
 "settings": {
 "index.merge.scheduler.max_thread_count": 8
 }
}
```

The value range of **index.merge.scheduler.max\_thread\_count** is [1, node.processors/2]. The default value is 4, and the recommended value is 8.

## 2.7 Searching Data in an Elasticsearch Cluster

### 2.7.1 Using DSL to Search for Data in Elasticsearch

DSL is the specified query language for Elasticsearch. It is the best language for interaction between Elasticsearch clusters and clients. Elasticsearch DSL is a JSON-based language. Other languages, such as SQL, are translated into Elasticsearch DSL before they can be used for interacting with Elasticsearch clusters.

This topic lists some of the most commonly used Elasticsearch DSL query statements. For more, see [Query DSL](#).

## DSL Usage Example

Compile the request content in JSON format on the Dev Tools page of Kibana and execute the search request.

For example, run the following command to retrieve all documents in the **test** index:

```
GET /test/_search
{
 "query": {
 "match_all": {}
 }
}
```

The search result is also in JSON format.

## Common DSL Query Statements

- Sets the query filters, which is equivalent to **where** in the SQL language. In the command below, there is no index filter in front of **\_search**, so all indexes are queried. A bool query allows you to combine multiple search queries with boolean conditions. **filter** forcibly filters documents whose **status** field is **published** and **publish\_date** is later than **2015-01-01**. **must** specifies that both **title** and **content** must include **Search**.

### NOTE

The difference between **must** and **filter** is that **filter** is equivalent to **where** in SQL but its results are not used for scoring. The **must** field is also a mandatory filter criteria, but the matching documents are scored based on relevance. The most relevant documents are displayed at the top.

```
GET /_search
{
 "query": {
 "bool": {
 "must": [
 {
 "match": {
 "title": "Search"
 }
 },
 {
 "match": {
 "content": "search"
 }
 }
],
 "filter": [
 {
 "term": {
 "status": "published"
 }
 },
 {
 "range": {
 "publish_date": {
 "gte": "2015-01-01"
 }
 }
 }
]
 }
 }
}
```

- **Aggregations** are similar to **Group by** in SQL.

An aggregation summarizes your data as metrics, statistics, or other analytics. In the example below, the results are aggregated based on the title field in the **test** index. If **title** is of the text (including keyword) type, use **title.keyword** for aggregation. By default, a cluster cannot directly aggregate data of the text type. **titles** is only an example name of the aggregation. You can name the aggregation **titleaggs** instead.

```
GET /test/_search
{
 "aggs": {
 "titles": {
 "terms": {
 "field": "title.keyword"
 }
 }
 }
}
```

The example above for query aggregation includes all documents in the **test** index. That is, **match\_all** is used. You can set search criteria to narrow the scope of the aggregation to specific documents.

## 2.7.2 Using SQL to Search for Data in Elasticsearch

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).

Users who are familiar with SQL can use SQL to search for data in Elasticsearch clusters.

### SQL Usage Example

- Use SQL to search for data on Kibana (recommended).

Log in to Kibana and, on the **Dev Tools** page, send requests to **\_opendistro/\_sql** using request parameters or the request body.

For example, run the following command to retrieve 50 records from the **my-index** index:

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

By default, the query result is returned in JSON format. To return data in CSV format, set the **format** parameter in the command as follows:

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

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

- Run cURL commands on an ECS to execute SQL statements to search for data.

For example, run the following command to retrieve 10 records from the **kibana\_sample\_data\_flights** index.

```
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'
```



**localhost** indicates the cluster access address. **username** and **password** indicate the username and password for accessing the security-mode cluster, respectively.

## Supported SQL Operations

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

- Statements

**Table 2-56** 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 2-57** 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                 |

| Condition | Example                                               |
|-----------|-------------------------------------------------------|
| In        | SELECT * FROM my-index WHERE name IN ('Bob', 'David') |
| Not       | SELECT * FROM my-index WHERE name NOT IN ('Bob')      |
| 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 2-58** 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 2-59** Include and exclude fields

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

- Functions

**Table 2-60** 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 |
| /           | 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 2-61** 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 constraints, see [Joins](#).

- Show

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

**Table 2-62** 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 each 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](#).

## 2.8 Enhancing Search Capabilities for Elasticsearch Clusters

### 2.8.1 Search Enhancement Features for Elasticsearch Clusters

Compared with open-source Elasticsearch, Elasticsearch clusters in CSS have many enhanced features. [Table 2-63](#) lists these enhanced features and the corresponding cluster versions.

**Table 2-63** Search enhancement features for CSS Elasticsearch clusters

| Enhanced Feature           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Cluster Version                           | Details                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|----------------------------------------------------------------------------------------|
| Vector search              | Unstructured data, such as images, videos, and language corpora, is converted into vectors, which are searched based on similarity using either an exact or approximate nearest neighbors algorithm.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Elasticsearch 7.6.2, Elasticsearch 7.10.2 | <a href="#">Configuring Vector Search for Elasticsearch Clusters</a>                   |
| Storage-compute decoupling | CSS stores new data as hot data on SSDs to ensure optimal query performance, and historical data as cold data in OBS to cut storage costs.<br><br>Compared with cold/hot storage switchover, storage-compute decoupling is a better option for use cases that are not particularly demanding in terms of search performance, as cold data is stored in OBS, which cuts storage costs.                                                                                                                                                                                                                                                                                                                                                                                                                                        | Elasticsearch 7.6.2, Elasticsearch 7.10.2 | <a href="#">Configuring Decoupled Storage and Compute for an Elasticsearch Cluster</a> |
| Flow Control 2.0           | Flow Control 2.0 controls traffic at the node level. You can configure blacklists and whitelists per node, the maximum concurrent HTTPS connections allowed, and the maximum HTTP connections allowed. You can also configure node heap memory-based flow control that uses a backpressure mechanism, or one-click traffic blocking. CSS can also collect statistics on node access IP addresses and URLs. When backpressure is enabled, large requests will be rejected by nodes when their heap memory usage is high. This mechanism prevents nodes from breaking down and reduces the risk of node unavailability.<br><br><b>NOTE</b><br>Elasticsearch 7.6.2 and Elasticsearch 7.10.2 clusters created after January 2023 support Flow Control 2.0 only, whereas those created before that support Flow Control 1.0 only. | Elasticsearch 7.6.2, Elasticsearch 7.10.2 | <a href="#">Configuring Flow Control 2.0 for an Elasticsearch Cluster</a>              |

| Enhanced Feature      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Cluster Version                           | Details                                                                        |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|--------------------------------------------------------------------------------|
| Flow Control 1.0      | <p>Flow Control 1.0 controls traffic at the node level. You can configure blacklists and whitelists per node, the maximum concurrent HTTP connections allowed, the maximum HTTP connections allowed, the maximum heap memory used by specific request paths, and the maximum CPU usage. You can block access in one click, and collect statistics on IP addresses and URLs accessing the nodes. If flow control is enabled, requests will be blocked at the entry, which alleviates the cluster pressure in high-concurrency scenarios and reduces the likelihood of unavailability issues.</p> <p><b>NOTE</b><br/>Elasticsearch 7.6.2 and Elasticsearch 7.10.2 clusters created after January 2023 support Flow Control 2.0 only, whereas those created before that support Flow Control 1.0 only.</p> | Elasticsearch 7.6.2, Elasticsearch 7.10.2 | <a href="#">Configuring Flow Control 1.0 for an Elasticsearch Cluster</a>      |
| Large query isolation | <p>You can isolate query requests that consume a large amount of memory or take a long period of time. This way, you ensure service availability for other requests. If the heap memory usage of a node is too high, an interrupt control program will be triggered to terminate a large query based on the policies you configured. You can also configure a global query timeout duration. Long queries will be intercepted by an Elasticsearch-native cancel API.</p>                                                                                                                                                                                                                                                                                                                                | Elasticsearch 7.6.2, Elasticsearch 7.10.2 | <a href="#">Configuring Large Query Isolation for an Elasticsearch Cluster</a> |
| Enhanced aggregation  | <p>Enhances aggregation performance in the face of large data volumes by leveraging vectorization and optimized clustering, enabling faster analytics and decision-making in complex situations.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Elasticsearch 7.10.2                      | <a href="#">Configuring Enhanced Aggregation for an Elasticsearch Cluster</a>  |

| Enhanced Feature                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Cluster Version                                                       | Details                                                                             |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| Read/write splitting                    | Data written to the primary cluster ( <b>Leader</b> ) is automatically synchronized to the secondary cluster ( <b>Follower</b> ). This allows data to be queried from the secondary cluster, improving query performance while alleviating the pressure of the primary cluster. If the primary cluster is unable to provide services, a primary/secondary switchover can be performed to use the secondary cluster to handle write and query requests, ensuring service continuity.                                                                                                                                                                                                                                                                                                                                                                                                 | Elasticsearch 7.6.2, Elasticsearch 7.10.2                             | <a href="#">Configuring Read/Write Splitting Between Two Elasticsearch Clusters</a> |
| Switchover between hot and cold storage | Hot/cold storage switchover is about allocating data to different types of data nodes in terms of their performance and costs based on how frequent data is expected to be accessed. Hot data nodes store real-time data that is frequently updated and queried. Typically, they use high-performance hardware (such as SSDs) to ensure fast read/write and retrieval. Cold data nodes store historical data that is seldom accessed. Typically, cold data nodes use lower-cost hardware to store data.<br><br>Compared with storage-compute decoupling, cold/hot storage switchover is a better option for use cases that are demanding in terms of search performance. Cold data is stored on local cold data nodes in the cluster. The storage capacity available depends on the number of cold data nodes and their disk capacity. The storage cost is higher than that of OBS. | This feature is supported as long as the cluster has cold data nodes. | <a href="#">Switching Between Hot and Cold Storage for an Elasticsearch Cluster</a> |
| Index recycle bin                       | Similar to any other type of recycle bin, the index recycle bin temporarily stores deleted indexes so that users can restore them before they are finally removed from the recycle bin. This helps to prevent misdeletion of data and improve cluster data reliability.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Elasticsearch 7.10.2                                                  | <a href="#">Configuring an Index Recycle Bin for an Elasticsearch Cluster</a>       |

## 2.8.2 Configuring Vector Search for Elasticsearch Clusters

### 2.8.2.1 About CSS Vector Search

In the age of AI, unstructured data—such as images, videos, audios, and text—is growing rapidly. Traditional keyword-based search cannot effectively handle unstructured data because it cannot capture deep semantic or visual features. To enable effective search across unstructured data, CSS provides a vector search solution that enables high-performance, high-accuracy nearest neighbor or approximate nearest neighbor search. Typical use cases include image search, video search, similar product recommendations, semantic text search, and cross-modal search (for example, searching for images using text).

#### Advantages

- **Efficient and reliable:** The built-in CSS vector search engine provides high search throughput and low latency. Based on the Elasticsearch native distributed architecture, CSS delivers enterprise-grade reliability guarantees, including multi-replica, snapshots, and permission control.
- **Flexible adaptation:** CSS supports multiple indexing algorithms, including brute-force search, graph-based algorithms (such as HNSW), product quantization (PQ), and IVF-HNSW; as well as multiple similarity measurement methods, including Euclidean, inner product, cosine, and Hamming.
- **Elasticsearch ecosystem compatibility:** CSS vector search is compatible with open-source Elasticsearch query language and APIs. It seamlessly integrates Elasticsearch ecosystem tools, such as Cerebro, Kibana, and Logstash; and it supports mainstream client languages such as Python, Java, Go, and C++, simplifying development and integration.

#### How It Works

CSS vector search uses approximate nearest neighbor (ANN) search to mitigate the intensive computational load of k-Nearest Neighbors (k-NN) search, perfectly balancing search efficiency and accuracy. Key points include:

- **Reducing the candidate set:** Traditional text search filters out irrelevant documents through inverted indexes, while vector search quickly retrieves potentially relevant vectors using index structures like HNSW graphs or IVF-PQ, avoiding scanning the full dataset. For example, HNSW uses a multi-layer structure to quickly find nearest neighbors to the query vector.
- **Reducing computational complexity:** The funnel model first performs coarse quantization (such as IVF-PQ) on vectors to quickly obtain the candidate set. Then, it performs fine retrieval or reranking (such as cosine similarity) on the candidate set. Quantization uses the product quantization (PQ) algorithm to encode high-dimensional vectors into compact codes, reducing storage and computational overheads.
- **Balancing performance and accuracy:** Indexing parameters (such as the number of HNSW layers and the number of IVF clusters) can be dynamically tuned to trade off between recall and query latency.



## Procedure

1. Data preparation: Use an AI model (such as CNN and Transformer) to process your unstructured data (such as images, videos, and text) and extract feature vectors.
2. Index creation: Create vector indexes in your Elasticsearch cluster and define vector field mappings, including specifying vector dimensions, indexing algorithms, and similarity measurement methods.
3. Data write: Store feature vectors (typically along with the original data or metadata) into these indexes.
4. Vector search: Use the standard Elasticsearch query DSL (such as KNN query) to provide the query vector (generated by the same model), and specify the number ( $k$ ) of nearest neighbors you expect to return.
5. Result: The CSS vector search engine performs an efficient ANN search and returns the  $k$  most relevant results and their similarity scores. Your application can then process these results (for example, showing similar images or recommending relevant products).

## Constraints

Only Elasticsearch 7.6.2 and 7.10.2 clusters offer the built-in CSS vector search engine. CSS's Elasticsearch clusters does not support open-source vector search.

## Related Documents

- To learn how to quickly get started with CSS's vector search service, see [Using Elasticsearch for Vector Search](#).
- To learn more about the CSS vector database performance, see [Testing the Performance of CSS's Elasticsearch Vector Search](#).

### 2.8.2.2 Creating an Elasticsearch Vector Cluster

Integrating efficient indexing techniques, the CSS vector database delivers a high-performance, low-cost, scalable solution for high-dimensional vector search. An Elasticsearch vector cluster converts unstructured data into high-dimensional vectors and uses vector indexing algorithms (such as HNSW graph-based indexing and product quantization) to enable approximate nearest neighbor (ANN) search. This significantly reduces computational complexity while ensuring a high recall rate.

This topic focus on the memory capacity requirements and planning for an Elasticsearch vector cluster, so as to provide guidance on how to choose cluster nodes of the appropriate specifications. All other procedures are the same as those for other regular search clusters. For details, see [Elasticsearch Cluster Planning Suggestions](#).

## Memory Planning

Before creating an Elasticsearch vector cluster, properly plan the cluster's memory capacity based on the data size, vector dimensions, and index types.

- **Node type**

The CSS vector search engine relies heavily on in-memory computing, and vector indexing relies on off-heap memory. For a vector cluster, we recommend **memory-optimized** nodes.

- **Memory calculation formula**

Estimate the required off-heap memory based on the index algorithms used, vector dimensions, and data volume.

- FLAT index:  $\text{mem\_size} = \text{dim} \times \text{dim\_size} \times \text{num} + \text{delta}$
- GRAPH index:  $\text{mem\_size} = (\text{dim} \times \text{dim\_size} + \text{neighbors} \times 4) \times \text{num} + \text{delta}$
- GRAPH\_PQ index:  $\text{mem\_size} = (\text{fragment\_num} + \text{neighbors} \times 4) \times \text{num} + \text{delta}$
- GRAPH\_SQ8 index:  $\text{mem\_size} = (\text{dim} \times 2 + \text{neighbors} \times 4) \times \text{num} + \text{delta}$
- GRAPH\_SQ4 index:  $\text{mem\_size} = (\text{dim} + \text{neighbors} \times 4) \times \text{num} + \text{delta}$
- IVF\_GRAPH and IVF\_GRAPH\_PQ indexes do not require resident memory, so there is need to estimate their memory requirements.

**Table 2-64** Parameter description

| Parameter    | Description                                                                                                                                                                                                                                                                                               |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| mem_size     | Size of off-heap memory required by vector indexes. (This estimate assumes there are no replicas. If index replicas are involved, the required memory size multiplies with the number of replicas.)                                                                                                       |
| dim          | Number of vector dimensions.                                                                                                                                                                                                                                                                              |
| dim_size     | Number of bytes required by each dimension. By default, each dimension is a 4-byte float value.                                                                                                                                                                                                           |
| neighbors    | Number of neighbors for each vector in a graph index. The default value is 64.                                                                                                                                                                                                                            |
| num          | Total number of vectors.                                                                                                                                                                                                                                                                                  |
| fragment_num | Number of segments a vector is split into when product quantization is used.<br>If fragment_num is not configured during index creation, the value is determined by <b>dim</b> .<br>if dim <= 256:<br>fragment_num = dim / 4<br>elif dim <= 512:<br>fragment_num = dim / 8<br>else :<br>fragment_num = 64 |
| delta        | Metadata size. This factor can be ignored.                                                                                                                                                                                                                                                                |

When selecting node 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's physical memory, and the maximum size is **31 GB**.

Example:

When creating a graph index based on the SIFT10M dataset (128-dimensional vector, 10 million records), if **neighbors** is set to 32, the off-heap memory required is as follows: **mem\_size = (128 x 4 + 32 x 4) x 10,000,000 + delta ≈ 6 GB**. If one index replica is needed, at least 12 GB (6 x 2) of off-heap memory is required. In this case, select one 8U32G or two 8U16G node for the cluster.

## Creating a Cluster

The procedure for creating a vector cluster is the same as that for creating any other regular search cluster. For details, see [Creating an Elasticsearch Cluster \(New Version\)](#).

Pay attention to the following key parameters:

- **Cluster Version:** Select 7.6.2 or 7.10.2. Only clusters of these two versions support the built-in CSS vector search engine.
- **Node Specifications:** Select the node specifications based on the result of cluster memory planning.

## (Optional) Configuring the Circuit Breaker

To mitigate out-of-memory (OOM) errors and maintain optimal vector query performance, a circuit breaker mechanism is employed. When the cluster's off-heap memory usage exceeds a predefined threshold, this mechanism automatically blocks vector data writes to the cluster. The purposes of this mechanism are as follows:

- Preventing memory overload: Write throttling lowers off-heap memory usage.
- Maintaining query performance: Optimal vector query performance can be maintained by preventing memory overload.

The off-heap memory circuit breaker is enabled by default. You can enable or disable it and adjust its threshold based on service requirements. The command is as follows:

```
PUT _cluster/settings
{
 "persistent": {
 "native.cache.circuit_breaker.enabled": "true",
 "native.cache.circuit_breaker.cpu.limit": "80%"
 }
}
```

**Table 2-65** Parameter description

| Parameter                              | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| native.cache.circuit_breaker.enabled   | Boolean | Whether to enable the off-heap memory circuit breaker.<br>Value range: <ul style="list-style-type: none"><li>• <b>true</b> (default value): Enable the off-heap memory circuit breaker. When the off-heap memory usage reaches the circuit breaker threshold, write requests are blocked.</li><li>• <b>false</b>: Disable the off-heap memory circuit breaker. OOM errors may occur in case of excessive off-heap memory usage.</li></ul>                                                                                  |
| native.cache.circuit_breaker.cpu.limit | String  | Circuit breaker threshold in terms of maximum off-heap memory usage.<br>This parameter is available only when <b>native.cache.circuit_breaker.enabled=true</b> .<br>Value range: a value in percentage<br>Default value: 80%<br>Assume a cluster uses 128 GB memory. The required heap memory is 31 GB, and the default circuit breaker threshold is 80%, then: $(128 - 31) \times 80\% = 77.6$ GB. This means when the off-heap memory usage exceeds 77.6 GB, the circuit breaker is triggered to block write operations. |

### 2.8.2.3 Creating a Vector Index

Create a vector index in your Elasticsearch cluster and define a mapping that contains vector fields, including vector dimensions, indexing algorithm, and similarity measurement methods. Then store vectors (typically along with the original data or metadata) into this index.

### Logging In to Kibana

Log in to Kibana and go to the command execution page. Elasticsearch clusters support multiple access methods. This topic uses Kibana as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.

4. In the left navigation pane, choose **Dev Tools**.

## Creating a Vector Index

Run the following command on Kibana to create a vector index.

For example, create an index named **my\_index**. This index contains a vector field named **my\_vector** and a text field named **my\_label**. A graph-based index is created for the vector field, and Euclidean distance is used for similarity measurement.

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

**Table 2-66** settings parameters

| Parameter                | Mandato<br>ry | Type    | Description                                                                                                                               |
|--------------------------|---------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------|
| index.vector             | Yes           | Boolean | Whether to enable vector indexes.<br>Set this parameter to <b>true</b> .<br>Otherwise, vector indexes cannot be created.                  |
| index.number_of_shards   | No            | Integer | Number of index shards. This value should be divisible by the number of cluster nodes.<br>Value range: 1–1024<br>Default value: <b>1</b>  |
| index.number_of_replicas | No            | Integer | Number of index replicas. Replicas improve data availability.<br>Value range: 0 to the number of nodes minus 1<br>Default value: <b>1</b> |

| Parameter                               | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------|-----------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.vector.exact_search_threshold     | No        | Integer | <p>Threshold for automatically switching from pre-filtering search to brute-force search. When the size of the intermediate result set in a segment is lower than this threshold, a brute-force search is performed.</p> <p>Value range: null (disables auto switchover from pre-filtering search to brute-force search) or a positive integer</p> <p>Default value: null (disables auto switchover from pre-filtering search to brute-force search)</p>                                                                                                                                                                                                                                                         |
| index.vector.search.concurrency.enabled | No        | Boolean | <p>Whether to enable concurrent vector searches across segments. In an Elasticsearch cluster, each index shard consists of multiple segments. By default, each segment is searched in sequence. Enabling concurrent searches across segments reduces query latency but does not increase the cluster's maximum query throughput. Furthermore, it may increase the average CPU usage of cluster nodes.</p> <p>This parameter is available only for Elasticsearch 7.10.2 clusters whose image version is 7.10.2_25.3.0_x.x.x or later.</p> <p>Value range:</p> <ul style="list-style-type: none"><li>• <b>true</b>: Enables concurrent search.</li><li>• <b>false</b> (default): Performs serial search.</li></ul> |

**Table 2-67** mappings parameters

| Parameter | Mandatory | Type    | Description                                                                                        |
|-----------|-----------|---------|----------------------------------------------------------------------------------------------------|
| type      | Yes       | String  | <p>Data type of a field.</p> <p>Set this parameter to <b>vector</b> to indicate vector fields.</p> |
| dimension | Yes       | Integer | <p>Number of vector dimensions.</p> <p>Value range: 1–4096</p>                                     |

| Parameter     | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| indexing      | No        | Boolean | <p>Whether to enable vector index acceleration.</p> <p>Value range:</p> <ul style="list-style-type: none"> <li><b>true</b>: Enables vector index acceleration. When this parameter is set to <b>true</b>, an extra vector index is created. The indexing algorithm is specified by the <b>algorithm</b> field and the index supports vector search.</li> <li><b>false</b>: Disables vector index acceleration. When this parameter is set to <b>false</b>, vector data is written only to docvalues, and only <b>ScriptScore</b> and <b>Rescore</b> can be used for vector search.</li> </ul>                                                                                                                                                                                                                                                                  |
| lazy_indexing | No        | Boolean | <p>Whether to enable delayed vector indexing, where vector indexes are not built immediately when data is ingested. The aim to speed up write operations.</p> <p>The configuration of lazy_indexing takes effect only when: <b>indexing</b> is set to <b>true</b>; the cluster version is Elasticsearch 7.10.2; and the cluster image version is no earlier than 7.10.2_24.3.3_x.x.x.</p> <p>Value range:</p> <ul style="list-style-type: none"> <li><b>true</b>: Enables delayed indexing. Vector indexes (such as graph-based indexes) are not built immediately when data is ingested. Instead, <b>(Optional) Post-processing after Data Ingestion: Offline Index Building</b> needs to be performed to build vector indexes. The vector indexes can then be used for vector search.</li> <li><b>false</b> (default): Disables delayed indexing.</li> </ul> |

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------|-----------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| algorithm | No        | String | <p>Vector indexing algorithm.</p> <p>This parameter is valid only when <b>indexing</b> is set to <b>true</b>. When this parameter is set to <b>IVF_GRAPH</b> or <b>IVF_GRAPH_PQ</b>, <b>(Optional) Pre-Building and Registering Centroid Vectors</b> is required.</p> <p>Value range:</p> <ul style="list-style-type: none"> <li>• <b>FLAT</b>: 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 can reach 100%. You can use this algorithm if you require high recall accuracy.</li> <li>• <b>GRAPH</b> (default): Hierarchical Navigable Small Worlds (HNSW) algorithm for graph-based indexes. This algorithm is mainly used when high performance and precision are required and the number of documents in a single shard reaches 10 million.</li> <li>• <b>GRAPH_PQ</b>: a combination of the HNSW algorithm and the PQ algorithm. The PQ algorithm reduces the storage overhead of the original vectors, so that HNSW can easily search through hundreds of millions of records.</li> <li>• <b>GRAPH_SQ8</b>: a combination of the HNSW algorithm and the scalar quantization (SQ) algorithm. By quantizing float32 values into int8, this algorithm reduces the storage overhead of the original vectors and improves build and query efficiency. The downside is a slightly decreased recall rate. Only Elasticsearch 7.10.2 clusters support this algorithm.</li> <li>• <b>GRAPH_SQ4</b>: a combination of the HNSW algorithm and the SQ algorithm. By quantizing float32 values into int4, this algorithm</li> </ul> |



| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------|-----------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           |           |        | <p>reduces the storage overhead of the original vectors and improves build and query efficiency. The downside is a slightly decreased recall rate. SQ4 has a higher quantization/compression ratio and higher computational efficiency than SQ8, but also a large decrease in recall rate. Only Elasticsearch 7.10.2 clusters support this algorithm.</p> <ul style="list-style-type: none"><li>• <b>IVF_GRAPH</b>: a 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 through hundreds of millions of records.</li><li>• <b>IVF_GRAPH_PQ</b>: a 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 when there are more than 1 billion documents in shards and high retrieval performance is required.</li></ul> <p>When the indexing algorithm is set to GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4, CSS provides additional parameters, as shown in <a href="#">Table 2-68</a> and <a href="#">Table 2-69</a>, that you can choose to configure to enhance query performance and accuracy.</p> |
| dim_type  | No        | String | <p>Vector data type.</p> <p>Value range:</p> <ul style="list-style-type: none"><li>• binary: binary value</li><li>• float (default): floating-point number</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| metric    | No        | String | <p>Vector distance metric, which measures the similarity or distance between vectors.</p> <p>Value range:</p> <ul style="list-style-type: none"><li>• <b>euclidean</b> (default): Euclidean distance</li><li>• <b>inner_product</b>: inner product distance</li><li>• <b>cosine</b>: cosine distance</li><li>• <b>hamming</b>: Hamming distance, which can be used only when <b>dim_type</b> is set to <b>binary</b>.</li></ul> |

**Table 2-68** Optional parameters for the GRAPH indexing algorithm

| Parameter | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                          |
|-----------|-----------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| neighbors | No        | Integer | <p>Number of neighbors of each vector in the graph index. A larger value results in higher query accuracy but slower index building and query.</p> <p>This parameter is available only when <b>indexing</b> is set to <b>true</b>, and <b>algorithm</b> is GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 20–255</p> <p>Default value: <b>64</b></p>                  |
| shrink    | No        | Float   | <p>How aggressively the HNSW graph removes redundant edges during construction. This setting affects the structure of the HNSW graph.</p> <p>Default value: 1.0f</p> <p>This parameter is available only when <b>indexing</b> is set to <b>true</b>, and <b>algorithm</b> is GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 0.1–10</p> <p>Default value: <b>1</b></p> |

| Parameter    | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                            |
|--------------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| scaling      | No        | Integer | <p>Scaling ratio for the number of upper-layer graph nodes in the HNSW graph. This setting affects the layers of the HNSW graph.</p> <p>This parameter is available only when <b>indexing</b> is set to <b>true</b>, and <b>algorithm</b> is GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 0-128</p> <p>Default value: 50</p>                                          |
| efc          | No        | Integer | <p>How many nearest neighbors to explore when inserting a new vector into the HNSW graph. A larger value results in higher accuracy but slower index building.</p> <p>This parameter is available only when <b>indexing</b> is set to <b>true</b>, and <b>algorithm</b> is GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 0-100000</p> <p>Default value: <b>200</b></p> |
| max_scan_num | No        | Integer | <p>Maximum number of nodes to be scanned during search or index building. A larger value results in higher query accuracy but slower indexing.</p> <p>This parameter is available only when <b>indexing</b> is set to <b>true</b>, and <b>algorithm</b> is GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 0 to 1000000</p> <p>Default value: <b>10000</b></p>           |

**Table 2-69** Optional parameters for the GRAPH\_PQ indexing algorithm

| Parameter    | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| centroid_num | No        | Integer | Number of centroids used during the coarse quantization stage of the algorithm. It affects quantization granularity and storage.<br><br>This parameter is available only when <b>indexing</b> is set to <b>true</b> , and <b>algorithm</b> is GRAPH_PQ.<br><br>Value range: 0–65535<br>Default value: 255                                                                                                                                                                              |
| fragment_num | No        | Integer | Number of fragments each vector is split into. It affects the PQ quantization granularity.<br><br>The default value is <b>0</b> . The plug-in automatically sets the number of fragments based on the vector length.<br><br>This parameter is available only when <b>indexing</b> is set to <b>true</b> , and <b>algorithm</b> is GRAPH_PQ.<br><br>Value range: 0–4096<br>Default value: <b>0</b> . The plug-in automatically sets the number of fragments based on the vector length. |

## (Optional) Pre-Building and Registering Centroid Vectors

When the IVF\_GRAPH or IVF\_GRAPH\_PQ algorithm is used for vector indexing, you need to pre-build and register centroid vectors.

IVF\_GRAPH and IVF\_GRAPH\_PQ help to accelerate indexing and queries in ultra-large-scale clusters. They allow you to narrow down the query scope by dividing a vector space into subspaces through clustering or random sampling. Before pre-build, you need to obtain all centroid vectors through clustering or random sampling. Centroid vectors are pre-built into a GRAPH or GRAPH\_PQ index and then registered with the Elasticsearch cluster. All nodes in the cluster can share this index. Reuse of the centroid index among shards can effectively reduce the training overhead and the number of centroid index queries, improving write and query performance.

1. Create a centroid index.

For example, run the following command on Kibana to create a centroid index named **my\_dict**:

```
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"
 }
 }
}
```

For detailed parameter configuration, see [Creating a Vector Index](#). Pay attention to the following mandatory parameters:

- **index.number\_of\_shards**: The number of index shards must be set to **1**. Otherwise, the centroid index cannot be registered.
- **indexing**: This parameter must be set to **true** to enable vector index acceleration.
- **algorithm**: Set the indexing algorithm. Set it to **GRAPH** for the IVF\_GRAPH algorithm, and GRAPH\_PQ if the IVF\_GRAPH\_PQ algorithm is used.

2. Write centroid vectors to the created index. Write the centroid vectors obtained through sampling or clustering into the newly created index **my\_dict**.
3. Call the registration API.

Run the following command on Kibana to register the centroid index as a Dict object with a globally unique name (**dict\_name**):

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

4. Create an IVF\_GRAPH or IVF\_GRAPH\_PQ vector index.

When creating the vector index, you do not need to specify **dimension** or **metric**. Rather, you specify the registered Dict object. [Table 2-70](#) describes key parameters for specifying a Dict object. For details about other parameters, see [Creating a Vector Index](#).

For example, run the following command to create an IVF\_GRAPH vector index:

```
PUT my_index
{
 "settings": {
 "index": {
 "vector": true,
 "sort.field": "my_vector.centroid" # Set the centroid subfield of each vector field as a ranking field.
 }
 },
 "mappings": {
 "properties": {
 "my_vector": {
 "type": "vector",
 "indexing": true,

```

```
"algorithm": "IVF_GRAPH",
 "dict_name": "my_dict",
 "offload_ivf": true
}
}
}
```

**Table 2-70** Field mappings parameters

| Parameter   | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------|-----------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dict_name   | Yes       | String  | Name of the centroid index. For example, <b>my_dict</b> . The vector dimensions and metrics of the index must be the same as those of the Dict index.                                                                                                                                                                                                                                                                                               |
| offload_ivf | Yes       | Boolean | Whether to offload the IVF inverted index to Elasticsearch.<br>Value: <b>true</b> or <b>false</b> . <ul style="list-style-type: none"><li>• <b>true</b> (recommended value): Offloads the IVF inverted index implemented by the underlying index to Elasticsearch. This reduces the use of off-heap memory and the overhead of write and merge operations.</li><li>• <b>false</b> (default value): Not to offload the IVF inverted index.</li></ul> |

#### 2.8.2.4 Importing Vector Data

Importing vector data is the process of ingesting data to the CSS vector database. When writing vector data to a vector index, you need to specify the vector field (for example, **my\_vector**) and the corresponding data format. The CSS vector database supports two common formats: floating-point arrays and Base64.

- Floating-point arrays: transmitting readable arrays directly.
- Base64: encoding vectors (little-endian byte order) into character strings to reduce network transmission overhead and improve efficiency in handling high-dimensional/binary vectors.

Choose a format based on the characteristics of your data. Also, choose an appropriate data importing method.

- Importing a single record: Use for small-scale applications or testing.
- Bulk import: Use for large-scale applications, where write requests are merged to reduce network overhead.

## Constraints

- Ensure that the vector field names and vector dimensions are consistent with those defined for the index.
- Base64 encoding must use the little-endian byte order. Otherwise, parsing errors may occur.
- In the case of bulk imports, you are advised to submit 100 to 1000 records per request. This balances throughput and memory overhead.

## Importing a Single Record

- Floating-point array

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

- Base64

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

## Bulk Import

- Floating-point array

```
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"}
```

- Base64

```
POST my_index/_bulk
{"index": {}}
{"my_vector": "AACAPwAAAAEA=", "my_label": "red"}
{"index": {}}
{"my_vector": "AAAAQAAAAEA=", "my_label": "green"}
{"index": {}}
{"my_vector": "AAAAQAAQEA=", "my_label": "red"}
```

For details about how to use the Bulk API, see [Bulk API](#).

## (Optional) Post-processing after Data Ingestion: Offline Index Building



### WARNING

- Use offline index building via an API only when: real-time data is not required or crucial; the cluster version is Elasticsearch 7.10.2 or later; and the cluster image version is no earlier than 7.10.2\_24.3.3\_x.x.x.
- If lazy\_indexing is enabled, offline index building must be performed after data ingestion. Otherwise, the system will return error code 500 for standard vector query, with the error message "Load native index failed exception." To solve this problem, perform offline index building before vectors queries.

Elasticsearch uses an LSM (Log-Structured Merge) tree-like model to accelerate write operations. As data is continuously written in and updated, numerous small

index segments are generated and later merged via a backend task to enhance query performance. As vector indexing is computationally intensive, frequent index merging while vector data is being written in consumes significant CPU resources. Therefore, where real-time data is not crucial, it is advisable to set **lazy\_indexing** to **true** for vector fields. This allows a final vector index to be created via a non-real time API after all data has been written in. This approach significantly reduces index merges, thereby improving overall write and index merging performance.

Offline index building consists of two steps:

1. Merge index segments.
2. Create the final vector index based on the final index segments.

The API used for offline index building is as follows:

```
POST _vector/indexing/{index_name}
{
 "field": "{field_name}"
}
```

where, {index\_name} indicates the name of the index to create. {field\_name} indicates the name of the vector field for which **lazy\_indexing** has been set to **true**.

### 2.8.2.5 Performing Vector Search

The CSS vector database supports a variety of query methods, including standard query, hybrid query, script\_score query, rescore query, and painless syntax extension. They enable efficient vector data search by accommodating diverse search needs.

- **Standard Query**: retrieves documents that are most similar to the query vector.
- **Hybrid Query**: combines vector search with traditional Elasticsearch queries, such as pre-filtering and Boolean queries.
- **Script Score Query**: enables custom similarity calculations for vector searches by executing a custom script
- **Rescore Query**: rescores and reranks the top results returned by an initial query to improve recall.
- **Painless Syntax Extension**: allows the use of vector distance or similarity calculation functions in custom scripts.

### Standard Query

Standard query is used to retrieve documents that are most similar to the query vector.

The following command will return  $k$  (specified by **size/topk**) records that are the closest matches to the query vector.

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



```
"topk":2
}
}
}
}
```

**Table 2-71** Parameters for standard query

| Parameter              | Mandatory | Type         | Description                                                                                                                                                                                                                                                                                                              |
|------------------------|-----------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| size                   | Yes       | Integer      | Number of search results to return.<br>Default value: <b>10</b>                                                                                                                                                                                                                                                          |
| _source                | No        | Boolean      | Whether to return the source text in documents. To reduce data transmission and improve query performance, set this parameter to <b>false</b> .<br>Value range: <ul style="list-style-type: none"><li>• <b>true</b> (default): Returns the source text.</li><li>• <b>false</b>: Not to return the source text.</li></ul> |
| query                  | Yes       | Map          | Specifies the query vector.<br>Parameters:<br><b>vector</b> (mandatory): indicates a vector query (vector similarity-based search), including the vector field and query vector value.<br><b>my_vector</b> (mandatory): queried vector field (for example, <b>my_vector</b> ).                                           |
| vector (sub-parameter) | Yes       | Array/String | Query vector value. It is used to calculate the similarity between indexed vectors and the query vector.<br>The value can be an array (for example, [1, 1]) or Base64-encoded value (for example, <b>AAABAAACAAAD</b> ).                                                                                                 |
| topk                   | Yes       | Integer      | The number of the most similar or relevant results to be returned.<br>Default value: same as <b>size</b> .                                                                                                                                                                                                               |

| Parameter    | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                       |
|--------------|-----------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ef           | No        | Integer | <p>How many nearest neighbors to explore when inserting a new vector into the graph. A larger value indicates a higher query accuracy yet slower query speed.</p> <p>This parameter is available only when <b>algorithm</b> is set to GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 0-100000</p> <p>Default value: <b>200</b></p> |
| max_scan_num | No        | Integer | <p>Maximum number of graph nodes to scan during search. A larger value indicates a higher query accuracy yet slower query speed.</p> <p>This parameter is available only when <b>algorithm</b> is set to GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 0-1000000</p> <p>Default value: <b>10000</b></p>                           |
| nprobe       | No        | Integer | <p>Number of centroids to explore during an IVF index query. A larger value indicates a higher query accuracy yet slower query speed.</p> <p>This parameter is available only when <b>algorithm</b> is IVF_GRAPH or IVF_GRAPH_PQ.</p> <p>Value range: 0-100000</p> <p>Default value: <b>100</b></p>                                               |

## Hybrid Query

Hybrid query combines vector search with traditional Elasticsearch queries, such as pre-filtering and Boolean queries.

---

### WARNING

Only Elasticsearch 7.10.2 clusters support pre-filtering queries.

---

In the following example, the top 10 records whose **my\_label** value is **red** are returned.

- **Pre-filtering query**

First, filters are applied to retrieve matching results. Then, vector search is performed on these results to retrieve the most relevant vectors based on similarity.

The following is an example:

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

**Table 2-72** Parameters for pre-filtering query

| Parameter | Mandatory | Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------|-----------|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| filter    | Yes       | Map  | Vector query filters. Standard Elasticsearch query filters are supported, such as term and range. If <b>filter</b> is too restrictive, leading to a small intermediate result set, you can set the <b>index.vector.exact_search_threshold</b> parameter, so that when the intermediate result set is smaller than this threshold, pre-filtering query automatically switches over to brute-force query (FLAT algorithm), which ensures a high recall rate. For more information, see <a href="#">Creating a Vector Index</a> . |
| term      | No        | Map  | Term query is a type of exact query. Documents that contain the exact term will be returned. For example, <b>{"term": {"my_label": "red"}}</b> means only to return documents whose <b>my_label</b> value is <b>red</b> .                                                                                                                                                                                                                                                                                                      |

- **Boolean query**

A Boolean query is in fact a post-filtering query method. Filtering and vector similarity-based search are performed separately. Then, the results of the two are combined using Boolean logic defined by clauses like **must**, **should**, and **filter**.

The following is an example:

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

**Table 2-73** Boolean query parameters

| Parameter | Mandatory | Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------|-----------|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| bool      | Yes       | Map  | <p>A compound query clause that combines subqueries using configured Boolean logic.</p> <p>Parameter description:</p> <ul style="list-style-type: none"><li>• <b>must</b>: Clauses that must match for documents to be included in the results.</li><li>• <b>filter</b>: It is similar to <b>must</b>, but do not contribute to the relevance score.</li><li>• <b>should</b>: Clauses that should match, but are not required. They are like nice-to-haves.</li><li>• <b>must_not</b>: Clauses that must not match for documents to be included in the results.</li></ul> |
| bool.must | Yes       | Map  | <p>Clauses that must match for documents to be included in the results. Parameter description:</p> <ul style="list-style-type: none"><li>• <b>vector</b>: query vector</li><li>• <b>my_vector</b>: vector field</li><li>• <b>topk</b>: number of results to return</li></ul>                                                                                                                                                                                                                                                                                              |

| Parameter   | Mandatory | Type | Description                                                                                                                                        |
|-------------|-----------|------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| bool.filter | Yes       | Map  | Clauses that must match, but do not contribute to the relevance score. Standard Elasticsearch query filters are supported, such as term and range. |

## Script Score Query

Script\_score query enables custom similarity calculations for vector searches by executing a user-defined script. It works as follows:

Pre-filtering works with any query. **script\_score** then calculates vector similarity on the pre-filtered results, and ranks the results. This query method does not use vector indexes. Its performance depends on the size of the intermediate result set after the pre-filtering. If the pre-filtering condition is set to **match\_all**, a brute-force search is performed on all data.

The following is an example:

```
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 2-74** script\_score query parameters

| Parameter    | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------|-----------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| script_score | Yes       | Map              | A root parameter for the script_score query.<br>Parameter description: <ul style="list-style-type: none"><li>• <b>query</b>: pre-filtering criteria. When it is set to <b>match_all</b>, a brute-force search is performed on all data.</li><li>• <b>script</b>: custom script that calculates similarity scores.</li></ul>                                                                                          |
| source       | Yes       | String           | Script name. The value is fixed to <b>vector_score</b> , indicating that a built-in script is used for calculating similarity.                                                                                                                                                                                                                                                                                       |
| lang         | Yes       | String           | Script language type. The value is fixed to <b>vector</b> .                                                                                                                                                                                                                                                                                                                                                          |
| field        | Yes       | String           | Queried vector field, for example, <b>my_vector</b> .                                                                                                                                                                                                                                                                                                                                                                |
| vector       | Yes       | Array/<br>String | Query vector value. It is used to calculate the similarity between indexed vectors and the query vector.<br>The value can be an array (for example, [1, 1]) or Base64-encoded value (for example, <b>AAABAAACAAAD</b> ).                                                                                                                                                                                             |
| metric       | Yes       | String           | Vector distance metric, which measures the similarity or distance between vectors.<br>Value range: <ul style="list-style-type: none"><li>• <b>euclidean</b> (default): Euclidean distance</li><li>• <b>inner_product</b>: inner product distance</li><li>• <b>cosine</b>: cosine distance</li><li>• <b>hamming</b>: Hamming distance, which can be used only when <b>dim_type</b> is set to <b>binary</b>.</li></ul> |

## Rescore Query

Rescore query rescores and reranks the top results returned by an initial query to improve recall.

When the GRAPH\_PQ or IVF\_GRAPH\_PQ indexing algorithm is used, query results are ranked based on the asymmetric distance calculated by PQ. Rescore query then rescores and reranks the initial search results to improve recall.

The following is an example of rescore query on a PQ index named **my\_index**:

```
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 2-75** Rescore query parameters

| Parameter   | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                               |
|-------------|-----------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rescore     | Yes       | Map     | Defines rescoring parameters.<br>Key parameters: <ul style="list-style-type: none"><li>• <code>window_size</code>: rescoring/reranking window size.</li><li>• <code>vector_rescore</code>: other vector rescoring settings.</li></ul>                                                                                     |
| window_size | Yes       | Integer | Rescoring/reranking window size. The vector search returns the top k results, but only the first <i>window_size</i> results are rescored and reranked. A larger value indicates a larger reranking scope and hence a higher recall rate, but it also leads to higher computational overhead.<br>Default value: <b>100</b> |
| field       | Yes       | String  | Queried vector field, for example, <b>my_vector</b> .                                                                                                                                                                                                                                                                     |

| Parameter | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------|-----------|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| vector    | Yes       | Array/<br>String | Query vector value. It is used to calculate the similarity between indexed vectors and the query vector.<br><br>The value can be an array (for example, [1, 1]) or Base64-encoded value (for example, <b>AAABAAACAAAD</b> ).                                                                                                                                                                                             |
| metric    | Yes       | String           | Vector distance metric, which measures the similarity or distance between vectors.<br><br>Value range: <ul style="list-style-type: none"><li>• <b>euclidean</b> (default): Euclidean distance</li><li>• <b>inner_product</b>: inner product distance</li><li>• <b>cosine</b>: cosine distance</li><li>• <b>hamming</b>: Hamming distance, which can be used only when <b>dim_type</b> is set to <b>binary</b>.</li></ul> |

## Painless Syntax Extension

Painless syntax extension allows the use of vector distance or similarity calculation functions in custom scripts. CSS extension supports several vector distance/similarity calculation functions, which users can use readily in custom Painless scripts to build flexible rescoring formulas.

The following is an example:

```
POST my_index/_search
{
 "size": 10,
 "query": {
 "script_score": {
 "query": {
 "match_all": {}
 },
 "script": {
 "source": "1 / (1 + euclidean(params.vector, doc[params.field]))",
 "params": {
 "field": "my_vector",
 "vector": [1, 2]
 }
 }
 }
 }
}
```



**Table 2-76** Supported vector distance/similarity calculation functions

| Function Signature               | Description                                                                                                                                           |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| euclidean(Float[], DocValues)    | Euclidean distance                                                                                                                                    |
| cosine(Float[], DocValues)       | Cosine similarity                                                                                                                                     |
| innerproduct(Float[], DocValues) | Inner product                                                                                                                                         |
| hamming(String, DocValues)       | Hamming distance Only vectors whose <b>dim_type</b> is <b>binary</b> are supported. The input query vector must be a Base64-encoded character string. |

### 2.8.2.6 Optimizing Vector Cluster Performance

This topic explains how to optimize the performance of a CSS vector database from two aspects—write and query.

#### Optimizing Write Performance

Writing vector data incurs three major overheads: replica synchronization, index refresh, and segment merging. When index data is written in real time, frequent index refresh operations generate a large number of small segments. This triggers frequent vector index build and merge operations, which consume excessive CPU/IO resources. You can try the following solutions to optimize write performance.

Solution 1: temporarily disable replicas

- **Description**  
Temporarily disable replicas during data ingestion and enable them after data ingestion is complete. Use this solution when importing historical data in batches or performing a full update (for example, when initializing a vector database).
- **Operation**  
Set the number of replicas:

```
PUT my_index/_settings
{
 "number_of_replicas": 0
}
```
- **Result**  
Write performance is enhanced by avoiding real-time vector index building on replica nodes.

Solution 2: adjust the refresh interval

- **Description**  
Set the index refresh interval to 120s or longer to reduce the number of small segments generated during frequent index refreshes and also reduce the

vector index building overhead caused by segment merges. You can also disable automatic index refresh by setting the refresh interval to -1. Use this solution in high-throughput write scenarios (for example, when writing vectorized log data).

- Operation

Set **refresh\_interval**.

```
PUT my_index/_settings
{
 "refresh_interval": "120s"
}
```

- Result

Index refreshes occur less frequently. This reduces the number of small segments and also the overhead of segment merges, leading to enhanced write performance.

Solution 3: increase indexing threads

- Description

Increasing the number of threads for vector index building accelerates the indexing process. However, too many such threads will compete for query resources. Use this solution when there are sufficient CPU resources but the write latency is high—such as in GPU-accelerated environments.

- Operation

The default value of **native.vector.index\_threads** is 4. Change this value as needed.

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

- Result

Vector index building is accelerated, and the performance of concurrent writes is enhanced.

## Optimizing Query Performance

Query performance is affected by the following factors: the number of segments, the memory circuit breaker mechanism, and field recall. An excessively large number of segments impacts search efficiency; when off-heap memory becomes insufficient, vector index data is frequently swapped in and out of the memory; recalling all fields increases the load during the fetch phase. You can optimize query performance by addressing these factors.

Solution 1: perform force merge

- Description

After batch data ingestion, perform the force merge operation to forcibly merge segments, thus reducing the number of segments. Typically, you should perform this operation after data ingestion and before data query (for example, after a scheduled batch ingestion).

- Operation

Perform the force merge operation:

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

- **Result**  
Multiple segments are merged into a single segment. This reduces the file scanning overhead and accelerates the query speed.

Solution 2: adjust the upper limit of the segment size

- **Description**  
During batch writes, the maximum size of segments generated by the system is 5 GB. You can increase this upper limit to reduce the number of segments generated after automatic merging. Typically, you should perform this operation before batch data ingestion starts.

- **Operation**  
Increase the maximum segment size:

```
PUT my_index/_settings
{
 "index.merge.policy.max_merged_segment": "10gb"
}
```

- **Result**  
Increasing the maximum segment size helps to reduce the number of segments and thus accelerate query performance.

Solution 3: adjust the circuit breaker limit for off-heap memory

- **Description**  
When the off-heap memory required by vector indexes exceeds the circuit breaker limit, the index cache manager frequently swaps in and out index data from the cache, which slows down queries. You can raise the circuit breaker limit to reduce circuit breaking (indicated by `CircuitBreakingException` recorded in the log) resulted from insufficient memory.

- **Operation**  
The default circuit breaker limit for off-heap memory is 80%. You can adjust this limit as required.

```
PUT _cluster/settings
{
 "persistent": {
 "native.cache.circuit_breaker.cpu.limit": "85%"
 }
}
```

- **Result**  
It is less likely for vector index data to be swapped out from the memory, and query jitter is reduced.

Solution 4: optimize field recall

- **Description**  
If the query result needs to return only a few fields that are either keywords or numeric values, you can use the `docvalue_fields` parameter to fetch them. Use this method if only numeric or enumerated metadata (such as product IDs and class labels) needs to be fetched. It can significantly reduce overhead during the fetch phase.

- **Operation**  
Use the `docvalue_fields` parameter to fetch only specific fields:

```
POST my_index/_search
{
 "size": 2,
 "stored_fields": ["_none_"],
 "docvalue_fields": ["my_label"],
 "query": {
 "vector": {
 "my_vector": {
 "vector": [1, 1],
 "topk": 2
 }
 }
 }
}
```

- **Result**  
There is no need to parse the entire `_source` document. Column-oriented storage (docvalues) reduces the overhead during the fetch phase and improves query performance.

## Setting Cache Timeout

When the cluster's memory resources are insufficient, data is frequently updated, or high data freshness is required, you can enable automatic cache expiration to have inactive data cleared from the cache. This helps to optimize system performance, ensure data consistency, and improve query stability. Use this approach where data updates frequently or memory resources are stretching thin.

Run the following command to set cache timeout:

```
PUT _cluster/settings
{
 "persistent": {
 "native.cache.expiry.enabled": "true",
 "native.cache.expiry.time": "30m"
 }
}
```

**Table 2-77** Parameter description

| Parameter                   | Type    | Description                                                                                                                                                                                                                                                                             |
|-----------------------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| native.cache.expiry.enabled | Boolean | Whether to enable automatic cache expiration.<br>Value range: <ul style="list-style-type: none"><li>• <b>true</b>: Enable automatic cache expiration. Inactive data in the cache will be cleared.</li><li>• <b>false</b> (default value): Disable automatic cache expiration.</li></ul> |
| native.cache.expiry.time    | String  | Timeout of inactive cache items.<br>This parameter takes effect only when <b>native.cache.expiry.enabled=true</b> .<br>Value: a time string, for example, <b>24h</b> (24 hours) or <b>30m</b> (30 minutes).<br>Default value: <b>24h</b> .                                              |

### 2.8.2.7 Managing the Vector Index Cache

The CSS vector search engine is built on C++. It uses off-heap memory to enhance performance and efficiency. To better manage and optimize the vector index cache, the CSS vector database provides a number of dedicated APIs that can be used to flexibly monitor and tune cache usage to ensure stable query performance.

#### Checking Cache Statistics

Use the following API to check the off-heap memory usage and return key metrics such as the current off-heap memory usage, cache hits, and data loading times. These metrics help you understand the cache status, so you can adjust the cache based on its usage. In the implementation of the vector search plug-in, each segment constructs and stores an index file. During query, all these files are loaded to the off-heap memory, and the plug-in uses a cache mechanism to manage the off-heap memory.

```
GET /_vector/stats
```

Example response:

```
{
 "_nodes" : { # Node information
 "total" : 1, # Total number of nodes
 "successful" : 1, # Number of successful nodes
 "failed" : 0 # Number of failed nodes
 },
 "cluster_name" : "css-d3a7", # Cluster name
 "cpu_circuit_breaker_triggered" : false, # Whether circuit breaking is triggered
 "nodes" : {
 "cAHmVUZTR9ON7t6jxcDCkg" : { # Node UUID
 "cpu_cache_capacity_reached" : false, # Whether the off-heap memory usage of the current node
 reaches the upper limit
 "cpu_eviction_count" : 0, # Number of segment-level cache swap-outs on the current node
 "cpu_hit_count" : 0, # Number of segment-level cache hits on the current node
 "cpu_load_exception_count" : 0, # Number of segment-level index loading failures on the current
 node
 "cpu_load_success_count" : 0, # Number of segment-level index loading successes on the current
 node
 "cpu_miss_count" : 0, # Number of segment-level cache misses on the current node
 "cpu_query_memory_usage" : 0, # Off-heap memory usage on the current node, in KB
 "cpu_total_load_time" : 0 # Total segment-level index loading time on the current node, in ms
 }
 }
}
```

#### Preloading Vector Indexes

Use the following API to preload specified vector indexes to the off-heap memory. This ensures that these indexes can be quickly accessed in subsequent queries, with high query efficiency and low latency.

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

where, {index\_name} specifies an index.

Example response:

```
{
 "_shards" : {
 "total" : 1,
 "successful" : 1,

```

```
"failed" : 0
}
}
```

## Clearing the Cache

When the cache is full, the system automatically swaps its data. However, frequent cache swapping may affect query performance. You can use the following API to manually clear the vector index cache. This way, you can reclaim off-heap memory resources and ensure hot data query performance.

- Clear the full cache:  
`PUT /_vector/clear/cache`
- Clear specified indexes from the cache:  
`PUT /_vector/clear/cache/{index_name}`

Example response:

```
{
 "acknowledged" : "true"
}
```

### 2.8.2.8 Using Nested Fields for Vector Search

Nested fields allow multiple vectorized records to be stored in a single document. For example, in an RAG scenario, documents usually need to be segmented by paragraph or by a fixed length, and then vectorized into multiple semantic vectors. By means of nested fields, these vectors can be written into a same Elasticsearch document. For a document that contains multiple vector records, if the query vector matches any of them, the document is returned.

## Constraints

Only Elasticsearch 7.10.2 clusters support this feature.

## Creating a Vector Index

Create a vector index with nested fields. The index contains an **id** field whose type is **keyword**, and an **embedding** field whose type is **nested**. The embedding field contains two subfields: **chunk** and **emb**. The **chunk** subfield is of the **keyword** type, and the **emb** subfield is of the **vector** type.

```
PUT my_index
{
 "settings": {
 "index.vector": true
 },
 "mappings": {
 "properties": {
 "id": {
 "type": "keyword"
 },
 "embedding": {
 "type": "nested",
 "properties": {
 "chunk": {
 "type": "keyword"
 },
 "emb": {
 "type": "vector",
 "dimension": 2,

```

```
 "indexing": true,
 "algorithm": "GRAPH",
 "metric": "euclidean"
 }
}
}
}
}
```

## Importing Vector Data

Use the bulk operation to write data in arrays. Each document contains two vector records.

```
POST my_index/_bulk
{ "index": {} }
{ "id": 1, "embedding": [{ "chunk": 1, "emb": [1, 1] }, { "chunk": 2, "emb": [2, 2] }] }
{ "index": {} }
{ "id": 2, "embedding": [{ "chunk": 1, "emb": [2, 2] }, { "chunk": 2, "emb": [3, 3] }] }
{ "index": {} }
{ "id": 3, "embedding": [{ "chunk": 1, "emb": [3, 3] }, { "chunk": 2, "emb": [4, 4] }] }
```

## Vector Search

The nested query is required for nested fields. To perform such a query, you need to set the path parameter to specify the nested path, and set **score\_mode** to **max**, indicating the maximum similarity between all vectors in the document and the query vector.

- Standard query

Query the top 10 documents that are most similar to vector [1, 1].

```
GET my_index/_search
{
 "_source": { "excludes": ["embedding"] },
 "query": {
 "nested": {
 "path": "embedding",
 "score_mode": "max",
 "query": {
 "vector": {
 "embedding.emb": {
 "vector": [1, 1],
 "topk": 10
 }
 }
 }
 }
 }
}
```

An example of the query result:

```
{
 "took" : 2,
 "timed_out" : false,
 "_shards" : {
 "total" : 1,
 "successful" : 1,
 "skipped" : 0,
 "failed" : 0
 },
 "hits" : {
 "total" : {
 "value" : 3,
 "relation" : "eq"
 }
 }
}
```

```
,
"max_score" : 1.0,
"hits" : [
 {
 "_index" : "my_index",
 "_type" : "_doc",
 "_id" : "Hc4Vc5QBSxCnghau22AE",
 "_score" : 1.0,
 "_source" : {
 "id" : 1
 }
 },
 {
 "_index" : "my_index",
 "_type" : "_doc",
 "_id" : "Hs4Vc5QBSxCnghau22AE",
 "_score" : 0.33333334,
 "_source" : {
 "id" : 2
 }
 },
 {
 "_index" : "my_index",
 "_type" : "_doc",
 "_id" : "H84Vc5QBSxCnghau22AE",
 "_score" : 0.11111111,
 "_source" : {
 "id" : 3
 }
 }
]
}
```

- Pre-filtering query

First retrieve documents whose ID is ["2", "3"], and then return the top 10 documents that are most similar to the query vector [1, 1].

```
GET my_index/_search
{
 "query": {
 "nested": {
 "path": "embedding",
 "score_mode": "max",
 "query": {
 "vector": {
 "embedding.emb": {
 "vector": [1, 1],
 "topk": 10,
 "filter": {
 "terms": {"id": ["2", "3"]}
 }
 }
 }
 }
 }
 }
}
```

An example of the query result:

```
{
 "took" : 3,
 "timed_out" : false,
 "_shards" : {
 "total" : 1,
 "successful" : 1,
 "skipped" : 0,
 "failed" : 0
 },
 "hits" : {
```



```
"total" : {
 "value" : 2,
 "relation" : "eq"
},
"max_score" : 0.33333334,
"hits" : [
 {
 "_index" : "my_index",
 "_type" : "_doc",
 "_id" : "3t0ZypcB-Tff59gMTZO2",
 "_score" : 0.33333334,
 "_source" : {
 "id" : 2,
 "embedding" : [
 {
 "chunk" : 1,
 "emb" : [
 2,
 2
]
 },
 {
 "chunk" : 2,
 "emb" : [
 3,
 3
]
 }
]
 }
 },
 {
 "_index" : "my_index",
 "_type" : "_doc",
 "_id" : "390ZypcB-Tff59gMTZO2",
 "_score" : 0.11111111,
 "_source" : {
 "id" : 3,
 "embedding" : [
 {
 "chunk" : 1,
 "emb" : [
 3,
 3
]
 },
 {
 "chunk" : 2,
 "emb" : [
 4,
 4
]
 }
]
 }
 }
]
```

### 2.8.2.9 Sample Python Code for Vector Search

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

This section provides a sample of Python code for creating vector indexes, and importing and querying vector data. It shows how to use the client to implement vector search.

## Prerequisites

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

```
Set the actual cluster version. 7.6 is used in this example.
pip install elasticsearch==7.6
```

## Sample Code

```
from elasticsearch import Elasticsearch
from elasticsearch import helpers

Create the Elasticsearch client.
def get_client(hosts: list, user: str = None, password: str = None):
 if user and password:
 return Elasticsearch(hosts, http_auth=(user, password), verify_certs=False, ssl_show_warn=False)
 else:
 return Elasticsearch(hosts)

Create an index table.
def create(client: Elasticsearch, index: str):
 # Index mapping information
 index_mapping = {
 "settings": {
 "index": {
 "vector": "true", # Enable the vector feature.
 "number_of_shards": 1, # Set the number of index shards as needed.
 "number_of_replicas": 0 # Set the number of index replicas as needed.
 }
 },
 "mappings": {
 "properties": {
 "my_vector": {
 "type": "vector",
 "dimension": 2,
 "indexing": True,
 "algorithm": "GRAPH",
 "metric": "euclidean"
 }
 }
 # Other fields can be added if necessary.
 }
 }
 res = client.indices.create(index=index, body=index_mapping)
 print("create index result: ", res)

Write data.
def write(client: Elasticsearch, index: str, vecs: list, bulk_size=500):
 for i in range(0, len(vecs), bulk_size):
 actions = [
 {
 "_index": index,
 "my_vector": vec,
 # Other fields can be added if necessary.
 }
 for vec in vecs[i:i+bulk_size]
]
 success, errors = helpers.bulk(client, actions, request_timeout=3600)
 if errors:
 print("write bulk failed with errors: ", errors) # Handle the error as needed.
 else:
 print("write bulk {} docs success".format(success))
 client.indices.refresh(index=index, request_timeout=3600)

Query a vector index.
def search(client: Elasticsearch, index: str, query: list, size: int):
 # Query statement. Select an appropriate query method.
```

```
query_body = {
 "size": size,
 "query": {
 "vector": {
 "my_vector": {
 "vector": query,
 "topk": size
 }
 }
 }
}
res = client.search(index=index, body=query_body)
print("search index result: ", res)

Delete an index.
def delete(client: Elasticsearch, index: str):
 res = client.indices.delete(index=index)
 print("delete index result: ", res)

if __name__ == '__main__':
 # For a non-security cluster, run the following:
 es_client = get_client(hosts=['http://xx.xx.xx.xx:9200'])

 # For a security-mode cluster that uses HTTPS, run the following:
 # es_client = get_client(hosts=['https://xx.xx.xx.xx:9200', 'https://xx.xx.xx.xx:9200'], user='xxxxx',
 password='xxxxx')

 # For a security-mode cluster with HTTPS disabled, run the following:
 # es_client = get_client(hosts=['http://xx.xx.xx.xx:9200', 'http://xx.xx.xx.xx:9200'], user='xxxxx',
 password='xxxxx')

 # Test the index name.
 index_name = "my_index"

 # Create an index.
 create(es_client, index=index_name)

 # Write data.
 data = [[1.0, 1.0], [2.0, 2.0], [3.0, 3.0]]
 write(es_client, index=index_name, vecs=data)

 # Query an index.
 query_vector = [1.0, 1.0]
 search(es_client, index=index_name, query=query_vector, size=3)

 # Delete an index.
 delete(es_client, index=index_name)
```

### 2.8.2.10 Sample Java Code for Vector Search

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

This section provides a sample of Java code for creating vector indexes, and importing and querying vector data. It shows how to use the client to implement vector search.

### Prerequisites

Add the following Maven dependency based on the actual cluster version. 7.10.2 is used in this example.

```
<dependency>
 <groupId>org.elasticsearch.client</groupId>
 <artifactId>elasticsearch-rest-high-level-client</artifactId>
 <version>7.10.2</version>
</dependency>
```

## Sample Code

```
package org.example;

import org.apache.http.HttpEntity;
import org.apache.http.HttpHost;
import org.apache.http.HttpStatus;
import org.apache.http.auth.AuthScope;
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.client.CredentialsProvider;
import org.apache.http.conn.ssl.NoopHostnameVerifier;
import org.apache.http.impl.client.BasicCredentialsProvider;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.elasticsearch.action.admin.indices.delete.DeleteIndexRequest;
import org.elasticsearch.action.admin.indices.refresh.RefreshRequest;
import org.elasticsearch.action.bulk.BulkRequest;
import org.elasticsearch.action.bulk.BulkResponse;
import org.elasticsearch.action.index.IndexRequest;
import org.elasticsearch.action.search.SearchResponse;
import org.elasticsearch.action.support.master.AcknowledgedResponse;
import org.elasticsearch.client.Request;
import org.elasticsearch.client.RequestOptions;
import org.elasticsearch.client.Response;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;
import org.elasticsearch.client.indices.CreateIndexRequest;
import org.elasticsearch.client.indices.CreateIndexResponse;
import org.elasticsearch.common.settings.Settings;
import org.elasticsearch.common.xcontent.DeprecationHandler;
import org.elasticsearch.common.xcontent.NamedXContentRegistry;
import org.elasticsearch.common.xcontent.XContentParser;
import org.elasticsearch.common.xcontent.XContentType;

import javax.net.ssl.SSLContext;
import javax.net.ssl.TrustManager;
import javax.net.ssl.X509TrustManager;
import java.io.IOException;
import java.security.KeyManagementException;
import java.security.NoSuchAlgorithmException;
import java.security.SecureRandom;
import java.security.cert.X509Certificate;
import java.util.Arrays;
import java.util.List;

public class ClientExample {
 private final RestHighLevelClient client;

 public ClientExample(RestHighLevelClient client) {
 this.client = client;
 }

 //Create a client for accessing non-security clusters.
 public static RestHighLevelClient getClient(List<String> hosts, int port, String scheme) {
 HttpHost[] httpHosts = hosts.stream().map(p -> new HttpHost(p, port,
scheme)).toArray(HttpHost[]::new);
 return new RestHighLevelClient(RestClient.builder(httpHosts));
 }

 //Create a client for accessing security clusters.
 public static RestHighLevelClient getClient(List<String> hosts, int port, String scheme, String user, String
password) {
 final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
 credentialsProvider.setCredentials(AuthScope.ANY, new UsernamePasswordCredentials(user,
password));
 SSLContext sc = null;
 try {
 sc = SSLContext.getInstance("SSL");
 sc.init(null, trustAllCerts, new SecureRandom());
 } catch (KeyManagementException | NoSuchAlgorithmException e) {
```

```

 e.printStackTrace();
 }
 HttpHost[] httpHosts = hosts.stream().map(p -> new HttpHost(p, port,
scheme)).toArray(HttpHost[]::new);
 final SSLIOSessionStrategy sessionStrategy = new SSLIOSessionStrategy(sc,
NoopHostnameVerifier.INSTANCE);
 RestClientBuilder builder = RestClient.builder(httpHosts).setHttpClientConfigCallback(httpClientBuilder
-> {
 httpClientBuilder.disableAuthCaching();
 httpClientBuilder.setSSLStrategy(sessionStrategy);
 return httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
 });
 return new RestHighLevelClient(builder);
}

public static TrustManager[] trustAllCerts = new TrustManager[] {
 new X509TrustManager() {
 @Override
 public void checkClientTrusted(X509Certificate[] chain, String authType) {
 }
 @Override
 public void checkServerTrusted(X509Certificate[] chain, String authType) {
 }
 @Override
 public X509Certificate[] getAcceptedIssuers() {
 return null;
 }
 }
};

Create an index.
public void create(String index) throws IOException {
 CreateIndexRequest request = new CreateIndexRequest(index);
 request.settings(Settings.builder()
 .put("index.vector", true) // Enable the vector feature.
 .put("index.number_of_shards", 1) //Set the number of index shards as needed.
 .put("index.number_of_replicas", 0) //Set the number of index replicas as needed.
);
 String mapping =
 "{\n" +
 " \"properties\": {\n" +
 " \"my_vector\": {\n" +
 " \"type\": \"vector\", \n" +
 " \"indexing\": \"true\", \n" +
 " \"dimension\": \"2\", \n" +
 " \"metric\": \"euclidean\", \n" +
 " \"algorithm\": \"GRAPH\" \n" +
 " } \n" +
 " } \n" +
 "}";
 request.mapping(mapping, XContentType.JSON);
 CreateIndexResponse response = client.indices().create(request, RequestOptions.DEFAULT);
 if (response.isAcknowledged()) {
 System.out.println("create " + index + " success");
 }
}

// Write data. You are advised to keep the batch size under 500 records.
public void write(String index, List<float[]> vectors) throws IOException {
 BulkRequest request = new BulkRequest();
 for (float[] vec : vectors) {
 request.add(new IndexRequest(index).source("my_vector", vec));
 }

 BulkResponse response = client.bulk(request, RequestOptions.DEFAULT);
 if (response.hasFailures()) {
 System.out.println(response.buildFailureMessage());
 } else {
 System.out.println("write bulk to " + index + " success");
 }
}

```

```
}

// Optional. Elasticsearch will refresh it by default.
client.indices().refresh(new RefreshRequest(index), RequestOptions.DEFAULT);
}

// Query vectors.
public void search(String index, float[] query, int size) throws IOException {
 String queryFormat =
 "{\n" +
 " \"size\":%d,\n" +
 " \"query\": {\n" +
 " \"vector\": {\n" +
 " \"my_vector\": {\n" +
 " \"vector\": %s,\n" +
 " \"topk\":%d\n" +
 " }\n" +
 " }\n" +
 " }\n" +
 "}";
 String body = String.format(queryFormat, size, Arrays.toString(query), size);
 Request request = new Request("POST", index + "/_search");
 request.setJsonEntity(body);
 Response response = client.getLowLevelClient().performRequest(request);
 if (response.getStatusLine().getStatusCode() != HttpStatus.SC_OK) {
 System.out.println(response.getEntity()); //Handle the error based on service requirements.
 return;
 }
 // Process the normally returned result based on service requirements.
 HttpEntity entity = response.getEntity();
 XContentType xContentType = XContentType.fromMediaTypeOrFormat("application/json");
 XContentParser parser = xContentType.xContent().createParser(NamedXContentRegistry.EMPTY,
 DeprecationHandler.IGNORE_DEPRECATIONS, entity.getContent());
 SearchResponse searchResponse = SearchResponse.fromXContent(parser);
 System.out.println(searchResponse);
}

// Delete an index.
public void delete(String index) throws IOException {
 DeleteIndexRequest request = new DeleteIndexRequest(index);
 AcknowledgedResponse response = client.indices().delete(request, RequestOptions.DEFAULT);
 if (response.isAcknowledged()) {
 System.out.println("delete " + index + " success");
 }
}

public void close() throws IOException {
 client.close();
}

public static void main(String[] args) throws IOException {
 // For a non-security cluster, run the following:
 RestHighLevelClient client = getClient(Arrays.asList("xx.xx.xx.xx"), 9200, "http");

 /*
 *For a security cluster with HTTPS enabled, run the following:
 * RestHighLevelClient client = getClient(Arrays.asList("xx.xx.xx.xx", "xx.xx.xx.xx"), 9200, "https",
 "user_name", "password");
 *For a security cluster with HTTPS disabled, run the following:
 * RestHighLevelClient client = getClient(Arrays.asList("xx.xx.xx.xx", "xx.xx.xx.xx"), 9200, "http",
 "user_name", "password");
 */

 ClientExample example = new ClientExample(client);
 String indexName = "my_index";

 //Create an index.
 example.create(indexName);
}
```

```
//Write data.
List<float[]> data = Arrays.asList(new float[]{1.0f, 1.0f}, new float[]{2.0f, 2.0f}, new float[]{3.0f, 3.0f});
example.write(indexName, data);

//Query an index.
float[] queryVector = new float[]{1.0f, 1.0f};
example.search(indexName, queryVector, 3);

//Delete an index.
example.delete(indexName);


//Close the client.
example.close();
}
}
```

### 2.8.3 Configuring Decoupled Storage and Compute for an Elasticsearch Cluster

As data grows rapidly, traditional monolithic storage architectures struggle to balance performance and cost efficiency. By decoupling storage and compute, CSS's Elasticsearch clusters allow separate storage of hot data (frequently accessed data) and cold data (rarely accessed historical data). This ensures real-time query performance for hot data while reducing long-term storage costs. The solution is ideal when cold data queries can tolerate a relatively low performance standard, such as in log analytics and auditing scenarios.

#### How the Feature Works

**Figure 2-18** Decoupled storage and compute



The core of decoupled storage and compute is storage tiering based on data temperature (how frequent data is accessed).

- Hot data: frequently accessed data stored on high-speed SSDs, searchable in milliseconds.
- Cold data: rarely accessed data stored in OBS (invisible to users), using Standard storage.

In CSS, indexes may go through the following stages during their entire lifecycle.

- Hot index (open): Writes are allowed. Query latency is in milliseconds.

- Frozen index (freeze): Rarely accessed data is stored in OBS. The index becomes read-only, and query latency is in seconds or even minutes.
- Deleted index: Frozen indexes are deleted periodically to reclaim storage resources.

## Impact on Billing

Upon freezing, indexes are dumped to OBS, and the storage class is Standard. OBS buckets that store frozen indexes are invisible to users—users cannot view or manage them on the OBS console. Note that storing indexes in OBS buckets incurs additional fees, billed on a pay-per-use basis according to consumption of Standard storage capacity. For details, see [Cloud Search Service Price Calculator](#).

## Constraints

- Only Elasticsearch 7.6.2 and 7.10.2 clusters support decoupled storage and compute.
- While an index is being frozen, the system sets the index to the read-only state. Even after the data in the index is dumped to OBS, the index remains read-only and no data can be written into the index.
- During index freezing, the data in it can still be queried. After the freezing is complete, the index is closed and then re-opened. During this transition period, the index becomes unsearchable, and the cluster may be in the **red** state temporarily. The index becomes searchable again after being re-opened.
- After an index is frozen, its data is dumped to an OBS bucket, and any index data is deleted from local disks. A frozen index has an increased query latency. During an aggregation operation, the latency becomes even longer because the query is complex and a large amount of data needs to be retrieved.
- A frozen index with data already dumped to OBS cannot be unfrozen. That is, a read-only index cannot be rolled back to writable.
- In the case of two clusters configured with read/write splitting, the secondary cluster cannot set storage-compute decoupling for indexes synchronized from the primary cluster.
- When an index is frozen, its data is dumped to an OBS bucket. This process will consume network bandwidth. Make sure data transmission between your cluster and OBS does not exhaust the maximum bandwidth and QPS supported by OBS. For details, see [OBS Use Constraints](#). If these constraints are violated, the performance of queries that involve OBS will deteriorate. For example, the speed of restoring shards and querying data will be slow.

## Logging In to Kibana

Log in to Kibana and go to the command execution page. Elasticsearch clusters support multiple access methods. This topic uses Kibana as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.



4. In the left navigation pane, choose **Dev Tools**.

## Freezing an Index

Freezing an index means to move cold data from solid state drives (SSDs) to OBS, freeing up SSD storage for hot, frequently accessed data. A frozen index becomes read-only.

1. Freeze an index.

Run the following command to dump the data of a specified index to OBS and set its storage class to Standard:

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

where, **index\_name** indicates the name of the index to be frozen.

Example response:

```
{
 "freeze_uuid": "pdsRgUtSTymVDWR_HoTGFw"
}
```

where, **freeze\_uuid** indicates the ID of the index freezing task. After an index freezing request is submitted, an asynchronous task is started, and an asynchronous task ID is returned, which you can use to query the progress of the index freezing task.

2. Check the progress of the index freezing task.

Run the following command to check the index freezing progress by specifying the task ID:

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

Example response:

```
{
 "stage": "STARTED",
 "shards_stats": {
 "INIT": 0,
 "FAILURE": 0,
 "DONE": 0,
 "STARTED": 1,
 "ABORTED": 0
 },
 "indices": {
 "data1": [
 {
 "uuid": "7OS-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
}
}
]
}
}
```

**Table 2-78** Response parameters

| Parameter    | Description                                                                                                                                                                                                                                                                                                                         |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| stage        | Status of the index freezing task.<br>Value range: <ul style="list-style-type: none"><li>• <b>INIT</b>: The task has just started or is being initialized.</li><li>• <b>FAILURE</b>: failed</li><li>• <b>DONE</b>: complete</li><li>• <b>STARTED</b>: started</li><li>• <b>ABORTED</b>: canceled. This field is reserved.</li></ul> |
| shards_stats | Number of shards in each state.                                                                                                                                                                                                                                                                                                     |
| indices      | Status details of each index, as described in <a href="#">Table 2-79</a> .                                                                                                                                                                                                                                                          |

**Table 2-79** indices parameters

| Parameter            | Description                                                              |
|----------------------|--------------------------------------------------------------------------|
| uuid                 | ID of the index freezing task.                                           |
| index                | Index and shard information.                                             |
| start_ms             | Task start time.                                                         |
| end_ms               | Task end time. If the task is still ongoing, <b>-1</b> is displayed.     |
| total_time           | How long the task has been ongoing.                                      |
| total_time_in_millis | How long the task has been ongoing, in milliseconds.                     |
| stage                | Shard status.                                                            |
| failure              | Task failure cause. If no failure has occurred, <b>null</b> is returned. |
| size.total_bytes     | Total bytes to be frozen.                                                |
| size.finished_bytes  | Number of bytes that have been frozen.                                   |

| Parameter               | Description                                                                                                          |
|-------------------------|----------------------------------------------------------------------------------------------------------------------|
| size.percent            | Freezing progress in terms of bytes that have been frozen vs total bytes to freeze.                                  |
| file.total_bytes        | Total number of documents to be frozen.                                                                              |
| file.finished_bytes     | Number of documents that have been frozen.                                                                           |
| file.percent            | Freezing progress in terms of the number of documents that have been frozen vs. total number of documents to freeze. |
| rate_limit.paused_times | Number of times freezing has been or was suspended due to rate limiting.                                             |
| rate_limit.paused_nanos | Duration of freezing task suspension due to rate limiting, in nanoseconds.                                           |

When index freezing is complete, the settings parameters will be returned, as described in [Table 2-80](#).

**Table 2-80** settings parameters

| Parameter             | Description                                                              |
|-----------------------|--------------------------------------------------------------------------|
| index.frozen_low_cost | Frozen index. The value is fixed to <b>true</b> .                        |
| index.blocks.write    | Writes disallowed to a frozen index. The value is fixed to <b>true</b> . |
| index.store.type      | Index storage type. The value is fixed to <b>obs</b> .                   |

### 3. Query the list of frozen indexes.

#### NOTE

Here, only some of the key parameters for querying the frozen index list are described. For more information about the API, see [cat indices API](#).

Run the following command to query the frozen index list based on index freezing status:

```
GET _cat/freeze_indices?stage={STAGE}
```

STAGE indicates the index freezing status. The value can be:

- start: Freezing has started and is ongoing.
- done: Freezing has completed.
- unfreeze: Freezing has not started.
- Empty or any other value: Freezing is ongoing or has completed.

Example response:

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

## Optimizing Cache Settings for Frozen Indexes

After an index is frozen, its data is dumped to OBS. To reduce direct data retrieval from OBS and thus improve query performance, some data is cached in the cluster. Data that is requested for the first time is directly retrieved from OBS. The retrieved data is then cached in the cluster memory. In response to subsequent queries, the system searches the cache first. CSS allows you to query cache statistics for frozen indexes that are stored in OBS buckets. You can also reset the cache status and modify cache configuration.

### 1. Query cache statistics for frozen indexes stored in OBS buckets.

- Run the following command to query cache statistics for frozen indexes on all nodes:

```
GET _frozen_stats
```

- Run the following command to query cache statistics for frozen indexes on specified nodes:

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

where, **node\_id** indicates a cluster node ID.

The following is an example of the returned information for querying frozen index cache statistics on all nodes:

```
{
 "_nodes": {
 "total": 1, // Total number of nodes
 "successful": 1, // Successful nodes
 "failed": 0 // Failed nodes
 },
 "cluster_name": "css-zzz1", //Cluster name
 "nodes": {
 "7uwKO38RRoaON37YsXhCYw": {
 "name": "css-zzz1-ess-esn-2-1", //Node name
 "transport_address": "10.0.0.247:9300", //Node transport address
 "host": "10.0.0.247", //Node host
 "ip": "10.0.0.247", //Node IP address
 "block_cache": {
 "default": {
 "type": "memory", //Cache type. memory indicates in-memory cache.
 "block_cache_capacity": 8192, //Cache capacity
 "block_cache_blocksize": 8192, //Single block size in the cache, in bytes. In the example, the
 block size is 8 KB.
 "block_cache_size": 12, //Cache capacity used
 "block_cache_hit": 14, //Number of cache hits
 "block_cache_miss": 0, //Number of cache misses
 "block_cache_eviction": 0, //Number of cache evictions
 "block_cache_store_fail": 0 //Number of cache storage failures, which occur when the
 cache is full.
 }
 },
 },
 "obs_stats": {
 "list": {
 "obs_list_count": 17, //Number of times the OBS list API was called.
 "obs_list_ms": 265, //Total length of time spent calling the OBS list API.
 "obs_list_avg_ms": 15 //Average time spent calling the OBS list API.
 },
 "get_meta": {
 "obs_get_meta_count": 79, //Number of times the OBS get metadata API was called.
 "obs_get_meta_ms": 183, //Total length of time spent calling the OBS get metadata API.
 "obs_get_meta_avg_ms": 2 //Average time spent calling the OBS get metadata API.
 },
 "get_obj": {
 "obs_get_obj_count": 12, //Number of times the OBS get object API was called.
 "obs_get_obj_ms": 123, //Total length of time spent calling the OBS get object API.
 "obs_get_obj_avg_ms": 10 //Average time spent calling the OBS get object API.
 },
 },
 },
}
```

```
"put_obj" : {
 "obs_put_obj_count" : 12, //Number of times the OBS put object API was called.
 "obs_put_obj_ms" : 2451, //Total length of time spent calling the OBS put object API.
 "obs_put_obj_avg_ms" : 204 //Average time spent calling the OBS put object API.
},
"obs_op_total" : {
 "obs_op_total_ms" : 3022, //Total length of time spent calling OBS APIs.
 "obs_op_total_count" : 120, //Total number of times calling OBS APIs.
 "obs_op_avg_ms" : 25 //Average time spent calling OBS APIs.
}
},
"reader_cache" : {
 "hit_count" : 0,
 "miss_count" : 1,
 "load_success_count" : 1,
 "load_exception_count" : 0,
 "total_load_time" : 291194714,
 "eviction_count" : 0
}
}
}
```

2. Reset the cache status for frozen indexes stored in OBS buckets.

Run the following command to reset the cache status for frozen indexes:

```
POST _frozen_stats/reset
```

This command is used to debug performance issues. If you reset the cache status and then run the cache query command, you can check the accurate cache command status. It is not advisable to use this command during runtime.

Example response:

```
{
 "_nodes" : {
 "total" : 1,
 "successful" : 1,
 "failed" : 0
 },
 "cluster_name" : "Es-0325-007_01",
 "nodes" : {
 "mqTdk2YRSPyOSXfesREFSg" : {
 "result" : "ok"
 }
 }
}
```

3. Modify cache configuration for frozen indexes.

Elasticsearch accesses different types of files using different methods. The cache system supports multi-level caches and uses blocks of different sizes to cache different types of files. For example, typically, a large number of small blocks are used to cache .fdx and .tip files, while a small number of large blocks are used to cache .fdt files. You can modify the cache configuration based on service needs. [Table 2-81](#) describes the parameters that can be configured.

**Table 2-81** Cache configuration parameters

| Parameter                                            | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| low_cost.obs.blockcache.names                        | Array   | A list of names of multi-level caches. The cache system supports multi-level caches for data of different access granularities.<br>There is a default cache named <b>default</b> . If you configure this parameter, the value must include at least the <b>default</b> cache in addition to other custom cache names.<br>Default value: <b>default</b>                                                                         |
| low_cost.obs.blockcache.<NAME>.type                  | ENUM    | Storage type of the cache named <NAME>.<br>Value range: <ul style="list-style-type: none"><li>• <b>memory</b>: The cache uses in-memory storage, which means it consumes memory resources.</li><li>• <b>file</b>: The cache uses disk storage. Its capacity is limited by available disk space. In this case, you are advised to use ultra-high I/O disks to improve cache performance.</li></ul> Default value: <b>memory</b> |
| low_cost.obs.blockcache.<NAME>.blockshift            | Integer | Block size in the cache named <NAME> (in bytes).<br>Formula: $2^{\text{blockshift}}$ . For example, if the value is 16, the block size is $2^{16}$ bytes, that is, 65536 bytes or 64 KB.<br>Value: number of left-shifts<br>Default value: 13 (equivalent to 8 KB)                                                                                                                                                             |
| low_cost.obs.blockcache.<NAME>.bank.count            | Integer | Number of partitions in the cache named <NAME>. Each partition is an independent unit in the cache.<br>Default value: <b>1</b>                                                                                                                                                                                                                                                                                                 |
| low_cost.obs.blockcache.<NAME>.number.blocks.perbank | Integer | Number of blocks in each partition of the cache named <NAME>.<br>Default value: <b>8192</b>                                                                                                                                                                                                                                                                                                                                    |
| low_cost.obs.blockcache.<NAME>.exclude.file.types    | Array   | List of file name extensions for files that cannot be cached in the cache named <NAME>.<br>If a file's extension is not in <b>exclude.file.types</b> or <b>&lt;NAME&gt;.file.types</b> (if configured), the default cache policy applies.                                                                                                                                                                                      |

| Parameter                                         | Type   | Description                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| low_cost.obs.blockcache.<NAME>.file.types         | Array  | List of file name extensions for files that can be cached in the cache named <NAME>.<br><br>If a file's extension is not in <b>exclude.file.types</b> or <NAME>.file.types (if configured), the default cache policy applies.                                                                                                                                       |
| index.frozen.obs.max_bytes_per_sec                | String | Maximum rate at which the documents of frozen indexes are uploaded to OBS.<br>Format: <value><unit>. The unit can be B (byte), KB, MB, or GB.<br>The change takes effect immediately.<br>Default value: <b>150 MB</b>                                                                                                                                               |
| low_cost.obs.index.upload.threshold.use.multipart | String | Multipart upload threshold. In the process of dumping frozen indexes to OBS, if the size of a single document exceeds this threshold, a multipart upload is performed.<br>Format: <value><unit>. The unit can be B (byte), KB, MB, or GB.<br>Multipart upload is supported by OBS.<br>0 or not configured: Multipart upload may be disabled.<br>Default value: 1 GB |

The following is a typical cache configuration example. A two-level cache system is used: **default** and **large**. The **default** cache has a total of 30 x 4096 64-KB blocks. It is used to cache non-.fdt files. The **large** cache has 5 x 1000 2-MB 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"]
```

## Improving Query Performance for Frozen Indexes



### WARNING

Only Elasticsearch 7.6.2 and 7.10.2 clusters created after February 2023 support this feature.

When frozen indexes are queried on the **Discover** page of Kibana for the first time, all data needs to be retrieved from OBS because there is no data in the cache. If a large number of documents need to be returned, it takes a long time to retrieve the corresponding time fields and document metadata from OBS. By caching this part of data in the cluster, query performance can be improved significantly. This is how CSS improves query performance for frozen indexes. Local cache settings are preset. You can review them and modify them as needed.

1. Modify local cache settings for frozen indexes.

**Table 2-82** Local cache settings

| Configurat<br>ion Item                        | Typ<br>e    | scop<br>e | Can Be<br>Change<br>d<br>Dynam<br>ically | Description                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------|-------------|-----------|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| low_cost.lo<br>cal_cache.<br>max.capacit<br>y | Inte<br>ger | nod<br>e  | Yes                                      | <p>Maximum number of cold data caches that can be made available on each node. (Each shard corresponds to a cache.)</p> <p>Value range: 10–5000</p> <p>Default value: <b>500</b></p> <ul style="list-style-type: none"><li>• If the heap memory usage remains high, decrease this value.</li><li>• If the value of <b>load_overflow_count</b> keeps increasing rapidly, increase this value.</li></ul> |



| Configurat<br>ion Item                            | Typ<br>e    | scop<br>e | Can Be<br>Change<br>d<br>Dynam<br>ically | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------|-------------|-----------|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.low_c<br>ost.local_ca<br>che.thresho<br>ld  | Inte<br>ger | inde<br>x | Yes                                      | <p>Threshold for enabling local caching of cold data.</p> <ul style="list-style-type: none"><li>If in an index, the percentage of <b>date</b> fields is lower than this value, cold data of the <b>date</b> type will be cached locally. Otherwise, the local cache will not be used.</li><li>If date fields account for the vast majority of the total data volume of the current index, you should not use this setting.</li></ul> <p>Value range: 0–100<br/>Default value: 50<br/>Unit: %</p> |
| index.low_c<br>ost.local_ca<br>che.evict_ti<br>me | Strin<br>g  | inde<br>x | Yes                                      | <p>Retention duration for cold data in the local cache. The value is determined by <b>index.frozen_date</b> (time of index freezing). If <b>index.frozen_date</b> is unavailable, the value is determined by the index creation time.</p> <p>You are advised to adjust the retention duration based on your disk usage.</p> <p>Value range: 1 to 365 days<br/>Default value: <b>30d</b><br/>Unit: days</p>                                                                                       |

The following provides some examples, where **index\_name** indicates the modified index.

- 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 {index_name}/_settings
{
```

```
"index.low_cost.local_cache.threshold":20
}
```

- Run the following command to modify **index.low\_cost.local\_cache.evict\_time**:

```
PUT {index_name}/_settings
{
 "index.low_cost.local_cache.evict_time":"7d"
}
```

## 2. Query local cache statistics for frozen indexes.

- Query local cache statistics for frozen indexes on all nodes:  
GET /\_frozen\_stats/local\_cache
- Query local cache statistics for frozen indexes on specified nodes:  
GET /\_frozen\_stats/local\_cache/{node\_id}

where, **node\_id** indicates a cluster node ID.

Example response:

```
{
 "_nodes" : {
 "total" : 1,
 "successful" : 1,
 "failed" : 0
 },
 "cluster_name" : "test",
 "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.
 "get_hit_count" : 562, //Total number of hits in the local cold data cache.
 "get_miss_count" : 0, //Total number of local cold data cache misses.
 "get_total_ns" : 43849200, //Total duration for retrieving data from the local cold data
cache.
 "get_avg_ns" : 78023 //Average duration for retrieving data from the local cold
data cache.
 },
 "load_stats" : {
 "load_count" : 2, //Number of times cold data was loaded from the local cache
 "load_total_ms" : 29, //Total duration for loading cold data from the local cache
 "load_avg_ms" : 14, //Average duration for loading cold data from the local cache
 "load_fail_count" : 0, //Number of failures for loading cold data from the local
cache
 "load_overflow_count" : 0 //Number of times the local cold data cache exceeds the
cache pool size.
 },
 "reload_stats" : {
 "reload_count" : 0, //Number of times the local cold data cache was regenerated.
 "reload_total_ms" : 0, //Total duration for regenerating the local cold data cache.
 "reload_avg_ms" : 0, //Average duration for regenerating the local cold data cache.
 "reload_fail_count" : 0 //Number of failures in regenerating the local cold data
cache.
 },
 "init_stats" : {
 "init_count" : 0, //Number of times the local cold data cache was initialized.
 "init_total_ms" : 0, //Total duration for initializing the local cold data cache.
 "init_avg_ms" : 0, //Average duration for initializing the local cold data cache.
 "init_fail_count" : 0 //Number of failures in initializing the local cold data cache.
 }
 }
 }
 }
}
```

## Querying the Real-Time Rates of OBS in Handling Cold Data

### WARNING

Only Elasticsearch 7.6.2 and 7.10.2 clusters created after February 2023 support this feature.

To help you understand how the storage-compute decoupling plug-in is working with OBS, an API for collecting statistics on the real-time rates of OBS has been added, and the real-time rates are recorded in the index **.freeze\_obs\_rate-YYYY.mm.dd**.

Calculation method: The average OBS operation rates in the last 5 seconds are calculated every 5 seconds.

The system index **.freeze\_obs\_rate-YYYY.mm.dd** records statistics on OBS real-time operation rates, helping you understand relevant trends about the OBS resources that store cold data. The default retention period of this index is 30 days.

#### 1. Querying the real-time rates of OBS in handling cold data.

- Run the following command to query the real-time OBS rates on all nodes:

```
GET _frozen_stats/obs_rate
```

- Run the following command to query the real-time OBS rates on specified nodes:

```
GET _frozen_stats/obs_rate/{node_id}
```

where, **node\_id** indicates a cluster node ID.

Example response:

```
{
 "_nodes": {
 "total": 1,
 "successful": 1,
 "failed": 0
 },
 "cluster_name": "test",
 "nodes": {
 "dfLDvcSwTJ-fkillT2zE3A": {
 "name": "node-1",
 "transport_address": "127.0.0.1:9300",
 "host": "127.0.0.1",
 "ip": "127.0.0.1",
 "update_time": 1671777600482, // Time when the current statistics were updated.
 "obs_rate": {
 "list_op_rate": 0.0, // Rate of OBS list operations. Unit: times/s.
 "get_meta_op_rate": 0.0, // Rate of OBS get meta operations. Unit: times/s.
 "get_obj_op_rate": 0.0, // Rate of OBS get operations. Unit: times/s.
 "put_op_rate": 0.0, // Rate of OBS put operations. Unit: times/s.
 "obs_total_op_rate": 0.0, // Rate of all OBS operations. Unit: times/s.
 "obs_upload_rate": "0.0 MB/s", // OBS data upload rate. Unit: MB/s.
 "obs_download_rate": "0.0 MB/s" // OBS data download rate. Unit: MB/s.
 }
 }
 }
}
```

- Modify the retention period of the **.freeze\_obs\_rate-YYYY.mm.dd** index that stores the OBS real-time rates. The default retention period of this index is 30 days.

Run the following command to change the index retention period to seven days:

```
PUT _cluster/settings
{
 "persistent": {
 "low_cost.obs_rate_index.evict_time": "7d"
 }
}
```

**Table 2-83** Configuration items

| Configuration Item                 | Type   | scope | Can Be Changed Dynamically | Description                                                                                                                                                                                           |
|------------------------------------|--------|-------|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| low_cost.obs_rate_index.evict_time | String | node  | Yes                        | The retention period of the <b>.freeze_obs_rate-YYYY.mm.dd</b> index. <ul style="list-style-type: none"><li>Value range: 1 to 365 days</li><li>Default value: <b>30d</b></li><li>Unit: days</li></ul> |

## Related Documents

- To ensure acceptable performance for cold data queries (for example, when historical data is queried periodically), use a tiered storage solution. For details, see [Switching Between Hot and Cold Storage for an Elasticsearch Cluster](#).
- Decoupled storage and compute can be automated via index lifecycle management policies. For details, see [Decoupling Index Storage and Compute in an Elasticsearch Cluster Through Index Lifecycle Management](#).

## 2.8.4 Configuring Flow Control 2.0 for an Elasticsearch Cluster

### Scenario

Flow Control 2.0 controls traffic at the node level. You can configure blacklists and whitelists per node, the maximum concurrent HTTPS connections allowed, and the maximum HTTP connections allowed. You can also configure node heap memory-based flow control that uses a backpressure mechanism, or one-click traffic blocking. CSS can also collect statistics on node access IP addresses and URLs. When backpressure is enabled, large requests will be rejected by nodes when their heap memory usage is high. This mechanism prevents nodes from breaking down and reduces the risk of node unavailability.

**Table 2-84** Flow control policies

| Policy                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Details                                                   |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| HTTP/HTTPS flow control    | <ul style="list-style-type: none"><li>You can control cluster access by client IP address or subnet through the HTTP/HTTPS blacklist or whitelist. If an IP address is in the blacklist, the client is disconnected right away and all its requests are rejected. The whitelist takes precedence over the blacklist. If a client IP address is on both the blacklist and whitelist, requests from it will not be rejected.</li><li>Flow control based on concurrent HTTP/HTTPS connections limits the total number of HTTP/HTTPS connections to a node per second.</li><li>Flow control based on new HTTP/HTTPS connections limits the number of new connections to a node.</li></ul> | <a href="#">Enabling HTTP/HTTPS Flow Control per Node</a> |
| Memory flow control        | Memory flow control limits write traffic based on the node heap memory. It uses a backpressure mechanism to ask the client to slow down or stop sending requests. In the meantime, it triggers garbage collection to reclaim resources, and continues to process requests based on heap memory available.                                                                                                                                                                                                                                                                                                                                                                             | <a href="#">Enabling Memory Flow Control</a>              |
| Request sampling           | Request sampling can record the access of client IP addresses and the type of requests from the client. Based on such statistics, you can identify the access traffic of specific client IP addresses and analyze write and query traffic from them.                                                                                                                                                                                                                                                                                                                                                                                                                                  | <a href="#">Enabling Request Sampling</a>                 |
| One-click traffic blocking | One-click traffic blocking blocks all client connections to a node. This, however, does not include connections for Kibana access, CSS O&M, or monitoring APIs. The purpose is to protect cluster nodes in the face of traffic spikes or to quickly restore clusters.                                                                                                                                                                                                                                                                                                                                                                                                                 | <a href="#">Enable One-Click Traffic Blocking</a>         |
| Flow control               | Flow control provides an independent API for checking traffic statistics, including the number of existing client connections as well as client connections where backpressure has been applied. You can evaluate the flow control threshold and analyze the cluster load based on these statistics.                                                                                                                                                                                                                                                                                                                                                                                  | <a href="#">Viewing Flow Control Information</a>          |

| Policy                  | Description                                                                                                                                                                 | Details                                          |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------|
| Access logging          | 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 load based on the access logs. | <a href="#">Enabling and Viewing Access Logs</a> |
| Access logging in files | Any cluster access is recorded in the <b>{<i>Cluster name</i>_access_log.log}</b> file. You can use the log backup function to view detailed access logs on OBS.            | <a href="#">Enabling Access Logging in Files</a> |

## Constraints

- Elasticsearch 7.6.2 and Elasticsearch 7.10.2 clusters created after January 2023 support Flow Control 2.0 only, whereas those created before that support Flow Control 1.0 only.
- Flow control may hurt the performance of some nodes.
- If flow control is enabled, user requests that exceed the flow control threshold will be rejected.
- Enabling memory flow control may hurt the performance of some search requests or cause some Kibana search requests to fail.
- Enabling access logging may hurt cluster performance.
- Memory flow control is based on request paths. Avoid configuring too many paths or paths that are too long, as they may hurt cluster performance.

## Logging In to Kibana

Log in to Kibana and go to the command execution page. Elasticsearch clusters support multiple access methods. This topic uses Kibana as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
4. In the left navigation pane, choose **Dev Tools**.

## Enabling HTTP/HTTPS Flow Control per Node

1. Run the following command to enable HTTP/HTTPS flow control for cluster nodes:

```
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
 }
}
```

```
}
}
```

**Table 2-85** Configuration items for HTTP/HTTPS flow control

| Configuration Item          | Type         | Description                                                                                                                                                                                                                                                             |
|-----------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.http.enabled    | Boolean      | Whether to enable HTTP/HTTPS flow control. HTTP/HTTPS flow control is disabled by default. Enabling it may affect node access performance.<br>Value: <b>true</b> or <b>false</b><br>Default value: <b>false</b>                                                         |
| flowcontrol.http.allow      | List<String> | IP address whitelist.<br>It can contain multiple IP addresses and subnet masks, or lists of IP addresses. Use commas (,) to separate different items. Example: <b>xx.xx.xx.xx/24,xx.xx.xx.xx/24</b> , or <b>xx.xx.xx.xx,xx.xx.xx.xx</b> .<br>The default value is null. |
| flowcontrol.http.deny       | List<String> | IP address blacklist.<br>It can contain multiple IP addresses and subnet masks, or lists of IP addresses. Use commas (,) to separate different items.<br>The default value is null.                                                                                     |
| flowcontrol.http.concurrent | Integer      | Maximum concurrent HTTP/HTTPS connections.<br>Default value: Number of available cores on a node x 600.                                                                                                                                                                 |
| flowcontrol.http.newconnect | Integer      | Maximum new connections that can be created for HTTP/HTTPS requests per second.<br>Default value: Number of available cores on a node x 200.                                                                                                                            |

| Configuration Item             | Type    | Description                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.http.warmup_period | Integer | Time required for the HTTP/HTTPS connection setup speed to reach the maximum. If <b>flowcontrol.http.newconnect</b> is set to <b>100</b> and <b>flowcontrol.http.warmup_period</b> is set to <b>5000ms</b> , it indicates the system can create up to 100 connections per second 5 seconds later.<br>Value range: 0–10000<br>Unit: ms<br>Default value: <b>0</b> |

 **NOTE**

If all parameters are set to **null**, they will be restored to their default values.

2. Run the following command to disable HTTP/HTTPS flow control for cluster nodes:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.http.enabled": false
 }
}
```

## Enabling Memory Flow Control

1. Run the following command to enable memory flow control:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.memory.enabled": true,
 "flowcontrol.memory.heap_limit": "80%"
 }
}
```

**Table 2-86** Configuration items for memory flow control

| Configuration Item         | Type    | Description                                                                                                                                                                                                         |
|----------------------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.memory.enabled | Boolean | Whether to enable memory flow control. After this function is enabled, the memory usage is continuously monitored.<br>Value: <ul style="list-style-type: none"><li>• true</li><li>• false (default value)</li></ul> |



| Configuration Item                   | Type      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.memory.heap_limit        | String    | <p>Maximum heap memory usage of a node that is used as a threshold for triggering backpressure for flow control.</p> <p>Value range: 10%–100%</p> <p>Default value: <b>90%</b></p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>The default value 90% of <b>flowcontrol.memory.heap_limit</b> is a conservative threshold. When the heap memory usage is greater than 90%, the system stops reading large requests that exceed 64 KB from the client until heap memory usage decreases. Once the heap memory usage decreases to 85%, client data equivalent to <b>5% x maximum heap memory capacity</b> can be read. If the heap memory usage stays above 90% for a long time, client requests cannot be processed. In this case, a GC algorithm is triggered to perform garbage collection until the heap memory usage drops below the threshold.</li> <li>Generally, you can set the <b>flowcontrol.memory.heap_limit</b> threshold to 80% or less to ensure that the node has reserved some heap memory for operations besides data writing, such as Elasticsearch query and segment merge.</li> </ul> |
| flowcontrol.holding.in_flight_factor | Float     | <p>Backpressure release factor, which works similarly to the circuit breaker parameter <b>network.breaker.inflight_requests.overhead</b>. With the memory usage exceeding the limit, a larger value of this parameter indicates stronger backpressure, in which case, write traffic will be limited.</p> <p>Value range: <math>\geq 0.5</math></p> <p>Default value: <b>1.0</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| flowcontrol.holding.max              | TimeValue | <p>Maximum delay of each request. If the delay exceeds the value of this parameter, backpressure may be stopped or the request connection may be disconnected. For details, see the setting of <b>flowcontrol.holding.max_strategy</b>.</p> <p>Value range: <math>\geq 15s</math></p> <p>Default value: <b>60s</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

| Configuration Item               | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.holding.max_strategy | String  | <p>The policy applied after the maximum delay time is exceeded.</p> <p>The value can be:</p> <ul style="list-style-type: none"><li>• <b>keep</b> (default value): If the heap memory is still high, continue the backpressure. The server determines when to execute the request based on the real-time memory.</li><li>• <b>soft</b>: The requests will be executed even if the heap memory usage is still high. The <b>inFlight</b> circuit breaker will determine whether to execute or reject the requests.</li><li>• <b>hard</b>: If the heap memory usage is still high, requests will be discarded and the client connections will be disconnected.</li></ul> |
| flowcontrol.memory.once_free_max | String  | <p>Maximum memory that can be made available at a time for a suspended request queue. This parameter helps prevent a cluster from becoming completely unavailable due to low memory availability under high pressure.</p> <p>Value range: 1%–50%</p> <p>Default value: 5%</p>                                                                                                                                                                                                                                                                                                                                                                                        |
| flowcontrol.memory.nudges_gc     | Boolean | <p>Whether to trigger garbage collection to ensure write stability when the write pressure is too high. (The backpressure connection pool is checked every second. The write pressure is considered high if all the existing connections are blocked and new write requests cannot be accepted.) The value can be:</p> <ul style="list-style-type: none"><li>• true (default value)</li><li>• false</li></ul>                                                                                                                                                                                                                                                        |

 **NOTE**

If all parameters are set to **null**, they will be restored to their default values.

2. Run the following command to disable memory flow control:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.memory.enabled": false
 }
}
```

## Enabling Request Sampling

1. Run the following command to enable request sampling:

```
PUT _cluster/settings
{
 "transient": {
 "flowcontrol.log.access.enabled": true
 }
}
```

**Table 2-87** Configuration items for request sampling

| Configuration Item             | Type    | Description                                                                                                                                                                                                                                                                                                     |
|--------------------------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.log.access.enabled | Boolean | Whether to collect statistics on client IP addresses that accessed the cluster recently and the number of requests from them, including the quantities of bulk write, search, and msearch requests.<br>The value can be: <ul style="list-style-type: none"><li>• true</li><li>• false (default value)</li></ul> |
| flowcontrol.log.access.count   | Integer | Number of client IP addresses that accessed a cluster recently. The IP address statistics switches control whether to collect request type statistics and whether to enable logging.<br>Value range: 0–100<br>Default value: <b>10</b>                                                                          |

When enabled, you can check relevant statistics through [Viewing Flow Control Information](#).

### NOTE

If all parameters are set to **null**, they will be restored to their default values.

2. Run the following command to disable request sampling:

```
PUT _cluster/settings
{
 "persistent": {
 "flowcontrol.log.access.enabled": false
 }
}
```

## Enable One-Click Traffic Blocking

1. Run the following command to enable one-click traffic blocking:

```
PUT _cluster/settings
{
 "persistent": {
 "flowcontrol.break.enabled": true
 }
}
```

2. Run the following command to disable one-click traffic blocking:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.break.enabled": false
 }
}
```

Viewing Flow Control Information

- Check the flow control status of all nodes.

```
GET /_nodes/stats/filter/v2
```

- View the flow control details of all nodes.

```
GET /_nodes/stats/filter/v2?detail
```

- View the flow control status of a specific node.

```
GET /_nodes/{nodeId}/stats/filter/v2
```

**{nodeId}** indicates the ID of the node you want to check.

Example response:

```
{
 "_nodes" : {
 "total" : 1,
 "successful" : 1,
 "failed" : 0
 },
 "cluster_name" : "css-xxxx",
 "nodes" : {
 "d3qnVipPTtSoadkV0LQEkA" : {
 "name" : "css-xxxx-ess-esn-1-1",
 "host" : "192.168.x.x",
 "timestamp" : 1672236425112,
 "flow_control" : {
 "http" : {
 "current_connect" : 52,
 "rejected_concurrent" : 0,
 "rejected_rate" : 0,
 "rejected_black" : 0,
 "rejected_breaker" : 0
 },
 "access_items" : [
 {
 "remote_address" : "10.0.0.x",
 "search_count" : 0,
 "bulk_count" : 0,
 "other_count" : 4
 }
]
 },
 "holding_requests" : 0
 }
 }
}
```

Table 2-88 Response parameters

| Parameter       | Description                                                                                                                                                                                                                                                  |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| current_connect | Number of HTTP connections of a node, which is recorded regardless of whether flow control is enabled. This value is equivalent to the <b>current_open</b> value of <b>GET /_nodes/stats/http</b> API. It shows the current client connections of each node. |

| Parameter           | Description                                                                                                                                                                                                                                       |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rejected_concurrent | Number of concurrent connections rejected during flow control.<br>This parameter is counted only when <b>flowcontrol.http.enabled</b> is set to <b>true</b> . This count will not be cleared when flow control is disabled.                       |
| rejected_rate       | Number of new connections rejected during flow control.<br>This parameter is counted only when <b>flowcontrol.http.enabled</b> is set to <b>true</b> . This count will not be cleared when flow control is disabled.                              |
| rejected_black      | Number of new connections rejected by a preconfigured blacklist during flow control.<br>This parameter is counted only when <b>flowcontrol.http.enabled</b> is set to <b>true</b> . This count will not be cleared when flow control is disabled. |
| rejected_breaker    | Number of new connections rejected during one-click traffic blocking.<br>This parameter is counted only when <b>flowcontrol.break.enabled</b> is set to <b>true</b> . This count will not be cleared when one-click traffic blocking is disabled. |
| access_items        | IP addresses of clients that recently accessed the cluster.<br>The value is determined by <b>flowcontrol.log.access.count</b> .                                                                                                                   |
| remote_address      | Remote access IP addresses and the number of requests from them.                                                                                                                                                                                  |
| search_count        | Number of times a client accessed a database using <b>_search</b> and <b>_msearch</b> .                                                                                                                                                           |
| bulk_count          | Number of times a client accessed a database using <b>_bulk</b> .                                                                                                                                                                                 |
| other_count         | Number of times a client accessed a database using other request methods.                                                                                                                                                                         |
| holding_requests    | Number of connections to the current node where writes are halted due to flow control.                                                                                                                                                            |

## Enabling and Viewing Access Logs

1. Run the following command to enable access logging:
  - Enable access logging for all nodes in a cluster.  
`PUT /_access_log?duration_limit=30s&capacity_limit=1mb`
  - Enable access logging for a specified node in a cluster.

```
PUT /_access_log/{nodeId}?duration_limit=30s&capacity_limit=1mb
```

**{nodeId}** indicates the node ID.

**Table 2-89** Configuration items for configuring access logging

| Configuration Item | Type   | Description                                                                                                                                                                           |
|--------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| duration_limit     | String | Maximum duration of access log records. When this duration is reached, the recording stops.<br>Value range: 10 to 120<br>Unit: s<br>Default value: <b>30</b>                          |
| capacity_limit     | String | Maximum memory size for recording access logs. When the size of an access log reaches this value, access logging stops.<br>Value range: 1 to 5<br>Unit: MB<br>Default value: <b>1</b> |

 **NOTE**

- Access logging stops when either **duration\_limit** or **capacity\_limit** is reached.
  - If all parameters are set to **null**, they will be restored to their default values.
2. Run the following command 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 specified node in a cluster  
GET /\_access\_log/{nodeId}

**{nodeId}** indicates the node ID.

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 2-90** Response parameters

| Parameter | Description                                                                                              |
|-----------|----------------------------------------------------------------------------------------------------------|
| name      | Node name                                                                                                |
| 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 <a href="#">Table 2-91</a> . |

**Table 2-91** access

| Parameter      | Description                                      |
|----------------|--------------------------------------------------|
| time           | Request time                                     |
| remote_address | Source IP address and port number in the request |
| url            | Original URL of the request                      |
| method         | Method corresponding to the request path         |
| content        | Request content                                  |

3. Run the following commands to delete access logs.

API for deleting access logs for all nodes:

```
DELETE /_access_log
```

## Enabling Access Logging in Files

Typically you record access logs in files to locate faults. After faults are rectified, you should disable it.

1. Run the following command to enable access logging in files:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.log.file.enabled": true
 }
}
```

**Table 2-92** Configuration items for enabling access logging in files

| Parameter                    | Type    | Description                                                                                                                                                                                                                                                                                |
|------------------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.log.file.enabled | Boolean | Whether to record the details of each request in the access log file. The log file name is <b>Cluster name_access_log.log</b> . You can check this file only through the log backup function.<br>Value: <ul style="list-style-type: none"><li>true</li><li>false (default value)</li></ul> |

2. Run the following command to disable access logging in files:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.log.file.enabled": false
 }
}
```

## 2.8.5 Configuring Flow Control 1.0 for an Elasticsearch Cluster

### Scenario

Flow Control 1.0 controls traffic at the node level. You can configure blacklists and whitelists per node, the maximum concurrent HTTP connections allowed, the maximum HTTP connections allowed, the maximum heap memory used by specific request paths, and the maximum CPU usage. You can block access in one click, and collect statistics on IP addresses and URLs accessing the nodes. If flow control is enabled, requests will be blocked at the entry, which alleviates the cluster pressure in high-concurrency scenarios and reduces the likelihood of unavailability issues.



**Table 2-93** Flow control policies

| Policy                                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Details                                                         |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|
| HTTP/HTTPS flow control                | <ul style="list-style-type: none"><li>You can control cluster access by client IP address or subnet through the HTTP/HTTPS blacklist or whitelist. If an IP address is in the blacklist, the client is disconnected right away and all its requests are rejected. The whitelist takes precedence over the blacklist. If a client IP address is on both the blacklist and whitelist, requests from it will not be rejected.</li><li>Flow control based on concurrent HTTP/HTTPS connections limits the total number of HTTP/HTTPS connections to a node per second.</li><li>Flow control based on new HTTP/HTTPS connections limits the number of new connections to a node.</li></ul> | <a href="#">Enabling HTTP/HTTPS Flow Control per Node</a>       |
| Memory flow control                    | Memory flow control limits request paths based on heap memory usage. You can configure a whitelist for memory flow control, a global memory usage threshold, and per-path heap memory thresholds. The global memory flow control threshold takes precedence over the memory threshold of a single path. Paths on the whitelist are exempt from memory flow control.                                                                                                                                                                                                                                                                                                                   | <a href="#">Enabling Memory Flow Control</a>                    |
| Global path whitelist for flow control | You can configure the global path whitelist for flow control as required when you need to use custom plug-ins.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <a href="#">Adding a Global Path Whitelist for Flow Control</a> |
| Request sampling                       | 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 and analyze access traffic by client IP addresses and request paths.                                                                                                                                                                                                                                                                                                                                                                                                                                             | <a href="#">Enabling Request Sampling</a>                       |
| One-click traffic blocking             | One-click traffic blocking blocks all client connections to a node. This, however, does not include connections for Kibana access or Elasticsearch monitor APIs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <a href="#">Enable One-Click Traffic Blocking</a>               |
| Flow control                           | Flow control provides an independent API for checking traffic statistics and records the number API calls. You can evaluate the flow control threshold and analyze the cluster load based on these statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <a href="#">Viewing Flow Control Information</a>                |

| Policy                  | Description                                                                                                                                                                 | Details                                          |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------|
| Access logging          | 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 load based on the access logs. | <a href="#">Enabling and Viewing Access Logs</a> |
| Access logging in files | Any cluster access is recorded in the <b>{<i>Cluster name</i>_access_log.log}</b> file. You can use the log backup function to view detailed access logs on OBS.            | <a href="#">Enabling Access Logging in Files</a> |
| CPU flow control        | You can configure the node CPU usage threshold to limit inbound traffic for each node.                                                                                      | <a href="#">Enabling CPU Flow Control</a>        |

## Constraints

- Elasticsearch 7.6.2 and Elasticsearch 7.10.2 clusters created after January 2023 support Flow Control 2.0 only, whereas those created before that support Flow Control 1.0 only.
- Flow control may hurt the performance of some nodes.
- If flow control is enabled, user requests that exceed the flow control threshold will be rejected.
- Enabling memory flow control may hurt the performance of some search requests or cause some Kibana search requests to fail.
- Enabling access logging may hurt cluster performance.
- Memory flow control and CPU flow control are based on request paths. Avoid configuring too many paths or paths that are too long, as they may hurt cluster performance.

## Enabling HTTP/HTTPS Flow Control per Node

1. Run the following command to enable HTTP/HTTPS flow control for cluster nodes:

```
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
 }
}
```

**Table 2-94** Configuration items for HTTP/HTTPS flow control

| Configuration Item             | Type         | Description                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.http.enabled       | Boolean      | Whether to enable HTTP/HTTPS flow control. HTTP/HTTPS flow control is disabled by default. Enabling it may affect node access performance.<br>Value: <b>true</b> or <b>false</b><br>Default value: <b>false</b>                                                                                                                                                  |
| flowcontrol.http.allow         | List<String> | IP address whitelist.<br>It can contain multiple IP addresses and subnet masks, or lists of IP addresses. Use commas (,) to separate different items. Example: <b>xx.xx.xx.xx/24,xx.xx.xx.xx/24</b> , or <b>xx.xx.xx.xx,xx.xx.xx.xx</b> .<br>The default value is null.                                                                                          |
| flowcontrol.http.deny          | List<String> | IP address blacklist.<br>It can contain multiple IP addresses and subnet masks, or lists of IP addresses. Use commas (,) to separate different items.<br>The default value is null.                                                                                                                                                                              |
| flowcontrol.http.concurrent    | Integer      | Maximum concurrent HTTP/HTTPS connections.<br>Default value: Number of available cores on a node x 600.                                                                                                                                                                                                                                                          |
| flowcontrol.http.newconnect    | Integer      | Maximum new connections that can be created for HTTP/HTTPS requests per second.<br>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 <b>flowcontrol.http.newconnect</b> is set to <b>100</b> and <b>flowcontrol.http.warmup_period</b> is set to <b>5000ms</b> , it indicates the system can create up to 100 connections per second 5 seconds later.<br>Value range: 0-10000<br>Unit: ms<br>Default value: <b>0</b> |

 **NOTE**

If all parameters are set to **null**, they will be restored to their default values.

2. Run the following command to disable HTTP/HTTPS flow control for cluster nodes:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.http.enabled": false
 }
}
```

## Enabling Memory Flow Control

1. Enable memory flow control.

- Run the following command to enable memory flow control:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.memory.enabled": true,
 "flowcontrol.memory.allow_path": "/index/_search",
 "flowcontrol.memory.heap_limit": "85%"
 }
}
```

- Run the following command to enable memory flow control for a single request path.

Configure the heap memory usage threshold for a request path. You can configure priorities for 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%"
 }
 }
 }
}
```

**Table 2-95** Configuration items for memory flow control

| Configuration Item            | 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.<br>Value: <b>true</b> or <b>false</b><br>Default value: <b>false</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| flowcontrol.memory.allow_path | List<String> | Request path whitelist for memory flow control.<br>Whitelisted paths are exempt from memory flow control. Wildcard characters are supported. By default, query APIs controlled by the cluster are exempt from memory flow control. This prevents the failure to query cluster information when the memory usage reaches the threshold.<br>Example: <ul style="list-style-type: none"><li>• "flowcontrol.memory.allow_path": "/index/_search",</li><li>• "flowcontrol.memory.allow_path": "/index*/_search",</li><li>• "flowcontrol.memory.allow_path": ["/index/_search", "/index1/_bulk"],</li></ul> A maximum of 10 paths can be configured. A path can contain up to 32 characters.<br>The default value is null. |

| Configuration Item               | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.memory.heap_limit    | String | <p>Maximum global heap memory usage of a node before flow control is triggered. The value cannot be less than 10% of the heap memory.</p> <p>Value range: 10%–100%</p> <p>Default value: <b>90%</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| flowcontrol.memory.*.filter_path | String | <p>Paths under memory flow control.</p> <p>The default value is <b>**</b>, indicating all paths. If <b>flowcontrol.memory.heap_limit</b> is configured and <b>flowcontrol.memory.*.filter_path</b> 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 <b>flowcontrol.memory.allow_path</b> and <b>flowcontrol.memory.*.filter_path</b>, the requests from the path will be allowed.</p> <p>For example, if <b>flowcontrol.memory.allow_path</b> and <b>flowcontrol.memory.*.filter_path</b> are both set to <b>abc/_search</b>, then <b>abc/_search</b> will not be under flow control.</p> <p>Maximum length: <b>32</b> characters</p> |

| Configuration Item              | Type   | Description                                                                                                                                                                       |
|---------------------------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.memory.*.heap_limit | String | Heap memory usage threshold of request paths. If the heap memory usage exceeds the threshold, flow control will be triggered.<br>Value range: 0–100%<br>Default value: <b>90%</b> |

 **NOTE**

If all parameters are set to **null**, they will be restored to their default values.

2. Run the following command to disable memory flow control:

- Run the following command to delete memory flow control for a single request path.

Before disabling memory flow control, you need delete per-request path memory flow control configuration.

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.memory.enabled": true,
 "flowcontrol.memory": {
 "flowcontrol_search": {
 "filter_path": null,
 "heap_limit": null
 }
 }
 }
}
```

- Run the following command to disable memory flow control:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.memory.enabled": false
 }
}
```

## Adding a Global Path Whitelist for Flow Control

Run the following command to add a global path whitelist for flow control:

```
PUT _cluster/settings
{
 "persistent": {
 "flowcontrol.path.white_list": "xxx"
 }
}
```

**Table 2-96** Configuration items for a global path whitelist for flow control

| Configuration Item          | Type         | Description                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.path.white_list | List<String> | <p>Paths that are exempt from 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><b>NOTE</b><br/>You are advised not to configure this parameter, unless required by plug-ins.</p> |

 **NOTE**

If all parameters are set to **null**, they will be restored to their default values.

## Enabling Request Sampling

1. Run the following command to enable request sampling:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.statics.enabled": true,
 "flowcontrol.statics.threshold": 100,
 "flowcontrol.statics.sample_frequency": 50
 }
}
```

**Table 2-97** Configuration items for request sampling

| Configuration Item          | Type    | Description                                                                                                                                                           |
|-----------------------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.statics.enabled | Boolean | <p>Whether to enable request sampling. Request sampling may affect node performance.</p> <p>Value: <b>true</b> or <b>false</b></p> <p>Default value: <b>false</b></p> |



| Configuration Item                   | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.statics.threshold        | Integer | <p>Number of recent access requests whose statistics are collected. The value <b>100</b> indicates that statistics will be collected on the 100 IP addresses and 100 URLs that are most frequently accessed.</p> <p>Minimum value: <b>10</b></p> <p>Maximum value: <b>1000</b></p> <p>Default value: <b>100</b></p> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>• The IP address statistics and URL sampling statistics are cached based on their access time. If the number of cached records reaches the threshold configured using <b>flowcontrol.statics.threshold</b>, the earliest records will be deleted.</li><li>• In URL sampling, an access path is uniquely identified by its URL hash.</li></ul> |
| flowcontrol.statics.sample_frequency | Integer | <p>Path sampling frequency. If this parameter is set to <b>100</b>, samples are collected from every 100 requests.</p> <p>Minimum value: <b>50</b></p> <p>Default value: <b>100</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

 **NOTE**

If all parameters are set to **null**, they will be restored to their default values.

2. Run the following command to disable request sampling:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.statics.enabled": false
 }
}
```

## Enable One-Click Traffic Blocking

1. Run the following command to enable one-click traffic blocking:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.break.enabled": true
 }
}
```

2. Run the following command to disable one-click traffic blocking:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.break.enabled": false
 }
}
```

```
}
}
```

## Viewing Flow Control Information

- Check the flow control status of all nodes.

```
GET /_nodes/stats/filter
```

- View the flow control status of a specific node.

```
GET /_nodes/{nodeId}/stats/filter
```

**{nodeId}** indicates the ID of the node you want to check.

Example response:

### NOTE

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.

```
{
 "_nodes" : {
 "total" : 1,
 "successful" : 1,
 "failed" : 0
 },
 "cluster_name" : "css-flowcontroller",
 "nodes" : {
 "ElBRNCMbTj6L1C-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",
 "count" : 1
 }
]
 }
 }
}
```

```
 "remote_address" : "/10.0.0.198:16763",
 "count" : 1
 }
]
}
}
```

**Table 2-98** 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 <b>current_open</b> of the <b>GET /_nodes/stats/http</b> API but is smaller, because whitelisted IP addresses and internal node IP addresses are not counted.                                              |
| rejected_concurrent | Number of concurrent connections rejected during HTTP flow control. Disabling HTTP flow control does not clear this record.                                                                                                                                                                                                                        |
| rejected_new        | Number of new connections rejected during HTTP flow control. Disabling HTTP flow control does not clear this record.                                                                                                                                                                                                                               |
| rejected_deny       | Number of requests rejected based on the blacklist during HTTP flow control. Disabling HTTP flow control does not clear this record.                                                                                                                                                                                                               |
| 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 <b>allow_path</b> whitelist are not recorded. If <b>allow_path</b> is set to <b>**</b> , 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 <b>allow_path</b> whitelist are not recorded. If <b>allow_path</b> is set to <b>**</b> , 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 <a href="#">Table 2-99</a> .                                                                                                                                                                                                                                                             |
| 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 <a href="#">Table 2-100</a> .                                                                                                                                                                            |

**Table 2-99** ip\_address

| Parameter | Description                                   |
|-----------|-----------------------------------------------|
| ip        | Source IP address for accessing the node.     |
| method    | Number of access requests from an IP address. |

**Table 2-100** url\_sample

| Parameter      | Description                                      |
|----------------|--------------------------------------------------|
| url            | Request URL                                      |
| method         | Method corresponding to the request path         |
| remote_address | Source IP address and port number in the request |
| count          | How many times a path is sampled                 |

## Enabling and Viewing Access Logs

1. Run the following command to enable access logging:
  - Enable access logging for all nodes in a cluster.  
`PUT /_access_log?duration_limit=30s&capacity_limit=1mb`
  - Enable access logging for a specified node in a cluster.  
`PUT /_access_log/{nodeId}?duration_limit=30s&capacity_limit=1mb`**{nodeId}** indicates the node ID.

**Table 2-101** Configuration items for configuring access logging

| Configuration Item | Type   | Description                                                                                                                                                                           |
|--------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| duration_limit     | String | Maximum duration of access log records. When this duration is reached, the recording stops.<br>Value range: 10 to 120<br>Unit: s<br>Default value: <b>30</b>                          |
| capacity_limit     | String | Maximum memory size for recording access logs. When the size of an access log reaches this value, access logging stops.<br>Value range: 1 to 5<br>Unit: MB<br>Default value: <b>1</b> |

 **NOTE**

- Access logging stops when either **duration\_limit** or **capacity\_limit** is reached.
- If all parameters are set to **null**, they will be restored to their default values.

## 2. Run the following command 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 specified node in a cluster

```
GET /_access_log/{nodeId}
```

**{nodeId}** indicates the node ID.

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 2-102** Response parameters

| Parameter | Description                                                                                               |
|-----------|-----------------------------------------------------------------------------------------------------------|
| name      | Node name                                                                                                 |
| 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 <a href="#">Table 2-103</a> . |

**Table 2-103** access

| Parameter      | Description                                      |
|----------------|--------------------------------------------------|
| time           | Request time                                     |
| remote_address | Source IP address and port number in the request |
| url            | Original URL of the request                      |
| method         | Method corresponding to the request path         |
| content        | Request content                                  |

3. Run the following commands to delete access logs.

API for deleting access logs for all nodes:

```
DELETE /_access_log
```

## Enabling Access Logging in Files

Typically you record access logs in files to locate faults. After faults are rectified, you should disable it.

1. Run the following command to enable access logging in files:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.log.file.enabled": true
 }
}
```

**Table 2-104** Configuration items for enabling access logging in files

| Parameter                    | Type    | Description                                                                                                                                                                                                                                                                                    |
|------------------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.log.file.enabled | Boolean | Whether to record the details of each request in the access log file. The log file name is <b>Cluster name_access_log.log</b> . You can check this file only through the log backup function.<br>Value: <ul style="list-style-type: none"><li>• true</li><li>• false (default value)</li></ul> |

2. Run the following command to disable access logging in files:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.log.file.enabled": false
 }
}
```

## Enabling CPU Flow Control

1. Run the following command to enable CPU flow control:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.cpu.enabled": true,
 "flowcontrol.cpu.percent_limit": 80,
 "flowcontrol.cpu.allow_path": ["index/_search"]
 }
}
```

**Table 2-105** Configuration items for configuring access logging

| Configuration Item            | Type         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.cpu.enabled       | Boolean      | Whether to enable CPU flow control.<br>The value can be: <ul style="list-style-type: none"><li>• true</li><li>• false (default value)</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| flowcontrol.cpu.percent_limit | Integer      | Maximum CPU usage (%) of a node before flow control is triggered.<br>Default value: <b>90</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| flowcontrol.cpu.allow_path    | List<String> | Path whitelist for CPU flow control.<br>The paths specified using this parameter are exempt from CPU flow control. By default, query APIs controlled by the cluster are exempt from CPU flow control. This prevents the failure to query cluster information when the CPU usage reaches the threshold.<br>Example: <ul style="list-style-type: none"><li>• "flowcontrol.memory.allow_path": "/index/_search",</li><li>• "flowcontrol.memory.allow_path": "/index*/_search",</li><li>• "flowcontrol.memory.allow_path": ["/index/_search", "/index1/_bulk"],</li></ul> A maximum of 10 paths can be configured.<br>A path can contain up to 32 characters.<br>The default value is null. |

### NOTE

If all parameters are set to **null**, they will be restored to their default values.

2. Run the following command to disable CPU flow control:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.cpu.enabled": false
 }
}
```

```
}
}
```

## 2.8.6 Configuring Large Query Isolation for an Elasticsearch Cluster

### Scenario

You can isolate query requests that consume a large amount of memory or take a long period of time. This way, you ensure service availability for other requests. If the heap memory usage of a node is too high, an interrupt control program will be triggered to terminate a large query based on the policies you configured. You can also configure a global query timeout duration. Long queries will be intercepted by an Elasticsearch-native cancel API.

Large query isolation can effectively solve the following problems and improve the search performance of clusters:

- A small number of queries occupy large chunks of node heap memory, resulting in frequent Garbage Collection (GC) and even out of memory (OOM) exceptions.
- Frequent GC causes node disconnections. As a result, queries cannot get response and may fail.
- The CPU usage is high due to heavy query load, affecting online services.

### Constraints

Only Elasticsearch 7.6.2 and Elasticsearch 7.10.2 clusters support large query isolation.

### Logging In to Kibana

Log in to Kibana and go to the command execution page. Elasticsearch clusters support multiple access methods. This topic uses Kibana as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
4. In the left navigation pane, choose **Dev Tools**.

### Enabling Large Query Isolation

Large query isolation is enabled by default, while global query timeout is disabled by default. If you enable them, the configuration will take effect immediately.

Run the following commands to enable large query isolation and global query timeout:

```
PUT _cluster/settings
{
 "persistent": {
 "search.isolator.enabled": true,
 "search.isolator.time.enabled": true
```



```
}
}
```

The two features each has an independent switch. [Table 2-106](#) describes their parameters.

**Table 2-106** Parameters for configuring large query isolation and global query timeout

| Switch                       | Parameter                                                                                             | Description                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------|-------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.enabled      | search.isolator.memory.task.limit<br>search.isolator.time.management                                  | Thresholds for identifying a single shard query task as a large query.                                                                                                                                                                                                                                                                                                                |
|                              | search.isolator.memory.pool.limit<br>search.isolator.memory.heap.limit<br>search.isolator.count.limit | Resource usage thresholds for isolation. If the resource usage of a query task exceeds one of these thresholds, the task will be paused.<br><b>NOTE</b><br><b>search.isolator.memory.heap.limit</b> defines the limit on the heap memory consumed by write, query, and other operations of a node. If this limit is exceeded, large query tasks in the isolation pool will be paused. |
|                              | search.isolator.strategy<br>search.isolator.strategy.ratio                                            | Policy for selecting query tasks to pause in the isolation pool.                                                                                                                                                                                                                                                                                                                      |
| search.isolator.time.enabled | search.isolator.time.limit                                                                            | Global timeout for query tasks.                                                                                                                                                                                                                                                                                                                                                       |

### Configuring Large Query Isolation Thresholds

- Run the following commands to configure thresholds for defining a large query task:

```
PUT _cluster/settings
{
 "persistent": {
 "search.isolator.memory.task.limit": "50MB",
 "search.isolator.time.management": "10s"
 }
}
```

**Table 2-107** Parameters for configuring large query isolation thresholds

| Parameter                         | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.memory.task.limit | String | <p>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.</p> <ul style="list-style-type: none"><li>Value range: <b>0b</b> to the maximum heap memory of a node</li><li>Default value: <b>50MB</b></li></ul> <p><b>NOTE</b><br/>You can run the following command to query the current heap memory and the maximum heap memory of a cluster:</p> <p><b>GET _cat/nodes?&amp;h=id,ip,port,r,ramPercent,ramCurrent,heapMax,heapCurrent</b></p> |
| search.isolator.time.management   | String | <p>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.</p> <ul style="list-style-type: none"><li>Value range: <b>≥ 0ms</b></li><li>Default value: <b>10s</b></li></ul>                                                                                                                                                                                                                                                                          |

- Configure the resource usage thresholds for triggering the isolation of query tasks.

```
PUT _cluster/settings
{
 "persistent": {
 "search.isolator.memory.pool.limit": "50%",
 "search.isolator.memory.heap.limit": "90%",
 "search.isolator.count.limit": 1000
 }
}
```

**Table 2-108** Parameters for configuring isolation pool thresholds

| Parameter                         | Type   | Description                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.memory.pool.limit | String | <p>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.</p> <ul style="list-style-type: none"><li>Value range: <b>0.0</b> to <b>100.0%</b></li><li>Default value: <b>50%</b></li></ul> |

| Parameter                         | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. <ul style="list-style-type: none"><li>• Value range: <b>0.0</b> to <b>100.0%</b></li><li>• Default value: <b>90%</b></li></ul>                                                                        |
| search.isolator.count.limit       | Integer | 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. <ul style="list-style-type: none"><li>• Value range: <b>10-50000</b></li><li>• Default value: <b>1000</b></li></ul> |

 **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.

- Configure a policy for selecting query tasks to pause in the isolation pool.

```
PUT _cluster/settings
{
 "persistent": {
 "search.isolator.strategy": "fair",
 "search.isolator.strategy.ratio": "0.5%"
 }
}
```

**Table 2-109** Parameters for configuring a query cancelation policy

| Parameter                      | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.strategy       | String | <p>Policy for selecting large queries when the interrupt control program is triggered. The selected queries will be interrupted.</p> <p><b>NOTE</b><br/>The large query isolation pool is checked every second until the heap memory is within the safe range.</p> <p>Values: <b>fair</b>, <b>mem-first</b>, or <b>time-first</b></p> <ul style="list-style-type: none"><li>• <b>mem-first</b>: The query task that uses the most heap memory in the isolation pool is interrupted.</li><li>• <b>time-first</b>: The query task that has been running for the longest time in the isolation pool is interrupted.</li><li>• <b>fair</b>: If the difference between the heap memory of shard queries is smaller than <math>Maximum\_heap\_memory \times search.isolator.strategy.ratio</math>, the query that takes the longest time should be interrupted. Otherwise, the query that uses the most heap memory is interrupted.</li></ul> <p>Default value: <b>fair</b></p> |
| search.isolator.strategy.ratio | String | <p>Threshold of the <b>fair</b> policy. This parameter takes effect only if <b>search.isolator.strategy</b> is set to <b>fair</b>. 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> <ul style="list-style-type: none"><li>• Value range: <b>0.0</b> to <b>100.0%</b></li><li>• Default value: <b>1%</b></li></ul>                                                                                                                                                                                                                                                                                                                                                                                               |

## Configuring the Global Query Timeout

Run the following command to set the global timeout of query tasks:

```
PUT _cluster/settings
{
 "persistent": {
 "search.isolator.time.limit": "120s"
 }
}
```

**Table 2-110** Description

| Parameter                  | Type   | Description                                                                                                                                                                                                             |
|----------------------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.time.limit | String | Global query timeout duration. Any query task that exceeds this duration will be canceled. <ul style="list-style-type: none"><li>Value range: <math>\geq 0\text{ms}</math></li><li>Default value: <b>120s</b></li></ul> |

## Configuring the Maximum Number of Log Records for Canceled Query Requests

Run the following command to set the maximum number of log records kept for canceled query requests:

```
PUT _cluster/settings
{
 "persistent": {
 "search.isolator.log.count": "100"
 }
}
```

| Parameter                 | Data Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.log.count | Integer   | <p>Maximum number of records of canceled query requests that can be recorded in the memory.</p> <ul style="list-style-type: none"><li>Value range: 0–5000</li><li>Default value: <b>100</b></li></ul> <p><b>NOTE</b></p> <p>You can use the following APIs to query canceled requests:</p> <ul style="list-style-type: none"><li>GET <code>/_isolator_metrics</code>: Queries all nodes.</li><li>GET <code>/_isolator_metrics/{nodeId}</code>: Queries a single node.</li><li>GET <code>/_isolator_metrics?detailed</code>: Queries request cancellation details of all nodes.</li><li>GET <code>/_isolator_metrics/{nodeId}?detailed</code>: Queries request cancellation details of a single node.</li></ul> <p>In the commands above, <b>nodeId</b> indicates the node ID.</p> |

## 2.8.7 Configuring Enhanced Aggregation for an Elasticsearch Cluster

CSS Elasticsearch clusters enhance aggregation performance in the face of large data volumes by leveraging vectorization and optimized clustering, enabling faster analytics and decision-making in complex situations.

### How the Feature Works

Enhanced aggregation works by pre-sorting and physically clustering data using carefully selected sorting and clustering keys, thereby minimizing data scanning and computational overheads during aggregation.

- **Sorting key:** A field (such as timestamps) used to physically order documents on disk, ensuring that documents with identical or similar sorting key values are stored contiguously.
- **Clustering key:** A subset of fields from the sorting key that groups related documents into contiguous physical blocks. This way, aggregations can process contiguous data blocks, instead of scattered documents.

**Table 2-111** Common scenarios for enhanced aggregation

| Scenario                                               | Description                                                                                                                                                                                                                                 |
|--------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Aggregation of low-cardinality fields                  | Aggregates fields that have a small number of unique values. In this case, grouping is used. One example is counting the number of orders per city. A clustering operation can easily achieve this purpose.                                 |
| Aggregation of high-cardinality fields                 | Aggregates fields that have a large number of unique values. Histogram aggregation is typically used. One example is counting the hourly visits. A clustering key can be used to accelerate data aggregation by a specified scope or range. |
| Hybrid aggregation of low- and high-cardinality fields | Groups and aggregates low-cardinality fields first (for example, orders per city), and then creates a histogram using high-cardinality fields (for example, timestamps). Multi-level clustering improves the efficiency of hybrid queries.  |

### Constraints

Only Elasticsearch 7.10.2 supports enhanced aggregation.

### Aggregation of Low-cardinality Fields

Generally, low-cardinality fields use grouping as a way of aggregation. With appropriate sorting keys, grouping prepares the data for batch vectorization.

For example, to aggregate two low-cardinality fields **city** and **product**, perform the following steps:

1. Run the following command to set the index **testindex**:

```
PUT testindex
{
 "mappings": {
 "properties": {
 "date": {
 "type": "date",
 "format": "yyyy-MM-dd"
 },
 "city": {
 "type": "keyword"
 },
 "product": {
 "type": "keyword"
 },
 "trade": {
 "type": "double"
 }
 }
 },
 "settings": {
 "index": {
 "search": {
 "turbo": {
 "enabled": "true"
 }
 }
 },
 "sort": {
 "field": [
 "city",
 "product"
]
 },
 "cluster": {
 "field": [
 "city",
 "product"
]
 }
 }
}
```

**Table 2-112** Parameters for low-cardinality field aggregation

| Parameter                  | Type    | Description                                                                                                                                                                                                                                                                                          |
|----------------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.search.turbo.enabled | Boolean | Whether to enable enhanced aggregation. Normally, enhanced aggregation must be enabled where aggregations are used.<br>The value can be: <ul style="list-style-type: none"><li>● <b>false</b> (default): Disable enhanced aggregation.</li><li>● <b>true</b>: Enable enhanced aggregation.</li></ul> |

| Parameter           | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.sort.field    | Array of strings | <p>Specify sorting keys. Sorting keys are fields used to sequence or rank documents.</p> <p>You can specify one or multiple fields as sorting keys. When multiple fields are specified, they will apply in the sequence in which they are specified. Documents are first ranked by the first field, then the initial result set is ranked by the second field, and so on.</p> <p>Value range: The value must be fields contained in the index.</p> |
| index.cluster.field | Array of strings | <p>Specify clustering keys. Clustering keys determine which documents are collected into the same clusters.</p> <p>During an aggregation operation, documents in the same cluster can be processed in batches, significantly enhancing aggregation performance.</p> <p>Constraint: The clustering keys must be a subset of the sorting keys.</p> <p>Value range: The value must be fields contained in the index.</p>                              |

2. Run the following command to import sample data.

```
PUT /_bulk
{ "index": { "_index": "testindex", "_id": "1" } }
{ "date": "2025-01-01", "city": "cityA", "product": "books", "trade": 3000.0 }
{ "index": { "_index": "testindex", "_id": "2" } }
{ "date": "2025-01-02", "city": "cityA", "product": "books", "trade": 1000.0 }
{ "index": { "_index": "testindex", "_id": "3" } }
{ "date": "2025-01-01", "city": "cityA", "product": "bottles", "trade": 100.0 }
{ "index": { "_index": "testindex", "_id": "4" } }
{ "date": "2025-01-02", "city": "cityA", "product": "bottles", "trade": 300.0 }
{ "index": { "_index": "testindex", "_id": "5" } }
{ "date": "2025-01-01", "city": "cityB", "product": "books", "trade": 7000.0 }
{ "index": { "_index": "testindex", "_id": "6" } }
{ "date": "2025-01-02", "city": "cityB", "product": "books", "trade": 1000.0 }
```

3. Run the following query command to aggregate low-cardinality fields.

```
POST testindex/_search
{
 "size": 0,
 "aggs": {
 "groupby_city": {
 "terms": {
 "field": "city"
 }
 },
 "aggs": {
 "groupby_product": {
 "terms": {
 "field": "product"
 }
 },
 "aggs": {
 "avg_trade": {
 "avg": {
 "field": "trade"
 }
 }
 }
 }
 }
}
```



```
}
}
}
}
}
}
}
```

## Aggregation of High-cardinality Fields

High-cardinality fields commonly use histogram aggregation, which facilitates data processing per range or scope.

For example, to aggregate the typical high-cardinality field **date**, perform the following steps:

1. Run the following command to set the index **testindex**:

```
PUT testindex
{
 "mappings": {
 "properties": {
 "timestamp": {
 "type": "date",
 "format": "yyyy-MM-dd"
 },
 "score": {
 "type": "double"
 }
 }
 },
 "settings": {
 "index": {
 "search": {
 "turbo": {
 "enabled": "true"
 }
 },
 "sort": {
 "field": [
 "timestamp"
]
 }
 }
 }
}
```

**Table 2-113** Parameters for high-cardinality field aggregation

| Parameter                  | Type    | Description                                                                                                                                                                                                                                                                                          |
|----------------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.search.turbo.enabled | Boolean | Whether to enable enhanced aggregation. Normally, enhanced aggregation must be enabled where aggregations are used.<br>The value can be: <ul style="list-style-type: none"><li>● <b>false</b> (default): Disable enhanced aggregation.</li><li>● <b>true</b>: Enable enhanced aggregation.</li></ul> |

| Parameter        | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.sort.field | Array of strings | <p>Specify sorting keys. Sorting keys are fields used to sequence or rank documents.</p> <p>You can specify one or multiple fields as sorting keys. When multiple fields are specified, they will apply in the sequence in which they are specified. Documents are first ranked by the first field, then the initial result set is ranked by the second field, and so on.</p> <p>Value range: The value must be fields contained in the index.</p> |

2. Run the following command to import sample data.

```
PUT /_bulk
{ "index": { "_index": "testindex", "_id": "1" } }
{ "date": "2025-01-01", "score": "12.0" }
{ "index": { "_index": "testindex", "_id": "2" } }
{ "date": "2025-01-02", "score": "24.0" }
{ "index": { "_index": "testindex", "_id": "3" } }
{ "date": "2025-01-01", "score": "53.0" }
{ "index": { "_index": "testindex", "_id": "4" } }
{ "date": "2025-01-02", "score": "22.0" }
{ "index": { "_index": "testindex", "_id": "5" } }
{ "date": "2025-01-01", "score": "99.0" }
{ "index": { "_index": "testindex", "_id": "6" } }
{ "date": "2025-01-02", "score": "26.0" }
```

3. Run the following query command to aggregate the high-cardinality field.  
This query groups the **date** field using a histogram and then calculates the average score.

```
POST testindex/_search?pretty
{
 "size": 0,
 "aggs": {
 "groupbytime": {
 "date_histogram": {
 "field": "date",
 "calendar_interval": "day"
 },
 "aggs": {
 "avg_score": {
 "avg": {
 "field": "score"
 }
 }
 }
 }
 }
}
```

## Hybrid Aggregation of Low- and High-cardinality Fields

Where low-cardinality and high-cardinality fields are mixed, first groups and aggregates low-cardinality fields, and then aggregates high-cardinality fields using histograms.

For example, to first group the low-cardinality field **city**, then group the low-cardinality field **product**, and then group the high-cardinality field **date** into a histogram, perform the following steps:

1. Run the following command to set the index **testindex**:

```
PUT testindex
{
 "mappings": {
 "properties": {
 "date": {
 "type": "date",
 "format": "yyyy-MM-dd"
 },
 "city": {
 "type": "keyword"
 },
 "product": {
 "type": "keyword"
 },
 "trade": {
 "type": "double"
 }
 }
 },
 "settings": {
 "index": {
 "search": {
 "turbo": {
 "enabled": "true"
 }
 }
 },
 "sort": {
 "field": [
 "city",
 "product",
 "date"
]
 },
 "cluster": {
 "field": [
 "city",
 "product"
]
 }
 }
}
```

**Table 2-114** Parameters for hybrid aggregation of low- and high-cardinality fields

| Parameter                  | Type    | Description                                                                                                                                                                                                                                                                                          |
|----------------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.search.turbo.enabled | Boolean | Whether to enable enhanced aggregation. Normally, enhanced aggregation must be enabled where aggregations are used.<br>The value can be: <ul style="list-style-type: none"><li>● <b>false</b> (default): Disable enhanced aggregation.</li><li>● <b>true</b>: Enable enhanced aggregation.</li></ul> |

| Parameter           | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.sort.field    | Array of strings | <p>Specify sorting keys. Sorting keys are fields used to sequence or rank documents.</p> <p>You can specify one or multiple fields as sorting keys. When multiple fields are specified, they will apply in the sequence in which they are specified. Documents are first ranked by the first field, then the initial result set is ranked by the second field, and so on.</p> <p>Constraint: High-cardinality fields must be among the sorting keys, and must follow the last low-cardinality field.</p> <p>Value range: The value must be fields contained in the index.</p> |
| index.cluster.field | Array of strings | <p>Specify clustering keys. Clustering keys determine which documents are collected into the same clusters.</p> <p>During an aggregation operation, documents in the same cluster can be processed in batches, significantly enhancing aggregation performance.</p> <p>Constraint: The clustering keys must be a subset of the sorting keys.</p> <p>Value range: The value must be fields contained in the index.</p>                                                                                                                                                         |

2. Run the following command to import sample data.

```
PUT /_bulk
{ "index": { "_index": "testindex", "_id": "1" } }
{ "date": "2025-01-01", "city": "cityA", "product": "books", "trade": 3000.0 }
{ "index": { "_index": "testindex", "_id": "2" } }
{ "date": "2025-01-02", "city": "cityA", "product": "books", "trade": 1000.0 }
{ "index": { "_index": "testindex", "_id": "3" } }
{ "date": "2025-01-01", "city": "cityA", "product": "bottles", "trade": 100.0 }
{ "index": { "_index": "testindex", "_id": "4" } }
{ "date": "2025-01-02", "city": "cityA", "product": "bottles", "trade": 300.0 }
{ "index": { "_index": "testindex", "_id": "5" } }
{ "date": "2025-01-01", "city": "cityB", "product": "books", "trade": 7000.0 }
{ "index": { "_index": "testindex", "_id": "6" } }
{ "date": "2025-01-02", "city": "cityB", "product": "books", "trade": 1000.0 }
```

3. Run the following query command to perform a hybrid aggregation of low- and high-cardinality fields.

```
POST testindex/_search
{
 "size": 0,
 "aggs": {
 "groupby_region": {
 "terms": {
 "field": "city"
 },
 "aggs": {
 "groupby_host": {
 "terms": {
 "field": "product"
 }
 }
 }
 }
 }
}
```

```
 },
 "aggs": {
 "groupby_timestamp": {
 "date_histogram": {
 "field": "date",
 "interval": "day"
 }
 },
 "aggs": {
 "avg_score": {
 "avg": {
 "field": "trade"
 }
 }
 }
 }
 }
}
```

## Performance Testing

### Test environment

- Dataset: esrally nyc\_taxi
- Cluster specifications: 4U16G, 100 GB, high I/O x 3 nodes

### Test Procedure

1. Create an index template in the cluster, specify sorting keys, and disable enhanced aggregation.

```
PUT /_template/nyc_taxi
{
 "template": "nyc_taxi*",
 "settings": {
 "index.search.turbo.enabled": false,
 "index.sort.field": "dropoff_datetime",
 "number_of_shards": 3,
 "number_of_replicas": 0
 }
}
```

2. Use esrally to test the nyc\_taxi dataset and obtain the result when enhanced aggregation is disabled.
3. Create another index template in the same cluster, specify sorting keys, and enable enhanced aggregation.

```
PUT /_template/nyc_taxi
{
 "template": "nyc_taxi*",
 "settings": {
 "index.search.turbo.enabled": true,
 "index.sort.field": "dropoff_datetime",
 "number_of_shards": 3,
 "number_of_replicas": 0
 }
}
```

4. Use esrally to test the nyc\_taxi dataset and obtain the result when enhanced aggregation is enabled.

### Test Result

This test focuses on the query result of **dropoff\_datetime** aggregation, that is, the results of tasks **autohisto\_agg** and **date\_histogram\_agg**. The following table

compares the test results between when enhanced aggregation is disabled and when it is enabled.

| Metric                        | Task           | Unit  | Enhanced Aggregation Disabled |              |              | Enhanced Aggregation Enabled |              |              | Enhanced Aggregation Disabled | Enhanced Aggregation Enabled | open / close | Conclusion                               |
|-------------------------------|----------------|-------|-------------------------------|--------------|--------------|------------------------------|--------------|--------------|-------------------------------|------------------------------|--------------|------------------------------------------|
|                               |                |       | Test Round 1                  | Test Round 2 | Test Round 3 | Test Round 1                 | Test Round 2 | Test Round 3 | Mean Value                    | Mean Value                   |              |                                          |
| Min Throughput                | autoh isto_agg | ops/s | 4.42                          | 4.44         | 4.43         | 11.66                        | 11.94        | 11.96        | 4.43                          | 11.85                        | 2.68         | Throughput improves more than 2.5 times. |
| Mean Throughput               | autoh isto_agg | ops/s | 4.50                          | 4.46         | 4.44         | 11.81                        | 11.99        | 12.00        | 4.47                          | 11.93                        | 2.67         |                                          |
| Median Throughput             | autoh isto_agg | ops/s | 4.51                          | 4.46         | 4.44         | 11.83                        | 11.98        | 12.00        | 4.47                          | 11.94                        | 2.67         |                                          |
| Max Throughput                | autoh isto_agg | ops/s | 4.54                          | 4.48         | 4.45         | 11.90                        | 12.07        | 12.02        | 4.49                          | 12.00                        | 2.67         |                                          |
| 100th percentile latency      | autoh isto_agg | ms    | 216.30                        | -            | -            | -                            | 84.56        | 80.38        | 216.30                        | 82.47                        | 0.38         | Latency decreases by more than 60%.      |
| 100th percentile service time | autoh isto_agg | ms    | 216.30                        | -            | -            | -                            | 84.56        | 80.38        | 216.30                        | 82.47                        | 0.38         |                                          |
| error rate                    | autoh isto_agg | %     | 0                             | 0            | 0            | 0                            | 0            | 0            | 0                             | 0                            | 0            | -                                        |

| Metric                   | Task               | Unit  | Enhanced Aggregation Disabled |              |              | Enhanced Aggregation Enabled |              |              | Enhanced Aggregation Disabled | Enhanced Aggregation Enabled | open / close | Conclusion                               |
|--------------------------|--------------------|-------|-------------------------------|--------------|--------------|------------------------------|--------------|--------------|-------------------------------|------------------------------|--------------|------------------------------------------|
|                          |                    |       | Test Round 1                  | Test Round 2 | Test Round 3 | Test Round 1                 | Test Round 2 | Test Round 3 | Mean Value                    | Mean Value                   |              |                                          |
| Min Throughput           | date_histogram_agg | ops/s | 4.72                          | 4.67         | 4.65         | 12.57                        | 12.40        | 12.59        | 4.68                          | 12.52                        | 2.68         | Throughput improves more than 2.5 times. |
| Mean Throughput          | date_histogram_agg | ops/s | 4.73                          | 4.67         | 4.67         | 12.61                        | 12.46        | 12.61        | 4.69                          | 12.56                        | 2.68         |                                          |
| Median Throughput        | date_histogram_agg | ops/s | 4.73                          | 4.67         | 4.67         | 12.62                        | 12.46        | 12.60        | 4.69                          | 12.56                        | 2.68         |                                          |
| Max Throughput           | date_histogram_agg | ops/s | 4.74                          | 4.67         | 4.67         | 12.64                        | 12.49        | 12.63        | 4.69                          | 12.59                        | 2.68         |                                          |
| 50th percentile latency  | date_histogram_agg | ms    | 202.61                        | 218.09       | 213.43       | 77.64                        | 76.02        | 82.63        | 211.38                        | 78.77                        | 0.37         | Latency decreases by more than 60%.      |
| 100th percentile latency | date_histogram_agg | ms    | 207.35                        | 223.88       | 246.63       | 77.99                        | -            | -            | 225.95                        | 77.99                        | 0.35         |                                          |

| Metric                        | Task               | Unit | Enhanced Aggregation Disabled |              |              | Enhanced Aggregation Enabled |              |              | Enhanced Aggregation Disabled | Enhanced Aggregation Enabled | open / close | Conclusion |
|-------------------------------|--------------------|------|-------------------------------|--------------|--------------|------------------------------|--------------|--------------|-------------------------------|------------------------------|--------------|------------|
|                               |                    |      | Test Round 1                  | Test Round 2 | Test Round 3 | Test Round 1                 | Test Round 2 | Test Round 3 | Mean Value                    | Mean Value                   |              |            |
| 50th percentile service time  | date_histogram_agg | ms   | 202.61                        | 218.09       | 213.43       | 77.64                        | 76.02        | 82.63        | 211.38                        | 78.77                        | 0.37         |            |
| 100th percentile service time | date_histogram_agg | ms   | 207.35                        | 223.88       | 246.63       | 77.99                        | -            | -            | 225.95                        | 77.99                        | 0.35         |            |
| error rate                    | date_histogram_agg | %    | 0                             | 0            | 0            | 0                            | 0            | 0            | 0                             | 0                            | 0            | -          |

**Test Conclusion**

Given the same cluster configuration, aggregation performance improves significantly when enhanced aggregation is enabled. Query throughput improves by more than 2.5 times, and latency decreases by more than 60%.

### 2.8.8 Configuring Read/Write Splitting Between Two Elasticsearch Clusters


**Scenario**

Data written to the primary cluster (**Leader**) is automatically synchronized to the secondary cluster (**Follower**). This allows data to be queried from the secondary cluster, improving query performance while alleviating the pressure of the primary cluster. If the primary cluster is unable to provide services, a primary/secondary



switchover can be performed to use the secondary cluster to handle write and query requests, ensuring service continuity.

**Figure 2-19** Two application scenarios for read/write splitting



Scenario 1 (left in the figure): Data is written into the primary cluster but queried from the secondary cluster, alleviating pressure for both clusters.

Scenario 2 (right in the figure): When the primary cluster fails, the secondary cluster takes over to ensure service continuity.

## Constraints

- Only Elasticsearch 7.6.2 and Elasticsearch 7.10.2 clusters support read/write splitting.
- The versions of the primary and secondary clusters must be kept consistent, or errors may occur.

## Prerequisites

Two clusters of the same version have been created. One functions as the primary cluster, and the other the secondary cluster. The secondary cluster must be able to access the REST API (default port: 9200) of the primary cluster.

## Logging In to Kibana

Log in to Kibana and go to the command execution page. Elasticsearch clusters support multiple access methods. This topic uses Kibana as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
4. In the left navigation pane, choose **Dev Tools**.

## Connecting the Primary and Secondary Clusters

1. Run the following command to configure information about the primary cluster in the secondary cluster:

```
PUT /_cluster/settings
{
 "persistent" : {
 "cluster" : {
```

```
"remote.rest" : {
 "leader1" : {
 "seeds" : [
 "http://10.0.0.1:9200",
 "http://10.0.0.2:9200",
 "http://10.0.0.3:9200"
],
 "username": "elastic",
 "password": "*****"
 }
}
```

**Table 2-115** Request body parameters

| Parameter      | Description                                                                                                                        |
|----------------|------------------------------------------------------------------------------------------------------------------------------------|
| <i>leader1</i> | Name of the primary cluster configuration task, which is user-defined and will be used for configuring read/write splitting later. |
| seeds          | Address for accessing the primary cluster. When HTTPS is enabled for the cluster, the URL schema must use HTTPS.                   |
| username       | Username of the primary cluster. This parameter is required only when security mode is enabled for the primary cluster.            |
| password       | Password of the primary cluster. This parameter is required only when security mode is enabled for the primary cluster.            |

Example response:

```
{
 "acknowledged" : true, //Whether the operation is successful
 "persistent" : {
 "cluster" : {
 "remote" : {
 "rest" : {
 "leader1" : {
 "seeds" : [
 "http://10.0.0.1:9200",
 "http://10.0.0.2:9200",
 "http://10.0.0.3:9200"
],
 "username": "elastic",
 "password": "*****"
 }
 }
 }
 }
 },
 "transient" : { }
}
```

2. After the configuration is complete, run the following command in the secondary cluster to check the connection between the secondary and primary clusters:

```
GET _remote/rest/info
```

Example response:

```
{
 "leader1" : {
 "connected" : true //The two clusters are connected.
 }
}
```

## Index Synchronization

There are two ways to synchronize indexes: specified index synchronization and matching index synchronization.

During synchronization, indexes in the secondary cluster become read-only. The synchronization is performed periodically. The default synchronization interval is 30 seconds. For how to change it, see [Changing the Synchronization Interval](#).

### Synchronizing Specified Indexes

- Run the following command in the secondary cluster to synchronize a single index from the primary cluster to the secondary cluster without modifying index settings:

```
PUT start_remote_sync
{
 "remote_cluster": "leader1",
 "remote_index": "data1_leader",
 "local_index": "data1_follower"
}
```

- Run the following command in the secondary cluster to synchronize a single index from the primary cluster to the secondary cluster while modifying some of the index settings—enabling synchronization of index settings:

```
PUT start_remote_sync
{
 "remote_cluster": "leader1",
 "remote_index": "data1_leader",
 "local_index": "data1_follower",
 "settings": {
 "number_of_replicas": 4
 },
 "settings_sync_enable": true,
 "settings_sync_patterns": ["*"],
 "settings_sync_exclude_patterns": ["index.routing.allocation.*"],
 "alias_sync_enable": true,
 "state_sync_enable": true
}
```

#### NOTE

The following index configuration items cannot be modified: **number\_of\_shards**, **version.created**, **uuid**, **creation\_date**, and **soft\_deletes.enabled**.

**Table 2-116** Request body parameters

| Parameter      | Description                                                                                                                                                             |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| remote_cluster | Name of the primary cluster configuration task, which was set in <a href="#">Connecting the Primary and Secondary Clusters</a> . <b>leader1</b> was set in our example. |
| remote_index   | Name of the index to be synchronized in the primary cluster                                                                                                             |

| Parameter                      | Description                                                                                                                                                                                                                                          |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| local_index                    | Index name in the secondary cluster                                                                                                                                                                                                                  |
| settings                       | Index settings to be synchronized                                                                                                                                                                                                                    |
| settings_sync_enable           | Whether to enable synchronization of index settings in the primary cluster. The default value is <b>false</b> .                                                                                                                                      |
| settings_sync_patterns         | Prefix of primary cluster index settings to be synchronized. The default value is *. This parameter takes effect when <b>settings_sync_enable</b> is set to <b>true</b> . The index settings configured in <b>settings</b> will not be synchronized. |
| settings_sync_exclude_patterns | Prefix of primary cluster index settings not to be synchronized. The default value is empty. This parameter is valid only when <b>settings_sync_enable</b> is set to <b>true</b> .                                                                   |
| alias_sync_enable              | Whether to enable index alias synchronization in the primary cluster. The default value is <b>false</b> .                                                                                                                                            |
| state_sync_enable              | Whether to enable index status synchronization in the primary cluster. The default value is <b>false</b> .                                                                                                                                           |

### Matching Index Synchronization

- Run the following command in the secondary cluster to create a pattern-matching index synchronization policy, which synchronizes matched indexes from the primary cluster to the secondary cluster:

```
PUT auto_sync/pattern/${PATTERN}
{
 "remote_cluster": "leader1",
 "remote_index_patterns": "log*",
 "local_index_pattern": "{{remote_index}}-sync",
 "apply_exist_index": true
}
```

- Run the following command in the secondary cluster to create a pattern-matching index synchronization policy, which synchronizes matched indexes from the primary cluster to the secondary cluster, with some of the index settings modified—enabling synchronization of index settings:

```
PUT auto_sync/pattern/${PATTERN}
{
 "remote_cluster": "leader1",
 "remote_index_patterns": "log*",
 "local_index_pattern": "{{remote_index}}-sync",
 "apply_exist_index": true,
 "settings": {
 "number_of_replicas": 4
 },
 "settings_sync_enable": true,
 "settings_sync_patterns": ["*"],
 "settings_sync_exclude_patterns": ["index.routing.allocation:*"],
 "alias_sync_enable": true,
 "state_sync_enable": true
}
```

 NOTE

The following index configuration items cannot be modified: **number\_of\_shards**, **version.created**, **uuid**, **creation\_date**, and **soft\_deletes.enabled**.

**Table 2-117** Request body parameters

| Parameter                      | Description                                                                                                                                                                                                                                          |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PATTERN                        | Name of the pattern for index matching.                                                                                                                                                                                                              |
| remote_cluster                 | Name of the primary cluster configuration task, which was set in <a href="#">Connecting the Primary and Secondary Clusters</a> . In our example, <b>leader1</b> is used.                                                                             |
| remote_index_patterns          | Pattern for matching indexes to be synchronized in the primary cluster. The wildcard (*) is supported.                                                                                                                                               |
| local_index_pattern            | Index pattern in the secondary cluster. The index template can be replaced. For example, if this parameter is set to <b>{{remote_index}}-sync</b> , the index <b>log1</b> changes to <b>log1-sync</b> after synchronization.                         |
| apply_exist_index              | Whether to synchronize existing indexes in the primary cluster. The default value is <b>true</b> .                                                                                                                                                   |
| settings                       | Index settings to be synchronized                                                                                                                                                                                                                    |
| settings_sync_enable           | Whether to enable synchronization of index settings in the primary cluster. The default value is <b>false</b> .                                                                                                                                      |
| settings_sync_patterns         | Prefix of primary cluster index settings to be synchronized. The default value is *. This parameter takes effect when <b>settings_sync_enable</b> is set to <b>true</b> . The index settings configured in <b>settings</b> will not be synchronized. |
| settings_sync_exclude_patterns | Prefix of primary cluster index settings not to be synchronized. The default value is empty. This parameter is valid only when <b>settings_sync_enable</b> is set to <b>true</b> .                                                                   |
| alias_sync_enable              | Whether to enable index alias synchronization in the primary cluster. The default value is <b>false</b> .                                                                                                                                            |
| state_sync_enable              | Whether to enable index status synchronization in the primary cluster. The default value is <b>false</b> .                                                                                                                                           |

## Stopping Index Synchronization

Run the following command in the secondary cluster to stop synchronization tasks for specified indexes. Subsequent changes to the indexes in the primary cluster will not be synchronized to the secondary cluster. The read-only state of the indexes in the secondary cluster will be cancelled, so that new data can be written into them.

```
PUT /log*/stop_remote_sync
```

In this command, **log\*** indicates the index name. You can specify multiple index names (separated by commas) or use a wildcard. In this example, synchronization tasks for all indexes that start with **log** are stopped.

## Querying and Deleting Created Patterns

1. Run the following command in the secondary cluster to query created patterns:

- Query the list of patterns.

```
GET auto_sync/pattern
```

- Query a specified pattern by name.

```
GET auto_sync/pattern/{PATTERN}
```

The following is an example of the response:

```
{
 "patterns" : [
 {
 "name" : "pattern1",
 "pattern" : {
 "remote_cluster" : "leader",
 "remote_index_patterns" : [
 "log*"
],
 "local_index_pattern" : "{{remote_index}}-sync",
 "settings" : { }
 }
 }
]
}
```

2. Run the following command in the secondary cluster to delete a specified pattern by name:

```
DELETE auto_sync/pattern/{PATTERN}
```

## Enabling Forcible Synchronization

By default, the plug-in determines whether to synchronize metadata based on whether the number of documents in the index of the primary cluster changes. If the primary cluster only updates documents and the number of documents remains unchanged, the plug-in does not synchronize the updates to the secondary cluster. The configuration can be modified. After forcible synchronization is enabled, the index metadata of the primary cluster is forcibly synchronized to the secondary cluster in each synchronization cycle.

The following is an example of enabling forcible synchronization:

```
PUT _cluster/settings
{
 "persistent": {
 "remote_sync.force_synchronize": true
 }
}
```

## Changing the Synchronization Interval

The default synchronization interval between the primary and secondary clusters is 30 seconds and can be changed.

The example request below changes the synchronization interval to 2 seconds:

```
PUT {index_name}/_settings
{
```

```
"index.remote_sync.sync_interval": "2s"
}
```

## Changing the Synchronization Speed

You can change the data synchronization speed between the primary and secondary clusters by configuring cluster-level settings.

The following is an example request: set the block size to 2 MB, the number of blocks to 20, and the maximum synchronization traffic to 100 MB per second.

```
PUT _cluster/settings
{
 "persistent": {
 "remote_sync.chunk_size": "2MB",
 "remote_sync.max_concurrent_file_chunks": 20,
 "remote_sync.max_bytes_per_sec": "100MB"
 }
}
```

**Table 2-118** Request body parameters

| Parameter                              | Description                                                                                                |
|----------------------------------------|------------------------------------------------------------------------------------------------------------|
| remote_sync.chunk_size                 | Block size for index synchronization.<br>Default value: 1 MB<br>Format: a string of characters             |
| remote_sync.max_concurrent_file_chunks | Number of blocks for concurrent index synchronization.<br>Default value: <b>10</b><br>Format: a number     |
| remote_sync.max_bytes_per_sec          | Maximum data synchronization traffic per second.<br>Default value: 40 MB<br>Format: a string of characters |

## Querying Index Synchronization Status

- **Obtaining the auto synchronization status of a specified index**

An example request is as follows:

```
GET {index_name}/sync_stats
```

An example response is as follows:

```
{
 "indices" : {
 "data1_follower" : {
 "shards" : {
 "0" : [
 {
 "primary" : false,
 "total_synced_times" : 27,
 "total_empty_times" : 25,
 "total_synced_files" : 4,
 "total_synced_bytes" : 3580,
 "total_paused_nanos" : 0,
 "total_paused_times" : 0,
 "current" : {
```

```
 "files_count" : 0,
 "finished_files_count" : 0,
 "bytes" : 0,
 "finished_bytes" : 0
 }
},
{
 "primary" : true,
 "total_synced_times" : 28,
 "total_empty_times" : 26,
 "total_synced_files" : 20,
 "total_synced_bytes" : 17547,
 "total_paused_nanos" : 0,
 "total_paused_times" : 0,
 "current" : {
 "files_count" : 0,
 "finished_files_count" : 0,
 "bytes" : 0,
 "finished_bytes" : 0
 }
}
]
}
}
```

## Switching the Roles of the Primary and Secondary Clusters

When the primary cluster becomes faulty, perform a primary/secondary switchover to have the secondary cluster take over services. The steps are as follows:

1. Determine the index synchronization method between the primary and secondary clusters. Check whether pattern-matching index synchronization policies have been configured in the secondary cluster. For the command to use, see [Querying and Deleting Created Patterns](#).
  - If there are no such policies, synchronization is performed for specified indexes between the primary and secondary clusters. In this case, go to [3](#).
  - If there are such policies, index synchronization between the primary and secondary clusters is based on index patterns. In this case, go to [2](#).
2. Delete pattern-matching index synchronization policies in the secondary cluster. For the command to use, see [Querying and Deleting Created Patterns](#).
3. Perform [Stopping Index Synchronization](#) in the secondary cluster. Then redirect read and write traffic to it. If the primary and secondary clusters synchronize indexes based on index patterns, use a wildcard to match indexes when running the command that stops index synchronization.
4. After the primary cluster recovers, configure information about the secondary cluster in the primary cluster, and connect the primary and secondary clusters again. For details, see [Connecting the Primary and Secondary Clusters](#).
5. Under the primary cluster, perform [Index Synchronization](#) to synchronize data from the secondary cluster to the primary cluster, and then perform a primary/secondary switchover to switch back.



## 2.8.9 Switching Between Hot and Cold Storage for an Elasticsearch Cluster


In an Elasticsearch cluster, switching between hot and cold data storage means to allocate data to nodes of different performance standards based on data temperature (that is, how often data is accessed). The goal is to achieve optimal storage costs and query performance.

- Hot data: frequently accessed data, low latency required, stored using high-performance hardware (such as SSDs) to ensure fast write and query performance.
- Cold data: rarely accessed data, cost-optimized storage (such as HDDs).

If your cluster stores data used for different purposes, such as real-time analytics, log analytics, and monitoring data archives, you can switch between hot and cold storage for specified indexes to balance performance and costs.

### How the Feature Works

**Figure 2-20** How cold/hot storage switchover works



The key is to allocate index data storage by node labels and index allocation policies.

- Node labels:
  - Data node (hot): Stores real-time data by default and supports high-concurrency read/write requests.
  - Cold data node (cold): Stores historical data. Cold data nodes use less expensive hardware and deliver lower query performance than data nodes.
- Data allocation:

You can configure an index template or directly configure specific indexes to allocate data to data nodes or cold data nodes.

Procedure: Enable cold data nodes when creating a cluster, and configure an index template or specific index settings to allocate index data storage. The cluster automatically allocates data based on your settings.

Constraints

- You cannot add cold data nodes to existing clusters that did not have such nodes enabled upon cluster creation.
- In comparison with data nodes, cold data nodes deliver lower query performance. Determine which type of node to use based on service needs.

Switching Over Between Hot and Cold Storage

- Step 1** Check whether cold data nodes are enabled in the target cluster.
1. Log in to the [CSS management console](#).
  2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
  3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  4. On the **Overview** tab, check whether the **Node Information** area contains cold data node information.

Figure 2-21 Cold data node information

| Node                        |              |
|-----------------------------|--------------|
| Node Specifications         | ess.s   8 GB |
| Node Storage Type           | 40 G         |
| Nodes                       | 1            |
| Cold Data Node Spec         | ess.s   8 GB |
| Cold Data Node Storage Type | 40 G         |
| Cold Data Node              | 3            |

- If there is information about cold data nodes, the cluster has cold data nodes. Go to the next step.
- Otherwise, the cluster does not have cold data nodes, in which case, you cannot switch between cold and hot data storage.

**Step 2** Log in to the Kibana console.

1. On the cluster information page, click **Kibana** in the upper-right corner to log in to Kibana.
2. In the left navigation pane, choose **Dev Tools**.

**Step 3** On the Kibana page, configure index allocation policies.

You can configure an index template or directly configure specific indexes to allocate data to data nodes or cold data nodes.

- **Configuring an index template**

Configure an index template to allocate indexes to cold or hot data nodes. For example, run the following command to store indexes whose name starts with **myindex** to cold data nodes. (The command varies slightly depending on the Elasticsearch version.)

- For Elasticsearch 6.x or later:

```
PUT _template/test
{
 "order": 1,
 "index_patterns": "myindex*",
 "settings": {
 "refresh_interval": "30s",
 "number_of_shards": "3",
 "number_of_replicas": "0",
 "routing.allocation.require.box_type": "cold"
 }
}
```

- For Elasticsearch earlier than 6.x:

```
PUT _template/test
{
 "order": 1,
 "template": "myindex*",
 "settings": {
 "index": {
 "refresh_interval": "30s",
 "number_of_shards": "3",
 "number_of_replicas": "0",
 "routing.allocation.require.box_type": "cold"
 }
 }
}
```

Parameter description:

- **template** or **index\_patterns**: rule for matching index names (for example, **myindex\***).
- **box\_type**: node type for matching indexes. **cold** indicates cold data nodes, and **hot** indicates regular data nodes.

- **Configuring specific indexes**

To change the node type for an existing index, run the following command:

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

Parameter description:

- **myindex**: index name.
- **box\_type**: node type for matching indexes. **cold** indicates cold data nodes, and **hot** indicates regular data nodes.

**Step 4** Verify the switchover between hot and cold storage.

Run the following command to check the distribution of index shards:

```
GET _cat/shards/myindex?v
```

When the data volume is large, the switchover may take a long time. There may be an intermediate state where the data of an index resides on both cold data nodes and data nodes.

As shown in the following figure, all shards of the **myindex** index are stored on the cold data node **css-e668-ess-cold-esn-1-1**.

| index   | shard | pri | rep | state   | docs     | store  | ip             | node                      |
|---------|-------|-----|-----|---------|----------|--------|----------------|---------------------------|
| myindex | 1     | p   |     | STARTED | 14085446 | 17.8gb | 192.168.91.188 | css-e668-ess-cold-esn-1-1 |
| myindex | 2     | p   |     | STARTED | 14094005 | 17.9gb | 192.168.91.188 | css-e668-ess-cold-esn-1-1 |
| myindex | 0     | p   |     | STARTED | 14094742 | 17.8gb | 192.168.91.188 | css-e668-ess-cold-esn-1-1 |

#### Step 5 Roll back the cold-hot switchover configuration.

To cancel the cold-hot switchover configuration, run the following command:

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

After the rollback, index data will be evenly and randomly allocated to both cold data nodes and data nodes.

----End

## Related Operations

- You can scale cold data nodes by adding or reducing nodes or their storage capacity. For details, see [Scaling an Elasticsearch Cluster](#).
- If your cluster does not have cold data nodes but you wish to cut storage costs, you can try using decoupled storage and compute. For details, see [Configuring Decoupled Storage and Compute for an Elasticsearch Cluster](#).

## 2.8.10 Configuring an Index Recycle Bin for an Elasticsearch Cluster

With Elasticsearch clusters, indexes are deleted without first being put into a recycle bin. To allow users to restore data after misdeletion, CSS provides an index recycle bin. Similar to other recycle bins, the index recycle bin temporarily stores deleted indexes so that users can restore them before they are finally removed from the recycle bin. This helps to improve cluster data reliability.

## Constraints

- Only Elasticsearch 7.10.2 clusters support the index recycle bin.
- With the index recycle bin enabled, after you use the DELETE API to move an index to the recycle bin, the cluster status may stay RED for a short period of time. This is because putting the index into the recycle bin causes the index to be closed, and the Elasticsearch cluster status becomes RED when there are closed indexes.
- When a restoration API is used to restore an index from the recycle bin, the cluster status may also stay RED for a short period of time. This is because restoring an index from the recycle bin causes that index to be reopened, and

when it happens, Elasticsearch needs to re-initialize shards. This causes the cluster status to become RED momentarily.

- Indexes stored in the recycle bin are still part of the cluster metadata. Before they are removed from the recycle bin, same-name indexes cannot be created for the cluster.

## Enabling the Index Recycle Bin

Run the following command to enable the index recycle bin:

```
PUT _cluster/settings
{
 "persistent": {
 "index.trash.enabled": true
 }
}
```

**Table 2-119** Configuration items

| Configuration Item             | Description                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>index.trash.enabled</b>     | Whether to enable the index recycle bin. The default value is <b>false</b> . This setting is compatible with the open-source Elasticsearch. Setting this parameter to <b>true</b> enables the index recycle bin. After an index is deleted, the index is stored in the recycle bin. You need to delete the index again to permanently delete the index data. |
| <b>indices.trash.keep.time</b> | Duration for retaining indexes in the recycle bin. The default value is <b>1d</b> , meaning indexes will be retained in the recycle bin for one day before they are permanently deleted. The minimum value is <b>1d</b> .                                                                                                                                    |

### NOTE

The following command deletes an index:

```
DELETE {INDEX_NAME}
```

**INDEX\_NAME** indicates the name of the index to be deleted. Wildcards can be used to specify indexes.

## Viewing Indexes in the Recycle Bin

Run the following command to view indexes in the index recycle bin:

```
GET _cat/trash?v=true&s=index
```

**Table 2-120** Parameters

| Parameter | Description                                                                                                                                                                                                                                 |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| v         | Whether to display the table header when the return format is a table. <ul style="list-style-type: none"><li>• true: display the table header.</li><li>• false: not display the table header.</li></ul> The default value is <b>false</b> . |
| s         | Index sorting fields. Options include <b>index</b> , <b>uuid</b> , <b>pri</b> , <b>rep</b> , <b>trash.ts</b> , and <b>delete.time</b> .                                                                                                     |
| format    | Return format of the command. The default format is table. Other options include json, yaml, coBr, and smile. coBr and smile are binary formats.                                                                                            |
| h         | Set the column names to be displayed. By default, all column names are displayed. To specify the column names you want to display, separate them using commas (,), for example, <b>h=index,uuid,delete.time</b> .                           |

The following is an example of the output.

```
index uuid pri rep trash.ts delete.time
index1 CMD3FCLzTOyTg4RUekWNNA 1 1 1714465116615 23.6h
index1 6ATijuu6SfqamVI-WMyOKg 1 1 1714466233898 23.9h
```

**Table 2-121** Parameters in the output

| Column      | Description                                                                                                                 |
|-------------|-----------------------------------------------------------------------------------------------------------------------------|
| index       | Index name                                                                                                                  |
| uuid        | Index UUID                                                                                                                  |
| pri         | Number of shards of an index                                                                                                |
| rep         | Number of replicas of an index                                                                                              |
| trash.ts    | Time when an index was moved to the recycle bin                                                                             |
| delete.time | Remaining retention duration of an index in the recycle bin. When the value changes to 0, the index is permanently deleted. |

## Restoring an Index from the Recycle Bin

Run the following command to restore an index from the recycle bin:

```
POST /trash/recover/{INDEX_NAME}
```

**INDEX\_NAME** indicates the name of the index to be restored. Wildcards can be used to specify indexes.

## Emptying the Recycle Bin

Run the following command to clear the recycle bin:

```
POST trash/empty
```


Emptying the recycle bin will permanently delete all data in it. Please exercise caution.

## 2.9 Configuring Networking for an Elasticsearch Cluster

### 2.9.1 Configuring Routes for an Elasticsearch Cluster

By default, a CSS Elasticsearch cluster cannot access a target service—such as a client program or a Lightweight Directory Access Protocol (LDAP) service—that is located in a different VPC. To enable it, you need to configure a route for the Elasticsearch cluster.

**Figure 2-22** Configuring a route for an Elasticsearch cluster



1. Establish a VPC peering connection to connect two VPCs.
2. On CSS, configure a route for the Elasticsearch cluster that points to the client.
3. Verify network connectivity after starting services.

Configuring routes to enable cluster communication offers the following advantages:

- Improved cross-VPC communication performance: A route table specifies the traffic forwarding paths, so that traffic does not have to pass through the public network.
- Secure communication: Private networks can communicate through VPC peering connections, which is more secure than transmitting data over the public network.

- Flexible access control policies: You can customize access permissions for specific subnets.
- Support for a hybrid cloud architecture: Services can communicate across VPCs in different environments.

## Constraints

- By default, routes can be configured only for Elasticsearch clusters whose image version is no earlier than x.x.x\_25.3.0\_x.x.x.
- Adding cluster routes will alter network connections, possibly causing disruptions in the network. Exercise caution.

## Prerequisites

- The Elasticsearch cluster is Available, and port 9200 is allowed in the inbound direction by the cluster's security group.
- The client is running properly, and the server's security group allows port 9200 in the outbound direction.

## Establishing a VPC Peering Connection

Before configuring routes for a CSS cluster, establish a connection between the CSS cluster's VPC and that of the client, that is, a VPC peering connection.

- If the CSS cluster and the client belong to the same account, see [Creating a VPC Peering Connection to Connect Two VPCs in the Same Account](#).
- If the CSS cluster and the client belong to different accounts, see [Creating a VPC Peering Connection to Connect Two VPCs in Different Accounts](#).

Set the peer VPC to the VPC where the CSS cluster is located and the local VPC to the VPC where the client is located.

When adding routes for a VPC peering connection, configure the destination addresses based on service requirements.

If the CSS cluster and the client share the same security group, there is no need to configure security group rules for the instances in the VPC at either end of the VPC peering connection.

## Configuring a Route for the Elasticsearch Cluster

Configure a route for the Elasticsearch cluster that points to the client, enabling the cluster to access the client.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Overview** tab. In the **Configuration** area, click **Add Route** next to **Cluster Route**.



**Figure 2-23** Adding a route

5. In the displayed dialog box, configure the route information.

**Table 2-122** Adding a route

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP Address  | Enter the IP address (or a CIDR block) of the client or server that the cluster needs to access, for example, <b>10.10.1.0</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Subnet Mask | <p>Enter the subnet mask used with the IP address. The subnet mask must align with the network part of the IP address. That is:</p> <ul style="list-style-type: none"><li>• The network bits (the 1s) of the subnet mask must encompass those of the IP address.</li><li>• The host bits (the 0s) of the subnet mask must accommodate the IP address's host bits.</li></ul> <p>The subnet mask should not cover too wide a range, avoiding unnecessary IP address exposure or waste of the address space. 255.255.255.0 (that is, /24) is recommended, which applies to most internal networking scenarios.</p> |

6. Click **OK** to add the route.  
After the route is added, click **View Route** to check the route information. You can click **Add Route** to add another route; or click **Delete Route** to delete one—In the **Delete Route** dialog box, select a route, enter **DELETE** manually, and then click **OK**.

## Testing Connectivity

Connectivity between the Elasticsearch cluster and client cannot be verified right away. After the route is configured, start services. If the search service can be accessed, the route is configured successfully.

## 2.9.2 Configuring Public Network Access for an Elasticsearch Cluster

When public network access is enabled for an Elasticsearch cluster, the cluster is automatically assigned a public IP address with dedicated dynamic BGP bandwidth, making it accessible from the Internet via HTTPS. You can configure public network access control by IP addresses or IP address ranges.

To enable public network access for Elasticsearch clusters, a shared load balancer is typically used for load balancing. If your workloads require quicker access, you

are advised to use a dedicated load balancer to connect to your clusters. For details about its configuration, see [Configuring a Dedicated Load Balancer for an Elasticsearch Cluster](#).

## Constraints

- Enabling public network access for a CSS cluster may incur some fees, as it will need to use EIP and bandwidth resources.
- To enable public network access for an Elasticsearch cluster, two conditions must be met: the cluster version is 6.5.4 or later; the security mode and HTTPS access are both enabled.
- Public network access and the VPC Endpoint service share a load balancer. If you configure a whitelist for public network access, and because this whitelist is deployed to the shared load balancer, it will control not only access from the public network, but also access using private IP addresses through VPCEP. In this case, you need to add IP address **198.19.128.0/17** to the public network access whitelist to allow traffic through VPCEP.

## Enabling Public Network Access

To enable public network access for an existing cluster, perform the following steps:

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, check whether **Security Mode** and **HTTPS Access** are enabled in the **Configuration** area.
  - If they are enabled, go to the next step to enable public network access.
  - If either one is disabled, public network access cannot be enabled for the cluster.

**Figure 2-24** Checking whether public network access can be enabled for the cluster

|                       |                                                |
|-----------------------|------------------------------------------------|
| Security Mode         | Enabled                                        |
| Reset Password        | <a href="#">Reset</a>                          |
| Cluster Route         | <a href="#">Add Route</a>                      |
| Enterprise Project    | <a href="#">default</a>                        |
| Public Network Access | -- <a href="#">Enable</a>                      |
| HTTPS Access          | Enabled   <a href="#">Download Certificate</a> |


5. Click **Enable** next to **Public Network Access**. In the displayed dialog box, configure the necessary settings.

**Table 2-123** Enabling public network access

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bandwidth           | Cluster bandwidth for public network access.<br>Value range: 1 Mbit/s to 200 Mbit/s                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Configure Whitelist | <p>Control public network access to the cluster using a whitelist.</p> <ul style="list-style-type: none"><li>• If a whitelist is configured, only IP addresses that are on this whitelist can access the cluster over the public network.<br/>Click <b>+Add</b>. In the displayed text box, enter IP addresses or CIDR blocks that are allowed to access the cluster from the public network. Separate them using commas (,). Each value must be unique. An example of valid values: <b>192.168.1.1,10.0.0.0/24</b>. Examples of invalid values: <b>0.0.0.0</b>, <b>xx.xx.xx.x0</b>, <b>172.16.0.0-172.16.255.255</b>, non-standard formats (e.g., <b>192.168.1</b>), and duplicate values.</li><li>• If no whitelist is configured, all public IP addresses can access the cluster. However, this can be a security risk and should be avoided.</li></ul> |

6. Click **OK** to enable public network access.


After public network access is enabled, the public IP address, public network access control, and bandwidth information is displayed.

**Figure 2-25** Enabling public network access

## Managing Public Network Access

When public network access is enabled, you can check the public IP address, and modify the bandwidth and access control settings.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, manage public network access settings in the **Configuration** area.

**Figure 2-26** Managing public network access

- **Checking the public IP address**  
Record the IP address and port number displayed next to **Public Network Access**.
- **Modifying public network access control settings**  
Click **Modify** next to **Public Network Access Control**. In the displayed dialog box, add or remove IP addresses or CIDR blocks to or from the whitelist. Click **OK** to save the change.
- **Modifying public network bandwidth**  
Click **Modify** next to **Bandwidth**. In the displayed dialog box, change the bandwidth. Click **OK** to save the change.

## Disabling Public Network Access

If public network access is no longer required for a cluster, disable it to release resources.

### WARNING

After the public IP address is disassociated, the cluster can no longer be accessed from the Internet through this IP address. If you disable public network access for a cluster and then re-enable it, the public IP address for accessing the cluster may change. Exercise caution.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, find **Public Network Access** in the **Configuration** area, and click **Disable** next to it. In the displayed dialog box, enter **CONFIRM** and click **OK**.

After public network access is disabled, the public IP address and the **Public Network Access Control** and **Bandwidth** parameters disappear.

## Accessing a Cluster Through the Public IP Address

After public network access is enabled, the cluster is assigned a public IP address. You can use this IP address plus a port number to access this cluster.

For example, if the public IP address is **10.62.xxx.xxx** and the port number is **9200**, run the following cURL command to view indexes in the cluster.

```
curl -u username:password -k 'https://10.62.xxx.xxx:9200/_cat/indices'
```

where, **username** and **password** indicate the username and password of the HTTPS-enabled security-mode cluster.

## 2.9.3 Configuring VPC Endpoints for an Elasticsearch Cluster

VPC Endpoint enables you to access resources across VPCs using a dedicated gateway, without exposing the network information of servers. A VPC endpoint can be accessed via an IPv4 address or private domain name.

- An IPv4 address is automatically allocated when VPC Endpoint is enabled.
- A private domain name is allocated only when you enable private domain names.

VPC Endpoint uses a shared load balancer for internal network access. If your workloads require faster access, we recommend that you use a dedicated load balancer to handle access to your cluster. For details, see [Configuring a Dedicated Load Balancer for an Elasticsearch Cluster](#).

### Impact on Billing

VPC endpoints, if created for the cluster, will incur extra fees, depending on the resource usage. For details, see [VPC Endpoint Billing](#).

### Constraints


- If a shared VPC and a subnet within this shared VPC were selected earlier for the cluster on the **Network** page, **VPC Endpoint** cannot be enabled for the cluster.
- VPC endpoint creation requires specific permissions. For details, see [VPCEP Permissions](#).
- Public network access and the VPC Endpoint service share a load balancer. If you configure a whitelist for public network access, and because this whitelist is deployed to the shared load balancer, it will control not only access from the public network, but also access using private IP addresses through VPCEP. In this case, you need to add IP address **198.19.128.0/17** to the public network access whitelist to allow traffic through VPCEP.
- After VPCEP is enabled, access to CSS through a VPCEP IP address or private domain name from within the internal network is not controlled by any cluster security group rules. Rather, you need to configure a VPCEP whitelist to implement access control. For details, see [Configuring Access Control for an Interface VPC Endpoint](#).

### Enabling VPC Endpoint

If VPC Endpoint was not enabled when the cluster was created, you can enable it as follows:

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.

3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, check the cluster's subnet information in the **Configuration** area. When you add a VPC endpoint, an IP address that belongs to the current subnet of the cluster is automatically assigned to it.

**Figure 2-27** Checking the current subnet

If you want to use another subnet, switch the subnet first, and then enable VPC Endpoint. For details, see [Can I Expand the Subnet for an Elasticsearch or OpenSearch Cluster?](#)


5. Choose **Cluster Access > VPC Endpoint**.
6. Toggle on **VPC Endpoint**. In the displayed dialog box, select relevant options.

**Table 2-124** Options for enabling VPC Endpoint

| Option                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create Private Domain Name | <p>Whether to create a private domain name for the VPC endpoint.</p> <ul style="list-style-type: none"><li>• Enable: The system automatically assigns a private domain name to the VPC endpoint. After cluster creation, you can check this private domain name on the <b>VPC Endpoint</b> tab of the cluster details page.</li><li>• Disable: No private domain name will be configured for the VPC endpoint. The cluster can only be accessed through an IP address assigned to the VPC endpoint.</li></ul> |

7. Click **OK** to enable VPC Endpoint. After the VPC endpoint is enabled, its information is displayed below. When its status changes to **Accepted**, the current cluster can be accessed through the VPC endpoint.

Figure 2-28 Checking the VPC endpoint status



Managing VPC Endpoints

After VPC Endpoint is enabled for a cluster, you can set access control, check VPC endpoint information, and deny access from specific VPC endpoints.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Cluster Access > VPC Endpoint**.
5. **Configure access control for a VPC endpoint.**
  - a. Click **Modify** on the right of **VPC Endpoint Whitelist**. In the displayed dialog box, add accounts that are allowed to access the cluster through the VPC endpoint. If no account is added or the account ID is set to \*, all users are allowed to access the cluster through the VPC endpoint.
    - Click **Add** to add accounts in **Account ID**. To obtain your authorized account ID, point to your username in the upper right corner, and choose **My Credentials**. Copy the value of **Account ID**.
    - Click **Delete** in the **Operation** column to delete an authorized account.
  - b. Click **OK**.
6. **Check VPC endpoint information.**

The VPC endpoint list shows the VPC endpoints created for the current cluster. You can obtain their **Status**, **Service Address**, and **Private Domain Name**.

Figure 2-29 Managing VPC endpoints

| VPC Endpoint ID | Status   | Service Address | Private Domain Name | Max. Connections | Owner      | Created          | Operation     |
|-----------------|----------|-----------------|---------------------|------------------|------------|------------------|---------------|
| 6c6e-03be-4...  | Accepted | 192...          | vpcep-03be-439...   | 3000             | c4b2-08... | Aug 30, 2024 ... | Accept Reject |

7. **Modify VPC endpoint status** to make the cluster accessible or inaccessible through specific endpoints.

- In the VPC endpoint list, select the target endpoint and click **Reject** in the **Operation** column. If the endpoint status changes to **Rejected**, it means the cluster is no longer accessible through this endpoint.
- Select an endpoint whose status is **Rejected**. Click **Accept** in the **Operation** column. If the endpoint status changes to **Accepted**, it means the cluster is accessible again through this endpoint.

## Disabling VPC Endpoint

If the cluster no longer requires cross-VPC access via VPC endpoints, disable VPC Endpoint to release resources.

### WARNING

After VPC Endpoint is disabled, the cluster is no longer accessible through a VPCEP IP address or private domain name. If you disable VPC Endpoint and then re-enable it, the VPCEP IP address or private domain name for accessing the cluster may change. When it happens, you may need to update the client connection. If you just need to temporarily disable VPC Endpoint (rather than permanently releasing its resources), do so by rejecting specific VPC endpoints. For details, see [7](#).

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Cluster Access > VPC Endpoint**.
5. Toggle off **VPC Endpoint**. In the displayed dialog box, enter **CONFIRM** and click **OK**.

## Accessing a Cluster Through a VPC Endpoint

1. Obtain the private domain name or IP address of a VPC endpoint.  
In the VPC endpoint list, check **Service Address** or **Private Domain Name**.

**Figure 2-30** Checking the service address or private domain name of a VPC endpoint

| VPC Endpoint ID  | Status   | Service Address | Private Domain Name  | Max. Connections | Owner        | Created          | Operation     |
|------------------|----------|-----------------|----------------------|------------------|--------------|------------------|---------------|
| 6c5e...03be-4... | Accepted | 192...          | vpcep-...03be-439... | 3000             | c4b2...08... | Aug 30, 2024 ... | Accept Reject |

2. On the client (e.g., an ECS), run the curl command to access the cluster.  
For example, run the following command to check the cluster's index information:
  - For a cluster with the security mode disabled:  

```
curl "http://<host>:9200/_cat/indices"
```
  - For a security-mode cluster that uses HTTP:  

```
curl -u <user>:<password> "http://<host>:9200/_cat/indices"
```
  - For a security-mode cluster that uses HTTPS:  

```
curl -u <user>:<password> -k "https://<host>:9200/_cat/indices"
```



where, **user** and **password** are the username and password used for access the cluster, and **host** indicates the VPC endpoint's private domain name or IP address.

## 2.9.4 Configuring a Dedicated Load Balancer for an Elasticsearch Cluster

CSS integrates shared load balancers, through which you can enable access to a cluster from the public network as well as through the VPC Endpoint service. Dedicated load balancers provide higher performance and more diverse features than shared load balancers. This topic describes how to configure a dedicated load balancer for a cluster.

### Scenarios

Advantages of connecting to a cluster through a dedicated load balancer:

- A non-security cluster can also use the capabilities of the Elastic Load Balance (ELB) service.
- You can use custom certificates for HTTPS two-way authentication.
- Seven-layer traffic monitoring and alarm configuration are supported, allowing you to keep close track of the cluster status.

There are eight different ELB service forms for clusters in different security modes to connect to a dedicated load balancer. [Table 2-125](#) describes the ELB capabilities for different cluster configurations. [Table 2-126](#) describes the configurations for different ELB service forms.

**Table 2-125** 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                            | Yes                | Yes                    | Yes                        |
|                      | Two-way authentication                            |                    |                        |                            |
| Security mode + HTTP | Password authentication                           | Yes                | Yes                    | No                         |
|                      | One-way authentication + Password authentication  | Yes                | Yes                    | Yes                        |
|                      | Two-way authentication + Password authentication  |                    |                        |                            |

| Security Mode         | Service Form Provided by ELB for External Systems                                                        | ELB Load Balancing | ELB Traffic Monitoring | ELB Two-way Authentication |
|-----------------------|----------------------------------------------------------------------------------------------------------|--------------------|------------------------|----------------------------|
| Security mode + HTTPS | One-way authentication + Password authentication<br><br>Two-way authentication + Password authentication | Yes                | Yes                    | Yes                        |

**Table 2-126** Configurations for different ELB service forms depending on the cluster

| Security Mode        | Service Form Provided by ELB for External Systems | ELB Listener      | ELB Listener  | ELB Listener           | Backend Server Group | Backend Server Group | Backend Server Group          |
|----------------------|---------------------------------------------------|-------------------|---------------|------------------------|----------------------|----------------------|-------------------------------|
|                      |                                                   | Frontend Protocol | Frontend Port | SSL Authentication     | 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                 |                               |

| Security Mode         | Service Form Provided by ELB for External Systems | ELB Listener      | ELB Listener  | ELB Listener           | Backend Server Group | Backend Server Group | Backend Server Group |
|-----------------------|---------------------------------------------------|-------------------|---------------|------------------------|----------------------|----------------------|----------------------|
|                       |                                                   | Frontend Protocol | Frontend Port | SSL Authentication     | Backend Protocol     | Health Check Port    | Health Check Path    |
|                       | 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                 |                      |

## Constraints

- You are not advised to connect a load balancer that has been associated with a public IP address to a non-security mode cluster. Allowing public network access through such a load balancer may cause security risks because a non-security mode cluster can be accessed via HTTP without security authentication.
- HTTPS-enabled security-mode clusters do not support HTTP-based frontend authentication. If the frontend uses HTTP, disable security mode for the clusters first. For details, see [Changing the Security Mode of an Elasticsearch Cluster](#). Before changing the security mode, disable load balancing first. After the security mode is changed, enable load balancing again.

## Prerequisites

- A dedicated load balancer has been created. For details, see [Creating a Dedicated Load Balancer](#). This load balancer must meet the following requirements:
  - Its VPC is the same as that of the CSS cluster. The network between the two are connected.
  - IP as a Backend** is enabled. This is necessary to connect a dedicated load balancer to a CSS cluster.

- Determine whether to configure an EIP based on service needs. A public IP address is displayed for the load balancer connecting the CSS cluster only if an EIP is configured. This will enable public network access to the cluster through this load balancer.
- If the ELB listener uses HTTPS, upload a server certificate or CA certificate to the ELB console. For details, see [Adding a Certificate](#).
  - If one-way authentication is used, upload a server certificate.
  - If two-way authentication is used, upload a server certificate and a CA certificate.

## Connecting a Cluster to a Load Balancer

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Cluster Access** tab, and then click the **Load Balancing** tab. On the **Elasticsearch** tab, toggle on **Load Balancing**. In the displayed dialog box, set the parameters.

**Table 2-127** Configuring load balancing

| Parameter     | Description                                                                                                                                                       |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Load Balancer | Select the dedicated load balancer you have created earlier.<br><br>To create a dedicated load balancer, see <a href="#">Creating a Dedicated Load Balancer</a> . |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>To configure a load balancer, you must have the permission to access ELB resources. By configuring an IAM agency, you can authorize CSS to access its ELB resources through an associated account.</p> <ul style="list-style-type: none"><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-elb-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>ELB Administrator</b> role or the <b>ELB FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<pre>"elb:loadbalancers:list", "elb:loadbalancers:get", "elb:certificates:list", "elb:healthmonitors:*", "elb:members:*", "elb:pools:*", "elb:listeners:*"</pre></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<pre>"iam:agencies:listAgencies", "iam:roles:listRoles", "iam:agencies:getAgency", "iam:agencies:createAgency", "iam:permissions:listRolesForAgency", "iam:permissions:grantRoleToAgency", "iam:permissions:listRolesForAgencyOnProject", "iam:permissions:revokeRoleFromAgency", "iam:roles:createRole"</pre></li><li>To use an IAM agency, the following minimum permissions are required:<pre>"iam:agencies:listAgencies", "iam:agencies:getAgency", "iam:permissions:listRolesForAgencyOnProject", "iam:permissions:listRolesForAgency"</pre></li></ul> |

**Figure 2-31** Enabling load balancing**Basic Configuration**

Load Balancer  [Create Load Balancer](#) [?](#)

Agency  [Create Agency](#) [?](#)

[Cancel](#) [OK](#)

- Click **OK** to enable load balancing.  
Load balancer information is displayed.

6. In the **Listener Configuration** area, click  to configure listener information.

**Table 2-128** Listener configuration

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Frontend Protocol  | Protocol used by the client and listener to distribute traffic.<br>Select <b>HTTP</b> or <b>HTTPS</b> .<br>Select this protocol based on your connectivity needs.                                                                                                                                                                                                                                                 |
| Frontend Port      | Port used by the client and listener to distribute traffic.<br>Set this parameter based on site requirements.                                                                                                                                                                                                                                                                                                     |
| SSL Authentication | Authentication mode for the client to access the server. Set this parameter only when <b>Frontend Protocol</b> is set to <b>HTTPS</b> .<br>Both one-way and two-way authentication are supported.<br>Select an authentication mode that suits your needs.                                                                                                                                                         |
| Server Certificate | The server certificate is used for SSL handshake. The certificate content and private key must be provided. It is required only when <b>Frontend Protocol</b> is set to <b>HTTPS</b> .<br>Select the server certificate created on ELB.                                                                                                                                                                           |
| CA Certificate     | Also called client CA public key certificate. It is used to verify the issuer of a client certificate. It is required only when <b>SSL Authentication</b> is set to <b>Two-way authentication</b> .<br>Select the CA certificate created on ELB.<br>When HTTPS two-way authentication is enabled, an HTTPS connection can be established only when the client can provide the certificate issued by a trusted CA. |

**Figure 2-32** Listener configuration  
monitor configuration

Frontend Protocol ☐ HTTP ☒ HTTPS

Frontend Port

SSL Authentication ☐ One-way authentication ☒ Two-way authentication

Server Certificate   [Create CA Certificate](#) 

CA Certificate   [Create CA Certificate](#) 

7. (Optional) In the **Listener** area, click **Configure** next to **Access Control** to go to the listener list of the load balancer. Click **Configure** in the Access Control column of a listener to configure access control for that listener. For more information, see [What Is Access Control?](#).

Without access control policies, all IP addresses are allowed to access the CSS cluster through this load balancer, which may create security risks.

8. In the **Health Check** area, you can view the health check result for each node IP address.

**Table 2-129** Health check result description

| Health Check Result | Description                          |
|---------------------|--------------------------------------|
| Normal              | The node IP address is connected.    |
| Abnormal            | The node IP address is disconnected. |

9. If the cluster no longer needs a dedicated load balancer, disassociate it to release resources.

Choose **Load Balancing** > **Elasticsearch**, toggle off **Load Balancing**. In the displayed dialog box, click **OK**.

---

 **CAUTION**

After the load balancer is disassociated, any listener or backend server group configurations will be permanently deleted.

---

## Accessing a Cluster Through a Load Balancer by Executing cURL Commands

1. In the navigation pane on the left, choose **Clusters**.
2. Log in to the [CSS management console](#).
3. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
4. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
5. Click the **Cluster Access** tab, and then click the **Load Balancing** tab. On the **Elasticsearch** tab, record the private or public IP address or IPv6 address of the load balancer, as well as the frontend protocol/port of the listener.

### CAUTION

You are not advised to connect a load balancer that has been associated with a public IP address to a non-security mode cluster. Allowing public network access through such a load balancer may cause security risks because a non-security mode cluster can be accessed via HTTP without security authentication.

6. Run the following cURL commands on an ECS to check whether the dedicated load balancer can connect to the cluster.

**Table 2-130** Commands for accessing different types of 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:port</code>                                                           |
|                      | One-way authentication                            | <code>curl --cacert ./ca.crt https://IP:port</code>                                        |
|                      | Two-way authentication                            | <code>curl --cacert ./ca.crt --cert ./client.crt --key ./client.key https://IP:port</code> |
| Security mode + HTTP | Password authentication                           | <code>curl http://IP:port -u user:pwd</code>                                               |
|                      | One-way authentication + Password authentication  | <code>curl --cacert ./ca.crt https://IP:port -u user:pwd</code>                            |



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

**Table 2-131** Variables

| Variable | Description                                                                                  |
|----------|----------------------------------------------------------------------------------------------|
| IP       | IP address of a load balancer instance.                                                      |
| port     | Frontend protocol and port configured for the listener.                                      |
| user     | Username of the cluster. This parameter is required only for a security-mode cluster.        |
| pwd      | Password of the username above. This parameter is required only for a security-mode cluster. |

If cluster information is returned, the connection is successful.

See also: [Sample Code for Accessing a Cluster Using HTTPS Two-way Authentication Through a Load Balancer \(Java\)](#).

## Sample Code for Accessing a Cluster Using HTTPS Two-way Authentication Through a Load Balancer (Java)

This section provides sample code that can be used on a Java client to access a cluster using HTTPS two-way authentication through a load balancer.

## Sample Code for ES Secured Client With Cer Demo

```
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 ES Secured Client With Cer Demo {
 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 by using 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 {
 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, 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;
}
}
```

## Sample Code for SecuredHttpClientConfigCallback

```
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)
credentials 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 Sample 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>
```

## 2.10 Backing up and Restoring the Data of an Elasticsearch Cluster

### 2.10.1 Creating Snapshots to Back Up the Data of an Elasticsearch Cluster

CSS allows you to use snapshots to back up and restore Elasticsearch cluster data. By storing a snapshot in an OBS bucket, you save a point-in-time copy of the cluster's data. By restoring this snapshot, you can restore the cluster to a previous state. There are two ways to create snapshots to back up a CSS cluster: automatic and manual.

- **Automatic snapshot creation:** Snapshots are automatically created daily or weekly based on a preset time schedule, ensuring ongoing data protection. By configuring an automatic snapshot policy, you reduce manual operations while improving backup reliability and efficiency.
- **Manual snapshot creation:** You create snapshots manually for special occasions, for example, prior to a high-risk operation (such as a cluster upgrade). This ensures you can restore data using snapshots in case anything should go wrong. Manual snapshots provide additional flexibility.

#### Impact on Billing

When you use snapshots for data backup, the generated snapshots are stored in OBS buckets, which will result in additional costs. For details, see [Object Storage Service Billing](#).

#### Constraints

- Snapshots cannot be created for clusters that were created before March 10, 2018.
- Cluster snapshots will increase CPU usage and disk I/O. You are advised to take cluster snapshots during off-peak hours.
- When creating or restoring a snapshot, you can only perform certain operations, including accessing Kibana, viewing metrics, and deleting other snapshots. However, you cannot perform the following operations: modifying configuration, restarting or deleting the cluster, deleting the snapshot that is being created or restored, and creating or restoring another snapshot. While a snapshot is being created or restored for a cluster, any automatic snapshot creation task initiated for that cluster will be canceled.
- When you create the first snapshot in a cluster, a full backup is performed on all data. Subsequent snapshots are all incremental, storing only the changes since the previous snapshot. Because each incremental snapshot relies on its predecessors, you must restore them in the correct sequence.

- In a cluster where storage and compute are decoupled, an index that is both frozen and closed cannot be backed up.

## Prerequisites

- The OBS bucket used for storing cluster snapshots has been created. The OBS bucket must be in the same region as the target Elasticsearch cluster, and the storage class must be **Standard** or **Warm**.
- The cluster status is **Available**.

## Enabling Cluster Snapshots

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Cluster Snapshots** tab, and check whether cluster snapshots are enabled.
  - If they are enabled, skip this step.
  - If they are disabled, go to the next step.
5. On the **Cluster Snapshots** tab, click **Enable Snapshot**. In the displayed dialog box, configure the necessary settings.

**Figure 2-33** Enabling cluster snapshots




**Table 2-132** Enabling cluster snapshots

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBS Bucket | <p>From the drop-down list, select an OBS bucket for storing snapshots.</p> <p>If no OBS buckets meet your requirements, click <b>Create Bucket</b> to go to the OBS console and create one. For details, see <a href="#">Creating a Bucket</a>.</p> <p>To grant an IAM user access to an OBS bucket, you need to grant the <b>GetBucketStoragePolicy</b>, <b>GetBucketLocation</b>, <b>ListBucket</b>, and <b>ListAllMyBuckets</b> permissions to that user.</p> |

| Parameter                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup Path                      | <p>Snapshot storage path in the OBS bucket.</p> <p>The backup path cannot:</p> <ul style="list-style-type: none"><li>• Contain the following characters: \:*?"&lt;&gt; '{}.</li><li>• Start with a slash (/).</li><li>• Start or end with a period (.).</li><li>• Contain more than two consecutive slashes (/) or periods (.).</li><li>• Exceed 512 characters.</li></ul>                                                                                                                                                                                                                                                                                                                                                                |
| Maximum Backup Rate (per Second) | <p>The parameter sets the maximum backup rate per node. When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual backup rate may not reach the configured value, as it depends on factors such as OBS performance and disk I/O.</p> <p><b>Value format:</b> number + unit</p> <ul style="list-style-type: none"><li>• Number range: 0–9999</li><li>• Unit: KB, MB, GB, TB, PB, or B</li></ul> <p><b>Default value:</b> 40 MB</p> <p>The value 0MB means there is no limit on how fast data is backed up to snapshots. An overly high backup rate may lead to excessive resource usage, which may impact cluster stability. Therefore, set this parameter carefully.</p> |

| Parameter                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Maximum Recovery Rate (per Second) | <p>The parameter sets the maximum recovery rate per node. When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual recovery rate may not reach the configured value, as it depends on factors such as OBS performance and disk I/O.</p> <p><b>Value format:</b> number + unit</p> <ul style="list-style-type: none"><li>• Number range: 0–9999</li><li>• Unit: KB, MB, GB, TB, PB, or B</li></ul> <p><b>Default value:</b></p> <ul style="list-style-type: none"><li>• For Elasticsearch 7.6.2 clusters or earlier, the default value is 40 MB.</li><li>• For Elasticsearch clusters later than 7.6.2, the default value is <b>0 MB</b>, indicating no limit.</li></ul> <p>An overly high recovery rate may lead to excessive resource usage, which may impact cluster stability. Therefore, set this parameter carefully.</p> <p>For Elasticsearch clusters later than 7.6.2, the recovery rate is also limited by the <b>indices.recovery.max_bytes_per_sec</b> parameter.</p> <ul style="list-style-type: none"><li>• If <b>Maximum Recovery Rate (per Second)</b> is less than <b>indices.recovery.max_bytes_per_sec</b>, the former takes effect.</li><li>• If <b>Maximum Recovery Rate (per Second)</b> is greater than <b>indices.recovery.max_bytes_per_sec</b>, the latter takes effect.</li></ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>• To check the value of <b>indices.recovery.max_bytes_per_sec</b>, run the following command:<br/>GET _cluster/settings</li><li>• To modify <b>indices.recovery.max_bytes_per_sec</b>, run the following command:<br/>PUT _cluster/settings<br/>{<br/>  "transient": {<br/>    "indices.recovery.max_bytes_per_sec": "100mb"<br/>  }<br/>}</li></ul> |



| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IAM Agency | <p>Select an IAM agency to grant the current account the permission to access and use OBS. To store snapshots to an OBS bucket, you must have the required OBS access permissions.</p> <ul style="list-style-type: none"><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-obs-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the OBS Administrator permissions deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/>"obs:bucket:GetBucketLocation",<br/>"obs:object:GetObjectVersion",<br/>"obs:object:GetObject",<br/>"obs:object:DeleteObject",<br/>"obs:bucket:HeadBucket",<br/>"obs:bucket:GetBucketStoragePolicy",<br/>"obs:object:DeleteObjectVersion",<br/>"obs:bucket:ListBucketVersions",<br/>"obs:bucket:ListBucket",<br/>"obs:object:PutObject"</li><li>When OBS buckets use SSE-KMS encryption, the IAM agency must be granted KMS permissions. You can click <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b> to have the following custom policies created automatically.<br/>"kms:cmk:create",<br/>"kms:dek:create",<br/>"kms:cmk:get",<br/>"kms:dek:decrypt",<br/>"kms:cmk:list"</li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</li><li>To use an IAM agency, the following minimum permissions are required:<br/>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</li></ul> <p><b>WARNING</b><br/>The agency name can contain only letters (case-sensitive), digits, underscores (_), and hyphens (-). Otherwise, the backup will fail.</p> |

Determine whether to enable automatic snapshot creation based on service needs.

- To enable it, select **Automatic Snapshot Creation**, and configure necessary settings by referring to [Configuring Automatic Snapshot Creation](#).
- To disable it, unselect **Automatic Snapshot Creation**, in which case, you will need to manually create snapshots.

6. Click **OK** to enable cluster snapshots.  
The configuration information will be displayed on the **Cluster Snapshots** tab.

Configuring Automatic Snapshot Creation

You can enable automatic snapshot creation after cluster snapshots are already enabled.

1. On the **Cluster Snapshots** tab, click **Modify Settings**. In the displayed dialog box, select **Automatic Snapshot Creation**, and configure necessary settings.

Figure 2-34 Modifying Settings




Table 2-133 Configuring automatic snapshot creation

| Parameter            | Description                                                                                                                                                                                                                                                                                                                                                |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Snapshot Name Prefix | A snapshot name consists of the snapshot name prefix and timestamp, for example, <i>snapshot-1566921603720</i> .<br><br>The snapshot name prefix can contain 1 to 32 characters. It must start with a lowercase letter, and it can only contain lowercase letters, digits, hyphens (-), and underscores (_).<br><br>The default value is <b>snapshot</b> . |
| Time Zone            | Select a time zone for the backup start time.                                                                                                                                                                                                                                                                                                              |
| Backup Start Time    | Specify the start time of auto backup.<br><br>Select a value from the drop-down list. The interval can be <b>Daily</b> , <b>Hourly</b> , or weekly (by selecting a specific day of the week), and the backup time can be set to any hour from 00:00 to 23:00 (full hours only).                                                                            |

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Retained Snapshots | <p>Number of automatic snapshots to be retained. The value ranges from 1 to 90. The system automatically deletes excess snapshots every half hour. (The expiration deletion policy applies only to the snapshots that were automatically created at the same frequency as the current automatic snapshot creation policy.)</p> <p><b>NOTE</b></p> <p>If the snapshot creation interval is short or if the data size of indexes is large, the number of automatic snapshots retained may not reach the value set using this parameter.</p>                                                                                                                                                                                                                             |
| Index              | <p>Specify the name of the index to be backed up.</p> <ul style="list-style-type: none"><li>You can restore a specified index. To select multiple indexes, use commas (,) to separate them, for example, <b>index1,index2,index3</b>.</li><li>You can use an asterisk (*) to match multiple indexes. For example, <b>index*</b> indicates that all indexes with the prefix <b>index</b> will be backed up.</li></ul> <p>If you do not specify this parameter, all indexes in the cluster are backed up by default.</p> <p>The index name cannot contain spaces, uppercase letters, or special characters "<code>&lt; &gt;/?</code>".</p> <p><b>NOTE</b></p> <p>You can run the <b>GET /_cat/indices</b> command to query the names of all indexes in the cluster.</p> |

- Click **OK** to enable automatic snapshot creation.

After automatic snapshot creation is enabled, the **Automatic Snapshot Creation** parameter in the **Cluster Snapshots** area is shown as **Enabled**.

**Figure 2-35** Automatic snapshot creation enabled

#### Cluster Snapshots

Create snapshots to save point-in-time copies of the cluster's data to OBS buckets, enabling [Cluster snapshot user guide](#) 

OBS Bucket

12345678

IAM Agency

css\_obs\_agency

Backup Path

css\_repository/css-46e5-wangjing

Automatic Snapshot Creation

Enabled

Snapshots automatically created according to the snapshot policy as well as those manually created are displayed in the cluster snapshot task list.


When the snapshot status changes to **Available**, the snapshot is created. If snapshot creation fails, click on the right of the snapshot list to view the number of failed tasks and learn the failure causes. A maximum of 20 failed tasks can be displayed. When the snapshot function is disabled or the cluster is deleted, the failure records are also cleared.

3. Check whether data is successfully backed up.

In the cluster snapshot task list, click the snapshot name. The **View Details** dialog box is displayed. Check the shards and indexes that have been backed up to see if the backup is successful.

**Figure 2-36** View Details

### View Details




4. To disable automatic snapshot creation, click **Modify Settings**. In the displayed dialog box, deselect **Automatic Snapshot Creation**, and click **OK**. In the displayed dialog box, choose whether to delete snapshots created automatically. Then click **OK** to disable automatic snapshot creation.

### CAUTION

- When disabling automatic snapshot creation, if you select **Delete snapshots created automatically**, only existing snapshots will be deleted. Snapshots that are being created will still be created.
- If automatic snapshot creation is disabled without deleting automatically created snapshots, all such snapshots in the cluster cannot be manually deleted later once automatic snapshot creation is enabled again. Instead, they can only be deleted by the system based on the snapshot lifecycle policy. For example, if **Retained Snapshots** is set to **10** in the policy, the system automatically deletes excess snapshots (oldest first) on the next half hour when the number of existing snapshots exceeds this limit.

After automatic snapshot creation is disabled, the **Automatic Snapshot Creation** parameter in the **Cluster Snapshots** area is shown as **Disabled**.

Figure 2-37 Disabling automatic snapshot creation



Manually Creating a Snapshot

When snapshots are enabled for a cluster, you can manually create snapshots for it.

- 1. Under **Cluster Snapshot Tasks**, click **Manually Create Snapshot**. In the displayed dialog box, configure the snapshot policy.

Figure 2-38 Manually creating a snapshot




Table 2-134 Parameters for manually creating a snapshot

| Parameter     | Description                                                                                                                                                                                                                                                                                                                   |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Snapshot Name | <p>Set the snapshot name. Unlike automatically created snapshots, manually created snapshots do not have time information automatically added to their names.</p> <p>The value can contain 4 to 64 characters and must start with a letter. Only lowercase letters, digits, hyphens (-), and underscores (_) are allowed.</p> |

| Parameter            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Index                | <p>Specify the name of the index to be backed up.</p> <ul style="list-style-type: none"><li>You can restore a specified index. To select multiple indexes, use commas (,) to separate them, for example, <b>index1,index2,index3</b>.</li><li>You can use an asterisk (*) to match multiple indexes. For example, <b>index*</b> indicates that all indexes with the prefix <b>index</b> will be backed up.</li></ul> <p>If you do not specify this parameter, all indexes in the cluster are backed up by default.</p> <p>The index name cannot contain spaces, uppercase letters, or special characters "<code>&lt; &gt;/?</code>".</p> <p><b>NOTE</b><br/>You can run the <b>GET /_cat/indices</b> command to query the names of all indexes in the cluster.</p> |
| Snapshot Description | <p>Add a snapshot description.</p> <p>The value contains 0 to 256 characters and cannot contain special characters "<code>!&lt;&gt;=&amp;\</code>".</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

- Click **OK** to start creating snapshots.


Snapshots manually created as well as those automatically created according to the snapshot policy are displayed in the cluster snapshot task list.

When the snapshot status changes to **Available**, the snapshot is created. If snapshot creation fails, click on the right of the snapshot list to view the number of failed tasks and learn the failure causes. A maximum of 20 failed tasks can be displayed. When the snapshot function is disabled or the cluster is deleted, the failure records are also cleared.

- Check whether data is successfully backed up.

In the cluster snapshot task list, click the snapshot name. The **View Details** dialog box is displayed. Check the shards and indexes that have been backed up to see if the backup is successful.

Figure 2-39 View Details




Disabling Cluster Snapshots

If you no longer need the cluster snapshot function, you can disable it to reclaim storage resources.

On the **Cluster Snapshots** page, click **Disable Snapshot**. In the displayed dialog box, click **OK**. Disabling cluster snapshots does not automatically delete existing snapshots. Instead, you need to manually delete them on the OBS console.

Figure 2-40 Disable Snapshot



Related Documents

By creating and restoring snapshots, you can migrate data between clusters. For details, see [Migrating Data Between Huawei Cloud Elasticsearch Clusters Using Backup and Restoration](#).

2.10.2 Restoring the Data of an Elasticsearch Cluster Using a Snapshot

Use a snapshot to restore a point-in-time copy of an Elasticsearch cluster. You can restore it to the original cluster or to another eligible cluster. This process fetches

the backed-up index data, metadata, and shard structure from OBS and reconstruct them in the selected destination cluster, ensuring data consistency. This capability is essential for operations like data rollback, cluster migration, and disaster recovery.

## Constraints

- Cluster snapshots will increase CPU usage and disk I/O. You are advised to take cluster snapshots during off-peak hours.
- If the source cluster is unavailable, you can restore its snapshots to a designated destination cluster.
- When creating or restoring a snapshot, you can only perform certain operations, including accessing Kibana, viewing metrics, and deleting other snapshots. However, you cannot perform the following operations: modifying configuration, restarting or deleting the cluster, deleting the snapshot that is being created or restored, and creating or restoring another snapshot. While a snapshot is being created or restored for a cluster, any automatic snapshot creation task initiated for that cluster will be canceled.
- When you create the first snapshot in a cluster, a full backup is performed on all data. Subsequent snapshots are all incremental, storing only the changes since the previous snapshot. Because each incremental snapshot relies on its predecessors, you must restore them in the correct sequence.
- Before snapshot restoration is completed, the indexes being restored to the destination cluster cannot be queried.
- If the two clusters have indexes with the same names but different shard structures, such indexes cannot be restored.
- The version of the destination cluster must be the same as or later than that of the source cluster.

## Prerequisites

In the source cluster's snapshot task list, there are snapshots whose status is **Available**. The status of the destination cluster is also **Available**.

## Restoring Data

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Cluster Snapshots** tab.
5. In the cluster snapshot task list, select a snapshot, and click **Restore** in the **Operation** column. In the displayed dialog box, configure necessary settings.



Figure 2-41 Restoring a snapshot

| Created  | Operation                |
|-------------------------------------------------------------------------------------------|--------------------------|
| Jul 31, 2025 20:00:07 GMT+08:00                                                           | <div>RestoreDelete</div> |
| Jul 31, 2025 15:57:23 GMT+08:00                                                           | <div>RestoreDelete</div> |

Table 2-135 Restoration settings

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Index          | <p>Specify the name of the index you want to restore.</p> <p>Constraints:</p> <ul style="list-style-type: none"><li>• The value is a string of 0 to 1024 characters that cannot contain uppercase letters, spaces, or the following special characters: "<code>\&lt; &gt;/?</code>".</li><li>• When restoring an index whose name is prefixed with <b>.kibana</b>, the index name must be specified.</li><li>• The <b>.opendistro_security</b> index cannot be restored.</li></ul> <p>Value range:</p> <ul style="list-style-type: none"><li>• You can use an asterisk (*) to match multiple indexes. For example, <b>index*</b> indicates that all indexes with the prefix <b>index</b> will be restored. When an asterisk (*) is used for index matching, the <b>.opendistro_security</b> index and any system indexes whose name is prefixed with <b>.kibana</b> are filtered out by default.</li><li>• You can restore indexes by specifying their names, for example, <b>index1,index2,index3</b>.</li></ul> <p>Default value:</p> <p>By default, this parameter is left blank. That is, no index name is specified, and all indexes will be restored.</p> |
| Rename Pattern | <p>Index name matching rule. Enter a regular expression. Indexes that match the regular expression will be restored.</p> <p>The <b>Rename Pattern</b> and <b>Rename Replacement</b> take effect only when they are both configured at the same time.</p> <p>The value is a string of 0 to 1024 characters that cannot contain uppercase letters, spaces, or the following special characters: "<code>\&lt; &gt;/?</code>".</p> <p>The default value is <b>index_(.+)</b>, which indicates all indexes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

| Parameter                                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Rename Replacement                                 | <p>Rule for index renaming.</p> <p>The <b>Rename Pattern</b> and <b>Rename Replacement</b> take effect only when they are both configured at the same time.</p> <p>The value is a string of 0 to 1024 characters that cannot contain uppercase letters, spaces, or the following special characters: "<code>&lt;</code>"<code>&gt;</code>"<code>?</code>"<code>.</code>".</p> <p>The default value <b>restored_index_\$1</b> indicates that <b>restored_</b> will be added to the beginning of the names of all restored indexes.</p>                                                                                                                                                                                                                                                                                     |
| Cluster                                            | <p>Select the cluster where you want to restore the data to. You can select the current cluster or another.</p> <ul style="list-style-type: none"><li>You must specify a cluster whose status is <b>Available</b>. If the status of the current cluster is <b>Unavailable</b>, you cannot restore the snapshot to the current cluster.</li><li>When you choose to restore data to another cluster, make sure the target cluster has the same or later version than the current cluster, and that the two clusters are in the same region. After a frozen index is restored using a snapshot, the index is stored on the local node and is read-only by default. To make it writable, run the following command to configure the index:<br/><pre>PUT index_name/_settings {   "index.blocks.write": null }</pre></li></ul> |
| Overwrite same-name indexes in destination cluster | <p>Whether to overwrite same-name indexes in the destination cluster. By default, this option is unselected.</p> <p>Restoring a snapshot will overwrite existing data in the destination cluster. When there are same-name indexes in the destination cluster, you need to select this option in order to restore same-name, same-shard structure indexes. Indexes with a different shard structure cannot be restored.</p>                                                                                                                                                                                                                                                                                                                                                                                               |

- Click **OK** to start restoring the snapshot.

If restoration succeeds, **Task Status** of the snapshot in the snapshot task list will change to **Restoration succeeded**, and the index data captured in the snapshot is reconstructed.

**Figure 2-42** Restoration succeeded

| Name/ID           | Snapshot Status | Task Status           |
|-------------------|-----------------|-----------------------|
| snapshot-a671bb49 | Available       | --                    |
| snapshot-371bff61 | Available       | Restoration succeeded |

## Verifying Snapshot Restoration

After snapshot restoration, go to the destination cluster to check the restored index information.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
4. In the left navigation pane, choose **Dev Tools**.
5. Run the following command to check index information.

```
GET _cat/indices
```

## Helpful Links

By creating and restoring snapshots, you can migrate data between clusters. For details, see [Migrating Data Between Huawei Cloud Elasticsearch Clusters Using Backup and Restoration](#).

## 2.10.3 Deleting an Elasticsearch Cluster Snapshot

Delete snapshots that are no longer needed to reclaim storage spaces.

### Constraints

- Deleting a snapshot will permanently delete its data.
- The deletion of automatically generated snapshots is subject to the following constraints:
  - While automatic snapshot creation is enabled, the system automatically deletes snapshots on the nearest half hour after their retention period expires.
  - After automatic snapshot creation is disabled, existing auto-created snapshots will not be deleted automatically, but they can be deleted manually.
  - If automatic snapshot creation is re-enabled, the system resumes automatic deletion of all auto-created snapshots upon expiration, including these created before automatic snapshot creation was previously disabled.

## Manually Deleting a Snapshot

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Cluster Snapshots** tab.
5. In the cluster snapshot task list, select the snapshot you want to delete, and click **Delete** in the **Operation** column. In the displayed dialog box, enter **DELETE** and click **OK**.

## 2.11 Scaling an Elasticsearch Cluster

### 2.11.1 Scaling Out/Up an Elasticsearch Cluster

If an Elasticsearch cluster struggles to maintain performance in the face of rapid data growth or sustained high memory usage, you may scale it horizontally by adding more nodes and more node types, or do so vertically by increasing the storage capacity of existing nodes.

**Table 2-136** Scaling scenarios


| Type             | Scenario                                                                                                                                 | Change Process                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Adding new nodes | If a cluster experiences rapid data growth or sustained high memory usage, you may add more nodes to it to improve cluster availability. | <p>Procedure for adding data nodes or cold data nodes:</p> <ol style="list-style-type: none"><li>1. New nodes are added and cluster configuration is modified.</li><li>2. A shard redistribution is automatically triggered to reallocate certain shards to newly added nodes.</li><li>3. After the redistribution is complete, the new nodes start to handle query and indexing requests.</li></ol> <p>Adding master or client nodes does not trigger data migration.</p> <p>During the capacity expansion, the system ensures that each shard has at least one available replica to ensure service continuity.</p> |

| Type                             | Scenario                                                                                                                                                                                                                                                                                                   | Change Process                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Adding new node types            | For a cluster that has no master or client nodes, as the load on its data plane increases, you may add master or client nodes to it to share the load of the data nodes.                                                                                                                                   | <p>Procedure for adding a new node type: New nodes are added and cluster configuration is modified. Adding master or client nodes does not trigger data migration, and therefore does not interrupt services.</p> <p>When client nodes are added, the cluster's address changes from the data node address to the client node address. You must update the client configuration to use this new address. Otherwise, the client nodes cannot function.</p> |
| Increasing node storage capacity | <ul style="list-style-type: none"><li>• If a cluster experiences rapid data growth and its storage capacity becomes insufficient, you can expand the cluster nodes' storage capacity.</li><li>• If the disk usage of a cluster remains high, you can expand the cluster nodes' storage capacity.</li></ul> | <ol style="list-style-type: none"><li>1. A request is sent to the EVS service to expand the disk capacities for all cluster nodes.</li><li>2. After disk capacity is expanded by EVS, the size of existing disk volumes is expanded on all cluster nodes.</li></ol> <p>The disk capacity expansion will not interrupt ongoing services.</p>                                                                                                               |

Prerequisites

- The cluster status is **Available**, and there are no ongoing tasks.
- Your CSS resource quotas are sufficient for the capacity expansion you are about to perform. You can check available resources on the **Modify Configuration** page.

Figure 2-43 Checking available resources



## Impact on Billing

For a pay-per-use cluster, you can see its new price when confirming the scale-out or scale-up on the console. After the scale-out or scale-up is complete, the new price will apply. For pricing details, see [Cloud Search Service Price Calculator](#).

For a yearly/monthly cluster, a new order will be placed after you confirm the scale-out or scale-up. Then you can check the new price.

## Constraints

- The node storage capacity of a cluster can only be increased—not decreased. Select an appropriate node storage capacity based on the data volume and projected data growth.
- For a yearly/monthly cluster, you cannot increase the number of nodes and node storage capacity at the same time. For a pay-per-use cluster, however, you can do that.
- The storage capacity of the master and client nodes in a cluster cannot be expanded.
- The storage capacity of data nodes that use local disks cannot be expanded.
- For the range of node quantities supported by each node type, see [Table 2-137](#).

**Table 2-137** Node quantity ranges

| Node Type       | Value Range                                                                                                           |
|-----------------|-----------------------------------------------------------------------------------------------------------------------|
| Data nodes      | <ul style="list-style-type: none"><li>• Without master nodes: 1 to 32</li><li>• With master nodes: 1 to 200</li></ul> |
| Master nodes    | 3, 5, 7, or 9 (must be an odd number from 3 to 9)                                                                     |
| Client nodes    | 1–32                                                                                                                  |
| Cold data nodes | 1–32                                                                                                                  |

## Change Impact

Before the change, learn about possible impacts and operation suggestions, and develop a plan to minimize these impacts.

Expanding the storage capacity of cluster nodes (vertical scaling) does not impact services, whereas adding new nodes or new node types (horizontal scaling) may have the following impacts:

- **Impact on performance**  
Adding new nodes does not interrupt services. However, after new nodes are added, data shards need to be redistributed to these nodes to balance the load, and this process will consume I/O performance. This is why you are advised to perform the operation during off-peak hours.  
To minimize this impact, it is advisable to adjust the data migration rate based on the cluster's traffic cycle: increase the data migration rate during

off-peak hours to shorten the task duration, and decrease it **before** peak hours arrive to ensure optimal cluster performance. The data migration rate is determined by the **indices.recovery.max\_bytes\_per\_sec** parameter. The default value of this parameter is the number of vCPUs multiplied by 32 MB. For example, for four vCPUs, the data migration rate is 128 MB. Set this parameter to a value between 40 MB and 1000 MB based on your service requirements.

```
PUT /_cluster/settings
```

```
{
 "transient": {
 "indices.recovery.max_bytes_per_sec": "1000MB"
 }
}
```

- Characteristics of this process

Once started, a scaling task cannot be stopped until it succeeds or fails.

## Duration

- It takes 10 to 30 minutes to add new nodes or new node types (horizontal scaling), depending on the cluster's scheduling capabilities.
- It takes 10 to 15 minutes to expand the storage capacity of cluster nodes (vertical scaling).

## Adding More Nodes or Increasing Node Storage Capacity


If a cluster experiences rapid data growth or sustained high memory usage, you may add more nodes to it or increase the cluster nodes' storage capacity. This helps to improve service and data availability and durability.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. Click the **Scale Cluster** tab.
5. Set the necessary parameters.

**Table 2-138** Adding more nodes or increasing node storage capacity

| Parameter | Description                    |
|-----------|--------------------------------|
| Action    | Select <b>Scale out</b> .      |
| Resources | Quantities of resources added. |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nodes     | <p>Increase the number of nodes and node storage capacity in the <b>Nodes</b> and <b>Node Storage Type</b> columns. You can change multiple node types at the same time.</p> <ul style="list-style-type: none"><li>For the range of node quantities supported by each node type, see <a href="#">Constraints</a>.</li><li>The value range of node storage capacity is determined by the node flavor you select. The value must be divisible by 20.</li></ul> <p><b>NOTE</b><br/>For a yearly/monthly cluster, you cannot increase the number of nodes and node storage capacity at the same time. For a pay-per-use cluster, however, you can do that.</p> |

**Figure 2-44** Scaling out a cluster

6. Click **Next**.
7. Confirm the information and click **Submit**.
8. Click **Back to Cluster List** to go back to the **Clusters** page. The **Task Status** is **Scaling out**. When **Cluster Status** changes to **Available**, the cluster has been successfully scaled out.

## Adding New Node Types

For a cluster that has no master or client nodes, as the load on its data plane increases, you may add master or client nodes to it to share the load of the data nodes.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. On the **Modify Configuration** page, choose the **Add Master/Client Nodes** tab.  
If a cluster already has both master and client nodes, the **Add Master/Client Node** tab is unavailable.
5. On the Add Master/Client Nodes tab, configure the nodes.



**Table 2-139** Adding master or client nodes

| Parameter           | Description                                                                                                                                                                                                                                                                                                                          |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node types          | Select the type of nodes you want to add. <ul style="list-style-type: none"><li>Only one node type can be selected at a time. If you need to add both Master and Client nodes, you need to perform this task twice.</li><li>If the cluster already has Master or Client nodes, only the other node type is displayed here.</li></ul> |
| Node Specifications | Select node specifications based on site requirements.                                                                                                                                                                                                                                                                               |
| Nodes               | Set the number of nodes to add. For the value range, see <a href="#">Constraints</a> .                                                                                                                                                                                                                                               |
| Node Storage Type   | Set the node storage type. The storage capacity per master or client node is fixed to 40 GB.                                                                                                                                                                                                                                         |

**Figure 2-45** Adding master/client nodes

Master node ☒ Client node ☐

| Type        | Node Specifications     | Nodes | Node Storage Type |
|-------------|-------------------------|-------|-------------------|
| Master Node | ess.spec-4u8g(sold out) | 3     | Common I/O 40 GB  |

- Click **Next**.
- Confirm the information and click **Submit**.

Return to the cluster list page. The **Task Status** of the cluster is **Scaling out**.

- If you added master nodes, you can tell they are added successfully when **Cluster Status** changes to **Available**.

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 nodes take effect. Otherwise, the cluster may be reported as unavailable. (The cluster services still run properly.) For details, see [Restarting an Elasticsearch Cluster](#).

- If you added client nodes, you can tell they are added successfully when **Cluster Status** changes to **Available**. You can restart data nodes and cold data nodes to stop Cerebro and Kibana processes on these nodes.

## Related Documents

- You can also upgrade the node specifications and change the EVS disk type of an Elasticsearch cluster. For details, see [Changing the Node Specifications and Storage Type of an Elasticsearch Cluster](#).
- During a capacity expansion, if the resources in an AZ become insufficient, you can migrate the nodes in this AZ to another that has sufficient resources, and then perform the capacity expansion in the new AZ. For details, see [Switching AZs for an Elasticsearch Cluster](#).

## 2.11.2 Scaling In an Elasticsearch Cluster

If an Elasticsearch cluster has excess capacity—due to off-peak traffic or reduced data volume—you can reduce its nodes to optimize costs.

**Table 2-140** Scaling scenarios

| Type                     | Scenario                                           | Change Process                                                                                                                                                                                                                                       |
|--------------------------|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Removing nodes randomly  | Randomly removes cluster nodes to optimize costs.  | 1. Migrate the shards on the to-be-removed nodes to the remaining nodes.<br><br>2. After data migration, bring the nodes offline and modify the cluster configuration.<br><br>Nodes are removed one at a time, so as to avoid interrupting services. |
| Removing specified nodes | Removes specified cluster nodes to optimize costs. |                                                                                                                                                                                                                                                      |

### Prerequisites

- The cluster status is **Available**, and there are no ongoing tasks.
- All mission-critical data has been backed up.

### Impact on Billing

For a pay-per-use cluster, you can see its new price when confirming the scale-in on the console. After the scale-in is complete, the new price will apply. For pricing details, see [Cloud Search Service Price Calculator](#).

For a yearly/monthly cluster, a refund request will be triggered after you confirm the scale-in. Then you can check the new price.

### Constraints

- During a scale-in, the data on the to-be-removed nodes needs to be migrated to the remaining nodes. The timeout for data migration per node is 48 hours. Scale-in will fail if this timeout expires. When the cluster has large quantities of data, you are advised to manually adjust the data migration rate and avoid performing the migration during peak hours.
- For a cluster without master nodes, scale-in is allowed only if the number of data nodes plus cold data nodes is at least three. During a scale-in, you can only remove less than half of the data nodes plus cold data nodes, and the remaining number must be greater than the maximum number of index replicas.
- For a cluster with master nodes, scale-in is allowed only if the cluster has at least two data nodes. During a scale-in, you can only remove less than half of the cluster's master nodes.
- After the scale-in, the cluster nodes' disk usage must be less than 80%.

- After the scale-in, there has to be at least one of each type of node in each AZ. For a cross-AZ cluster, the difference between the numbers of the same-type nodes in different AZs cannot exceed 1.
- For the range of node quantities supported by each node type, see [Table 2-137](#).

**Table 2-141** Node quantity ranges

| Node Type       | Value Range                                                                                                           |
|-----------------|-----------------------------------------------------------------------------------------------------------------------|
| Data nodes      | <ul style="list-style-type: none"><li>• Without master nodes: 1 to 32</li><li>• With master nodes: 1 to 200</li></ul> |
| Master nodes    | 3, 5, 7, or 9 (must be an odd number from 3 to 9)                                                                     |
| Client nodes    | 1–32                                                                                                                  |
| Cold data nodes | 1–32                                                                                                                  |

## Change Impact

Before the change, learn about possible impacts and operation suggestions, and develop a plan to minimize these impacts.

- Impact on performance

During a scale-in, shards on the to-be-removed nodes are migrated to the remaining nodes. This process will consume I/O performance. This is why you are advised to perform the operation during off-peak hours.

To minimize this impact, it is advisable to adjust the data migration rate based on the cluster's traffic cycle: increase the data migration rate during off-peak hours to shorten the task duration, and decrease it **before** peak hours arrive to ensure optimal cluster performance. The data migration rate is determined by the **indices.recovery.max\_bytes\_per\_sec** parameter. The default value of this parameter is the number of vCPUs multiplied by 32 MB. For example, for four vCPUs, the data migration rate is 128 MB. Set this parameter to a value between 40 MB and 1000 MB based on your service requirements.

PUT /\_cluster/settings

```
{
 "transient": {
 "indices.recovery.max_bytes_per_sec": "1000MB"
 }
}
```

- Cluster load change

After a scale-in, the remaining nodes will need to handle all of the cluster's load. This may lead to higher CPU, memory, and disk I/O usage, impacting query and write performance. If shards are unevenly allocated, performance bottlenecks may occur. This is why before a scale-in, it is necessary to evaluate whether the remaining nodes have the capacity to handle the current cluster load.

- Characteristics of this process

Once started, a scaling task cannot be stopped until it succeeds or fails.

## Scale-in Duration

The following formula can be used to estimate how long a scale-in operation will take:

**Scale-in duration (min) = 5 (min) x Number of nodes to be removed + Data migration duration (min)**

where, 5 minutes indicates how long non-data migration operations (e.g., initialization) typically take per node. It is an empirical value.

**Data migration duration (min) = Total data size of the nodes to be removed (MB) ÷ [Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s)]**

where,

- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.
- The formulas above use estimates under ideal conditions. The actual migration speed depends on cluster load.

## Removing Nodes Randomly


1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. Click the **Scale Cluster** tab.
5. Click **Scale in** to set parameters.

**Table 2-142** Removing nodes randomly

| Parameter | Description                      |
|-----------|----------------------------------|
| Action    | Select <b>Scale in</b> .         |
| Resources | Quantities of resources reduced. |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>When a node is deleted, NICs are released. This means you need to have VPC permissions. Select an IAM agency to grant the current account the permission to access and use VPC.</p> <ul style="list-style-type: none"><li>This parameter is available only when the new IAM plane is connected.</li><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-upgrade-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>VPC Administrator</b> role or the <b>VPC FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/><pre>"vpc:subnets:get",<br/>"vpc:ports:*"</pre></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</pre></li><li>To use an IAM agency, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</pre></li></ul> |
| Nodes     | <p>Reduce the number of nodes in the <b>Nodes</b> column. You can change multiple node types at the same time.</p> <p>For the range of node quantities supported by each node type, see <a href="#">Constraints</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Figure 2-46 Scaling in a cluster



6. Click **Next**.
7. Confirm the information and click **Submit**.
8. Click **Back to Cluster List** to go back to the **Clusters** page. **Task Status** is **Scaling in**. When **Cluster Status** changes to **Available**, the cluster has been successfully scaled in.

## Removing Specified Nodes

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. On the **Modify Configuration** page, click the **Scale In Nodes** tab.
5. Set scale-in parameters.

**Table 2-143** Removing specified nodes (Scale In Nodes)

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>When a node is deleted, NICs are released. This means you need to have VPC permissions. Select an IAM agency to grant the current account the permission to access and use VPC.</p> <ul style="list-style-type: none"><li>• This parameter is available only when the new IAM plane is connected.</li><li>• If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-upgrade-agency</b>.</li><li>• If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>VPC Administrator</b> role or the <b>VPC FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/><pre>"vpc:subnets:get",<br/>"vpc:ports:*"</pre></li><li>• To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</pre></li><li>• To use an IAM agency, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</pre></li></ul> |
| Node Type | Expand the node type that needs be changed to show all nodes under it. Select the nodes you want to remove.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

6. Click **Next**.

7. Confirm the change information and click **Submit**. In the confirm dialog box, choose to migrate data, which helps to prevent data loss, and click **OK**.

During data migration, the system migrates all data from the to-be-removed nodes to the remaining nodes, and removes these nodes upon completion of the data migration. If the data on the to-be-removed nodes has replicas on other nodes, data migration can be skipped and the cluster change can be completed faster.

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

## Related Documents

- For an Elasticsearch cluster, you can also optimize costs by changing node specifications and EVS disk types. For details, see [Changing the Node Specifications and Storage Type of an Elasticsearch Cluster](#).
- If you want to reduce cluster nodes even though your cluster is not eligible for a scale-in operation, you can simply create a new cluster. Then you can migrate data to the new cluster using snapshots, and delete the old cluster after data is migrated.

## 2.12 Upgrading the Version of an Elasticsearch Cluster

In CSS, you can upgrade the version of an Elasticsearch cluster to use feature and performance enhancements or fix known issues.

**Table 2-144** Upgrade scenarios

| Upgrade Type         | Scenario                                                                                                                                                                                                                                                                                                                                                                                                                      | Upgrade Process                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Same-version upgrade | Update the cluster's kernel patches. The cluster is upgraded to the latest image of the current version to fix known issues or optimize performance. For example, if the current cluster version is <b>7.10.2(7.10.2_24.3.3_0102)</b> , upon a same-version upgrade, the cluster will be upgraded to the latest image <b>7.10.2(7.10.2_24.3.4_0109)</b> of version 7.10.2. (The version numbers used here are examples only.) | <ol style="list-style-type: none"><li>1. Bring a node offline, and migrate its data to other nodes.</li><li>2. Upgrade the node to the target version.</li><li>3. Add the node back into the cluster. The system automatically triggers a shard reallocation, migrating some of the shards to the new node.</li><li>4. Confirm that the node status is normal. Then upgrade the remaining nodes one by one.</li></ol> <p>Nodes in clusters are upgraded one by one to avoid interrupting services.</p> |



| Upgrade Type          | Scenario                                                                                                                                                                                                                                                                                                                                                                                                               | Upgrade Process |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Cross-version upgrade | Upgrade a cluster to the latest image of the target version to enhance functionality or incorporate versions. For example, if the current cluster version is <b>7.6.2(7.6.2_24.3.3_1224)</b> , upon a cross-version upgrade, the cluster will be upgraded to the latest image <b>7.10.2(7.10.2_24.3.4_0109)</b> of version 7.10.2. (The version numbers used here are examples only.)                                  |                 |
| Cross-engine upgrade  | Upgrade an Elasticsearch cluster to the latest image of the target OpenSearch version. For example, if the Elasticsearch cluster version is <b>7.10.2(7.10.2_24.3.3_0102)</b> , upon a cross-engine upgrade, the Elasticsearch cluster will be upgraded to the latest image <b>1.3.6(1.3.6_24.3.4_0109)</b> of OpenSearch 1.3.6. (The version numbers used here are examples only.) This feature is not yet supported. |                 |

Version Support

The supported target versions vary depending on the current version. For details, see [Table 2-145](#).

Table 2-145 Version upgrade paths

| Current Version       | Target Version                 |
|-----------------------|--------------------------------|
| Elasticsearch: 6.2.3  | Elasticsearch: 6.5.4 or 6.8.23 |
| Elasticsearch: 6.5.4  | Elasticsearch: 6.8.23          |
| Elasticsearch: 6.8.23 | Elasticsearch: 7.6.2 or 7.10.2 |
| Elasticsearch: 7.1.1  | Elasticsearch: 7.6.2 or 7.10.2 |
| Elasticsearch: 7.6.2  | Elasticsearch: 7.10.2          |
| Elasticsearch: 7.9.3  | Elasticsearch: 7.10.2          |

| Current Version                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Target Version |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| <p>Note:</p> <ul style="list-style-type: none"><li>Elasticsearch 7.6.2 and 7.10.2 are the current major cluster versions. You are advised to upgrade your clusters to these two versions. The supported target versions are displayed in the drop-down list of <b>Target Image</b>.</li><li>Elasticsearch clusters of version 5.X do not support cross-version upgrade. Elasticsearch clusters of versions 6.2.3 and 6.5.4 can be upgraded to 6.8.23 and then to 7.X.X.</li></ul> |                |

## Constraints

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

## Upgrade Impact

Before upgrading a cluster, it is essential to assess the potential impacts and review operational recommendations. This enables proper scheduling of the upgrade time, ensuring a smooth upgrade process and minimizing service interruptions.

- Impact on performance

The nodes of a cluster are upgraded one at a time to ensure service continuity. However, data migration that occurs during the upgrade consumes I/O performance, and taking individual nodes offline still has some impact on the overall cluster performance.

To minimize this impact, it is advisable to adjust the data migration rate based on the cluster's traffic cycle: increase the data migration rate during off-peak hours to shorten the task duration, and decrease it **before** peak hours arrive to ensure optimal cluster performance. The data migration rate is determined by the **indices.recovery.max\_bytes\_per\_sec** parameter. The default value of this parameter is the number of vCPUs multiplied by 32 MB. For example, for four vCPUs, the data migration rate is 128 MB. Set this parameter to a value between 40 MB and 1000 MB based on your service requirements.

```
PUT /_cluster/settings
```

```
{
 "transient": {
 "indices.recovery.max_bytes_per_sec": "1000MB"
 }
}
```

- Impact on request handling

During the upgrade, nodes are upgraded one at a time. Requests sent to the node that is being upgraded may fail. To mitigate this impact, the following measures may be taken:

- Use a VPC endpoint or a dedicated load balancer to handle access requests to your cluster, which makes sure that requests are automatically routed to available nodes.
- Enable an exponential backoff & retry mechanism on the client (configure three retries).
- Perform the upgrade during off-peak hours.
- Rebuilding Kibana and Cerebro  
Kibana and Cerebro will be rebuilt during the upgrade, making them temporarily unavailable. Additionally, due to cross-version compatibility issues, Kibana may become unavailable during the upgrade. These problems will go away once the upgrade is completed.
- Characteristics of the upgrade process  
Once started, an upgrade task cannot be stopped until it succeeds or fails. An upgrade failure only impacts a single node, and does not interrupt services if there are data replicas. If necessary, you can restore a node that has failed to be upgraded by performing [Replacing a Specified Node for an Elasticsearch Cluster](#).

## Upgrade Duration

The following formula can be used to estimate how long it takes to upgrade a cluster:

**Upgrade duration (min) = 15 (min) x Total number of nodes to be upgraded + Data migration duration (min)**

where,

- 15 minutes indicates how long non-data migration operations (e.g., initialization and image upgrade) typically take per node. It is an empirical value.
- The total number of nodes is the sum of the number of data nodes, master nodes, client nodes, and cold data nodes in the cluster.

**Data migration duration (min) = Total data size (MB)/[Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s) x Concurrency]**

where,

- Typically, when concurrency is higher than 1, the actual data migration rate can only reach half of the throughput intended by the configured concurrency, depending on the cluster load.
- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.

The formulas above use estimates under ideal conditions. In practice, it is advisable to add 20% to 30% of redundancy.

## Pre-Upgrade Check

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

**Table 2-146** 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 <b>green</b> or <b>yellow</b> can provide services properly and have no unallocated primary shards.                                                                                                                                                                     | The cluster is available and there are no ongoing tasks.                                                                                                                                                                                                                                                             |
| Node quantity  | System check | After an upgrade task is started, the system automatically checks node quantities. To ensure service continuity, there has to be at least two nodes of each type for a cluster in each of its AZs. For a cluster with master nodes, there has to be at least two data nodes. For a cluster without master nodes, the number of data nodes plus cold data nodes must be three at least. | <ul style="list-style-type: none"><li>• Number of nodes of each type in each AZ <math>\geq 2</math></li><li>• For a cluster with master nodes: Number of data nodes <math>\geq 2</math></li><li>• For a cluster without master nodes: Number of data nodes + Number of cold data nodes <math>\geq 3</math></li></ul> |
| 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 total disk capacity of the remaining nodes is sufficient for handling all the data of the cluster, and that the nodes' disk usage stays below 80%.                                       | After a single node is brought offline, the total disk capacity of the remaining nodes is sufficient for handling all the data of the cluster, and the nodes' disk usage stays below 80%.                                                                                                                            |
| Data replicas  | System check | Check that the remaining data nodes and cold data nodes can handle the maximum number of primary and standby shards of indexes in the cluster. There should not be unallocated replicas during the upgrade.                                                                                                                                                                            | Number of data nodes + Number of cold data nodes > Maximum number of index replicas + 1                                                                                                                                                                                                                              |

| Check Item     | Check Method            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Normal Status                                                                          |
|----------------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|
| Data backup    | System check            | Before the upgrade, back up data to prevent data loss caused by upgrade failures. When submitting an upgrade task, you can choose whether to check for full index snapshots.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Check whether data has been backed up.                                                 |
| Resources      | System check            | After an upgrade task is started, the system automatically checks resources. New resources will be created during the upgrade. Ensure that resources are available.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Resources are available and sufficient.                                                |
| Custom plugins | System and manual check | <p>Perform this check only if custom plugins are installed in the source cluster. If a cluster has a custom plugin, upload all plugin packages of the target version on the plugin management page before the upgrade. During the upgrade, install the custom plugin in the new nodes. Otherwise, the custom plugins will be lost after the cluster is successfully upgraded. 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><b>NOTE</b></p> <p>If the uploaded plugin package is incorrect or incompatible, the plugin package cannot be automatically installed during the upgrade. As a result, the upgrade task fails. To restore a cluster, you can terminate the upgrade task and restore the node that fails to be upgraded by performing <a href="#">Replacing a Specified Node for an Elasticsearch Cluster</a>.</p> <p>After the upgrade is complete, the status of the custom plugin is reset to <b>Uploaded</b>.</p> | The plugin package of the cluster to be upgraded has been uploaded to the plugin list. |

| Check Item              | Check Method            | Description                                                                                                                                                                                                                                                                                                                           | Normal Status                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Custom configurations   | System check            | During the upgrade, the system automatically synchronizes the content of the cluster configuration file <b>elasticsearch.yml</b> .                                                                                                                                                                                                    | Clusters' custom configurations are not lost after the upgrade.                                                                                                                                                                                                                                                                                                  |
| Non-standard operations | Manual check            | Check whether non-standard operations have been performed in the cluster. Non-standard operations refer to manual operations that are not recorded. These operations cannot be automatically passed on during the upgrade, for example, modification of the <b>kibana.yml</b> configuration file, system settings, and return routes. | 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 <b>kibana.yml</b> file, cannot be retained, and you must back up the file in advance. |
| 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.                                                                                                                                                                                                                                                                                        |

| Check Item          | Check Method            | Description                                                                                                                                                                                                                                                                                                                | Normal Status                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Check Cluster Loads | System and manual check | <p>If the cluster is heavily loaded, there is a high probability that the upgrade will get stuck or fail. You are advised to check the cluster load before the upgrade and perform the upgrade only during off-peak hours.</p> <p>You can also choose to check the cluster load while configuring upgrade information.</p> | <ul style="list-style-type: none"><li>• <b>nodes.thread_pool.search.queue &lt; 1000</b>: Check that the maximum number of requests in the search queue is less than 1000.</li><li>• <b>nodes.thread_pool.write.queue &lt; 200</b>: Check that the maximum number of requests in the write queue is less than 200.</li><li>• <b>nodes.process.cpu.percent &lt; 90</b>: Check that the maximum CPU usage is less than 90%.</li><li>• <b>nodes.os.cpu.load_average/Number of vCPUs &lt; 80%</b>: Check that the number of running processes plus the number of processes waiting for CPUs is less than 80% of the total number of vCPUs.</li></ul> |

## Performing the Upgrade Task

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Cluster Snapshots** tab, and perform a full data backup. For details, see [Manually Creating a Snapshot](#).

When creating an upgrade task, you can choose to check whether the full index data has been backed up using snapshots. This helps to prevent data loss in case of an upgrade failure.

5. Click the **Version Upgrade** tab, and set upgrade parameters.

**Table 2-147** Upgrade parameters

| Parameter    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Upgrade Type | Select an upgrade type. <ul style="list-style-type: none"><li>• Same-version upgrade: upgrade kernel patches to the latest images within the current cluster version.</li><li>• Cross-version upgrade: upgrade a cluster to the latest image of the target version.</li><li>• Cross-engine upgrade: Upgrade an Elasticsearch cluster to an OpenSearch cluster. This feature is not yet supported.</li></ul>                                                                                                                                                                                                             |
| Target Image | Image of the target version. When you select an image, the image name and target version details are displayed.<br>The supported target versions are displayed in the drop-down list of <b>Target Image</b> . If no target image is available, possible causes are as follows: <ul style="list-style-type: none"><li>• The current cluster is of the latest version.</li><li>• The current cluster was created before 2023 and has vector indexes.</li><li>• The new version images are unavailable in the current region.</li><li>• The current cluster does not support the upgrade type you have selected.</li></ul> |



| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>After a node is upgraded, NICs need to be reattached to it. To do that, you must have the permission to access VPC resources. By configuring an IAM agency, you can authorize CSS to access its VPC resources through an associated account.</p> <ul style="list-style-type: none"><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-upgrade-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>VPC Administrator</b> role or the <b>VPC FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/><pre>"vpc:subnets:get",<br/>"vpc:ports:*"</pre></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</pre></li><li>To use an IAM agency, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</pre></li></ul> |

6. Click **Submit**.
7. In the displayed dialog box, configure upgrade check and acceleration options.
  - a. Select whether to enable **Check Full Index Snapshot**.

When enabled, CSS checks for valid snapshots by matching indexes by name. Snapshots help prevent potential data loss caused by upgrade failures.

CSS cannot check the content or backup times of snapshots. You should manually check existing snapshots. If any of them is over one month old, create the latest snapshot.
  - b. Select whether to enable **Check Cluster Loads**.

During an upgrade, data migration and node restarts will consume cluster resources and increase loads. When this option is enabled, CSS evaluates overload risks for the cluster and reduce the likelihood of a cluster upgrade failure caused by an overload.

The check items are as follows:

- **nodes.thread\_pool.search.queue < 1000**: Check that the maximum number of requests in the search queue is less than 1000.
- **nodes.thread\_pool.write.queue < 200**: Check that the maximum number of requests in the write queue is less than 200.
- **nodes.process.cpu.percent < 90**: Check that the maximum CPU usage of cluster nodes is less than 90%.
- **nodes.os.cpu.load\_average/Number of vCPUs < 80%**: Check that the number of running processes plus the number of processes waiting for CPUs is less than 80% of the total number of vCPUs.

If any of the results is abnormal, wait until the load drops or actively optimize it before performing the upgrade.


c. Set **Data migration concurrency control**.

Increasing the data migration concurrency can accelerate the upgrade process, but faster data migration leads to higher I/O usage. A higher concurrency is likely to lead to higher cluster load, which may impact cluster performance. You are advised to retain the default value **1**. The value should not exceed half of the number of data nodes.

8. Click **OK** to start the pre-upgrade check. The system automatically performs the check based on **Pre-Upgrade Check**, and if the check is successful, proceeds to the upgrade.

If the upgrade check fails, error information is displayed in the upper right corner of the console. Adjust the cluster configuration accordingly and try again.

**Figure 2-47** Upgrade check errors



When **Task Status** in the task list below changes to **Successful**, the upgrade is completed.

9. Confirm the upgrade result.
- a. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  - b. On the **Overview** tab, check the cluster version in the **Cluster Information** area to determine whether the upgrade is successful.

**Figure 2-48** Checking the cluster version

## Checking the Update Task

On the **Version Upgrade** page, find the current upgrade task in the **Jobs** list.

Expand the task, and click **View Progress** to check the upgrade progress.

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.

After an upgrade task is terminated, nodes that have failed to be upgraded will remain in the current state, while those that have been successfully upgraded will not be rolled back to the old version. As a result, the nodes in the same cluster may have different versions. You are advised to restart the upgrade as soon as relevant issues are handled, ensuring that all nodes have the same version.

---

### WARNING

During a cross-version upgrade, if **Task Status** is **Failed**, you can terminate the upgrade task only if all nodes have not been upgraded.

---

## 2.13 Managing Elasticsearch Clusters

### 2.13.1 Viewing Elasticsearch Cluster Information

On the CSS management console, you can check an Elasticsearch cluster's status, version, nodes, configurations and settings, and resource allocation on the cluster list and the cluster information page.

#### Introduction to Cluster Lists

A cluster list displays all clusters of the same type under the current account. If there are a large number of clusters, the cluster list will consist of multiple pages, showing clusters of all states.

Clusters are listed in chronological order by default in the cluster list, with the most recent clusters displayed at the top.

You can click **Export** to download the full cluster list.

**Table 2-148** Cluster list parameters

| Parameter               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name/ID                 | Name and ID of a cluster. You can click a cluster name to switch to the <b>Cluster Information</b> page. The cluster ID is automatically generated by the system and uniquely identifies a cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Cluster Status          | <p>Current status of a cluster.</p> <ul style="list-style-type: none"><li>• Available: The cluster is running properly.</li><li>• Creating: The cluster is being created.</li><li>• Processing: The cluster is being restarted, scaled out, backed up, or restored.</li><li>• Abnormal: The cluster failed to be created or is unavailable.</li></ul> <p><b>WARNING</b></p> <p>If a cluster is unavailable, you can delete the cluster or use a snapshot created when the cluster was available to restore its data to another cluster. However, you cannot scale out this cluster, access it through Kibana, create snapshots for it, or use an existing snapshot to restore its data to itself. Do not try to import data into a cluster that is unavailable, as the data may get lost. For such a cluster, you can check its monitoring data or restart it to try to fix the problem. If the fault persists, contact technical support.</p> |
| Task Status             | Status of a task, such as cluster restart, scale-out, or scale-in.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Version                 | <p>Cluster version and image version.</p> <p>For example, if the version is <b>7.10.2(7.10.2_x.x.x_x.x.x)</b>, <b>7.10.2</b> is the cluster version and <b>7.10.2_x.x.x_x.x.x</b> enclosed in parentheses is the image version.</p> <p><b>CAUTION</b></p> <ul style="list-style-type: none"><li>• For a cluster created a long time ago, its image version may not be available because it is likely not recorded in the database.</li><li>• If <b>EOS</b> is displayed below the version number, the current version has reached EOS and hence poses certain service availability risks. You are advised to upgrade the cluster version as soon as possible. For details, see <a href="#">Upgrading the Version of an Elasticsearch Cluster</a>.</li></ul>                                                                                                                                                                                    |
| Private Network Address | Private network address and port number of a cluster. You can use it to access the cluster. If the cluster has multiple nodes, the private network addresses and port numbers of all nodes are displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| VPC Endpoint Address    | IP address or private domain name of the VPC endpoint associated with the cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Security Mode           | Whether the security mode is enabled for the cluster. Only authorized users can access a security-mode cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enterprise Project  | Enterprise project that a cluster belongs to.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Associated Cluster  | <p>Name and ID of a cluster associated with the current cluster.</p> <ul style="list-style-type: none"><li>• The associated cluster is the parent cluster of the current cluster. For example, when KooSearch uses an Elasticsearch cluster as a vector database, the KooSearch service is the associated cluster of this Elasticsearch cluster.</li><li>• If any of the following operations is performed on the associated cluster, it will also be performed on the current cluster in a synchronized manner: deletion, unsubscription, renewal, and change of billing mode.</li><li>• The following operations cannot be performed alone on an Elasticsearch cluster that has an associated cluster: deletion, unsubscription, renewal, and change of billing mode or security mode.</li><li>• You can click the name of an associated cluster to go to its details page and check its configuration and status.</li></ul> |
| Cluster Description | Provides supplementary information about the cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Created             | Cluster creation time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Billing Mode        | Billing mode of a cluster. It can be <b>Pay-per-use</b> or <b>Yearly/Monthly</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Operation           | Operations that can be performed on a cluster, such as checking metrics, restart, and deletion. If an operation is unavailable, the menu button is grayed out.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## Viewing Basic Cluster Information

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, check the cluster's basic information, configuration, and node information.

**Table 2-149** Cluster information

| Type                | Parameter | Description                                                                                                                                                                     |
|---------------------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Information | Name      | <p>Custom cluster name.</p> <p>You can click  on the right to change the cluster name.</p> |

| Type          | Parameter                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               | Initial Name                  | Name of the cluster when it was created for the first time, which is unique and cannot be changed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|               | ID                            | Unique ID of a cluster, which is automatically generated by the system.<br>Each cluster in the same region has a unique ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|               | Cluster Description           | Provides supplementary information about the cluster. You can click  on the right to modify the description.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|               | Version                       | Cluster version.<br>For example, if the version is <b>7.10.2(7.10.2_x.x.x_x.x.x)</b> , <b>7.10.2</b> is the cluster version and <b>7.10.2_x.x.x_x.x.x</b> enclosed in parentheses is the image version.<br><b>CAUTION</b> <ul style="list-style-type: none"><li>For a cluster created a long time ago, its image version may not be available because it is likely not recorded in the database.</li><li>If <b>EOS</b> is displayed below the version number, the current version has reached EOS and hence poses certain service availability risks. You are advised to upgrade the cluster version as soon as possible. For details, see <a href="#">Upgrading the Version of an Elasticsearch Cluster</a>.</li></ul> |
|               | Cluster Status                | Current status of a cluster                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|               | Task Status                   | Current task status of a cluster. If there is no ongoing task, -- 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 a cluster belongs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

| Type | Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|      | Current Subnet | <p>Subnet to which a cluster belongs.</p> <p>If the current subnet cannot meet your service requirements, click <b>Switch Subnet</b>. In the displayed dialog box, select another subnet. If no existing subnets meet your requirements, click <b>Create Subnet</b> to create a new subnet.</p> <p><b>CAUTION</b></p> <ul style="list-style-type: none"><li>• After switching the subnet, the new subnet is used for adding nodes, including scaling, adding dedicated master/client nodes, and enabling VPC Endpoint. IP addresses will be allocated to newly added nodes from the new subnet.</li><li>• A cluster can be associated with a maximum of two subnets.</li><li>• Subnet switching does not affect the network settings of existing nodes.</li></ul> |
|      | Used Subnet    | Subnet once used by the cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|      | Security Group | <p>Security group to which a cluster belongs.</p> <p>To modify the security group of a cluster, click <b>Change Security Group</b> on the right.</p> <p><b>WARNING</b></p> <p>The selected security group must allow port 9200 in the inbound direction. Otherwise, the cluster may be inaccessible to external services. Furthermore, to enable Kibana private network access, port 5601 must also be allowed.</p>                                                                                                                                                                                                                                                                                                                                               |
|      | Security Mode  | <p>Security mode of a cluster.</p> <ul style="list-style-type: none"><li>• Enabled: The security mode is enabled for the current cluster.</li><li>• Disabled: The security mode is disabled for the current cluster.</li></ul> <p>For details about how to change the security mode of a cluster, see <a href="#">Changing the Security Mode of an Elasticsearch Cluster</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                 |

| Type | Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|      | Reset Password     | <p>This parameter is available only for security-mode clusters.</p> <p>Click <b>Reset</b> to change the <b>admin</b> password for a security-mode cluster.</p> <p>Requirements for admin passwords:</p> <ul style="list-style-type: none"><li>• The password can contain 8 to 32 characters.</li><li>• 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 allowed: ~!@#\$%^&amp;*()-_+=\ []{};,:&lt;.&gt;/?</li><li>• Do not use the administrator name, or the administrator name spelled backwards, as the password.</li><li>• Change the password periodically.</li></ul> |
|      | Cluster Routing    | <p>Cluster route information. You can view, add, or modify routes for a cluster. For details, see <a href="#">Configuring Routes for an Elasticsearch Cluster</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|      | Enterprise Project | <p>Enterprise project that a cluster belongs to.</p> <p>You can click a project name to check the basic project information.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |



| Type | Parameter                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|      | Public IP Address             | <p>Public IP address that enables public network access. It is displayed only for security-mode clusters.</p> <ul style="list-style-type: none"><li>For a security-mode cluster with public network access enabled, the configured public IP address is displayed.</li><li>For a security-mode cluster with public network access disabled, -- is displayed.</li></ul> <p>When public network access is enabled for a cluster, you are advised to use a whitelist for access control. For details about how to configure public network access, see <a href="#">Configuring Public Network Access for an Elasticsearch Cluster</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                     |
|      | Public Network Access Control | <p>You can configure public network access control for a security-mode cluster by IP addresses or IP address ranges. This parameter is displayed only for clusters with public network access enabled.</p> <p>Click <b>Modify</b> to modify the access control whitelist.</p> <ul style="list-style-type: none"><li>If a whitelist is configured, only IP addresses that are on this whitelist can access the cluster over the public network. Click <b>+Add</b>. In the displayed text box, enter IP addresses or CIDR blocks that are allowed to access the cluster from the public network. Separate them using commas (.). Each value must be unique. An example of valid values: <b>192.168.1.1,10.0.0.0/24</b>. Examples of invalid values: <b>0.0.0.0, xx.xx.xx.x0, 172.16.0.0-172.16.255.255</b>, non-standard formats (e.g., <b>192.168.1</b>), and duplicate values.</li><li>If no whitelist is configured, all public IP addresses can access the cluster. However, this can be a security risk and should be avoided.</li></ul> |

| Type | Parameter            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|      | Bandwidth            | <p>The bandwidth for public network access. This parameter is displayed only for clusters with public network access enabled.</p> <p>Click <b>Edit</b> to change the bandwidth size.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|      | HTTPS Access         | <p>Whether to enable HTTPS for a cluster.</p> <ul style="list-style-type: none"> <li>If HTTPS is disabled, the cluster uses HTTP instead. For a cluster with the security mode disabled, HTTPS access is disabled by default.</li> <li>If it is enabled, HTTPS is used for cluster access. Only security-mode clusters support HTTPS. If <b>HTTPS Access</b> is enabled, you can click <b>Download Certificate</b> to obtain the certificate used for access control. Currently, the security certificate cannot be used in a public network environment.</li> </ul> <p>For details about how to switch the network protocol for a security-mode cluster, see <a href="#">Switching between HTTP and HTTPS in Security Mode</a>.</p> |
|      | Private IPv4 Address | <p>Private IPv4 address and port number of a cluster. If the cluster has only one node, the IP address and port number of the only node are displayed, for example, <b>10.62.179.32:9200</b>. If the cluster has multiple nodes, the IP addresses and port numbers of all nodes are displayed, for example, <b>10.62.179.32:9200,10.62.179.33:9200</b>.</p>                                                                                                                                                                                                                                                                                                                                                                          |
|      | Private IPv6 Address | <p>Private IPv6 address and port number of the cluster. They are displayed only when the VPC selected for the cluster supports IPv6 addresses and that an IPv6 address has been assigned automatically.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

| Type             | Parameter               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                  | Associated Cluster Name | <p>Name of the cluster associated with the current cluster.</p> <ul style="list-style-type: none"><li>• The associated cluster is the parent cluster of the current cluster. For example, when KooSearch uses an Elasticsearch cluster as a vector database, the KooSearch service is the associated cluster of this Elasticsearch cluster.</li><li>• If any of the following operations is performed on the associated cluster, it will also be performed on the current cluster in a synchronized manner: deletion, unsubscription, renewal, and change of billing mode.</li><li>• The following operations cannot be performed alone on an Elasticsearch cluster that has an associated cluster: deletion, unsubscription, renewal, and change of billing mode or security mode.</li><li>• You can click the name of an associated cluster to go to its details page and check its configuration and status.</li></ul> |
|                  | Associated Cluster ID   | ID of the cluster associated with the current cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Node Information | 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                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## 2.13.2 Creating Users for an Elasticsearch Cluster and Granting Cluster Access

CSS limits access to security-mode clusters to authorized users only. When creating a security-mode cluster, an administrator account must be created. This administrator account can later use Kibana to add new users for the cluster and grant them the required permissions. This topic uses Kibana 7.10.2 as an example to describe how to use Kibana to grant users access to a security-mode cluster.

## Background

CSS uses the `opendistro_security` plug-in to provide security cluster capabilities. The `opendistro_security` plug-in is built based on the RBAC model. RBAC involves three core concepts: user, action, and role. RBAC simplifies the relationship between users and actions, simplifies permission management, and facilitates permission expansion and maintenance. [Figure 2-49](#) shows the relationship between the three.

**Figure 2-49** User, action, and role



**Table 2-150** Concepts

| Concept      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| User         | A user can send operation requests to an Elasticsearch cluster. The user has credentials such as username and password, and zero or multiple backend roles and custom attributes.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Role         | A role is a combination of permissions or action groups, including operation permissions on clusters, indexes, documents, or fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Permission   | A single permission, for example, creating an index (for example, <b>indices:admin/create</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Role mapping | A user will be assigned one or multiple roles after successful authentication. Role mapping is to map a role to a user (or a backend role). For example, the mapping from <b>kibana_user</b> (role) to <b>Bob</b> (user) means that Bob obtains all permissions of <b>kibana_user</b> after authentication. Similarly, the mapping from <b>all_access</b> (role) to <b>admin</b> (backend role) means that any user with the backend role <b>admin</b> (from the LDAP/Active Directory server) has all the permissions of role <b>all_access</b> after being authenticated. You can map each role to multiple users or backend roles. |
| Action group | An action group is a group of permissions. For example, the predefined <b>SEARCH</b> action group grants roles permissions to use <b>_search</b> and <b>_msearch</b> APIs.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

In addition to the RBAC model, Elasticsearch also uses a concept called tenant. The RBAC model addresses the problem of user-level authorization, while the tenant model addresses the problem of data and resource sharing between different tenants. Within a tenant space, tenants can share information such as dashboards and index patterns.

By default, users can only see the index patterns and dashboards in their own private tenant space. When a new user **test** is added, the system automatically

generates an index named **.kibana\_xxx\_test**. The data in this user's private space will be stored in this index. Similarly, the data of the administrator's private tenant space is stored in the **.kibana\_xxx\_admin** index. To share an index pattern or dashboard with other tenants, you can create them in the global tenant space. Other users can access the shared resource only by switching to the global tenant space.

On the Kibana console, you can configure user permissions on an Elasticsearch cluster under **Security** to implement fine-grained access control at four levels: cluster, index, document, and field.

Users can be added or deleted for a cluster, and mapped to roles. This way, you assign roles to users.

With role mapping, you can configure the members of each role and assign roles to users based on usernames, backend roles, and host names. For each role, you can configure cluster, index, and document permissions, as well as the permission to use Kibana.

For more about security configuration for a security-mode cluster and the detailed guide, see the official Elasticsearch document [here](#).


## Constraints

- You can customize the username, role name, and tenant name in Kibana.
- The Kibana GUI varies depending on the Kibana version. Kibana 7.10.2 is used as an example here.

## Creating a User and Granting Permissions


1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to Kibana using the administrator account.
  - Username: **admin** (default administrator account name)
  - Password: Enter the administrator password you set when creating the cluster in security mode.
4. After a successful login, choose **Security** in the navigation tree on the left of the Kibana operation page. The **Security** page is displayed.
5. Add a new user **test** for the security-mode cluster.
  - a. Click **Internal users** in the navigation area.
  - b. On the **Internal users** page, click **Create internal user**. The page for creating a new user is displayed.

**Figure 2-50** Create internal user



- c. On the displayed page, set **Username**, **Password**, and **Re-enter password**. The username **test** is used as an example here.  
The following two parameters are optional. You can click **Learn more** on the page to learn more about them.
    - **Backend roles**: used to map external users (such as those from LDAP or SAML) to Open Distro security roles.
    - **Attributes**: used to further describe users. More importantly, they can be used to enable document-level access control on top of the index permissions of a role. This makes it possible to conduct dynamic document-level security (DLS) queries based on user attributes.
  - d. Click **Create**. Upon successful creation, the new user is displayed in the user list.
6. Create a role named **role\_test** and assign permissions to it.
- a. Click **Roles** under **Security**. The system has preset roles. For the permissions of each role, click **Learn more** on the page. If the preset roles can already meet your needs, you are advised to use these preset roles.
  - b. On the **Roles** page, click **Create Role**.
  - c. Set the role name, for example, **role\_test**.

**Figure 2-51** Setting the role name




- d. On the **Cluster Permissions** page, set cluster permissions based on service requirements. If they are not configured for a role, the role will not have any cluster-level permissions. The following uses the **cluster\_monitor** permission as an example.

**NOTE**

In Elasticsearch, the **cluster\_monitor** permission allows users to monitor and observe cluster status, but not to perform any operations that may alter the cluster status. Specifically, the **cluster\_monitor** permission enables users to perform the following operations:

- Check a cluster's status and health.
- Check the nodes of a cluster.
- View cluster statistics.
- Check the pending tasks of a cluster.
- Check information about cluster recovery, segments, and indexes.

**Figure 2-52 Cluster Permissions**


- e. Configure index permissions on the **Index Permissions** page. This configuration is optional. It allows you to define the permissions of users assigned this role on specific indexes.

- **Index:** Set the index name. For example, **my\_store**.

**NOTE**

Use different names for the index and the user.

- **Index permissions:** Set the index permissions to grant.
- f. **Tenant Permissions:** Set tenant permissions. This configuration is optional. Tenants in Kibana are spaces for saving index patterns, visualizations, dashboards, and other Kibana objects. By default, all Kibana users have access to two tenants: Private and Global. The global tenant is shared between every Kibana user. The private tenant is exclusive to each user and cannot be shared. For more on tenant permissions, click **Learn more** on the page.
  - g. Click **Create** to save the role settings. The new role is displayed in the **Roles** list.
7. Map a role to a user to assign permissions to that user.
    - a. Choose **Security > Roles**, and click **role\_test**. The role details page is displayed.
    - b. Click the **Mapped users** tab, then click **Map user**.
    - c. On the **Map user** page, select user **test** created earlier from the **Users** list.
    - d. Click **Map**.



8. Verify that the user permissions have taken effect.
  - a. Log in to Kibana as user **test**.
  - b. Click **Dev Tools** in the navigation tree on the left.
  - c. Run the **GET /\_cluster/health?pretty** command to check the cluster health. The code 200 is returned. Basic information about the cluster can

be queried, indicating that the user has the permission to check cluster status.

- d. Run the **PUT /my\_test** command to create an index. The code 403 is returned, indicating that the user is not authorized to create indexes.

We can see that user **test** only has the permission check cluster status but cannot create indexes. The configuration is successful.

If necessary, you can add the index creation permission for the role later. The returned error message provides tips on adding role permissions.

### 2.13.3 Setting Tags for an Elasticsearch Cluster

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

Tags can be added during or after cluster creation. If you want to use the same tag to identify multiple cloud resources for better resource grouping, we recommend that you predefine tags in Tag Management Service (TMS). For details, see [Predefined Tags Overview](#).

#### Constraints

- Each cluster can have a maximum of 20 tags.
- If your organization has configured tag policies for CSS, add tags to clusters based on these policies. If a tag does not comply with existing tag policies, cluster creation may fail. Contact the administrator to learn more about these policies.

#### Managing Tags

In CSS, you can add, delete, modify, and query cluster tags.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Select the **Tags** tab, where 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 **Edit Tag** in the upper left corner. In the displayed dialog box, click **Add Tag**, enter the tag key and value, and click **OK**.

- Modify

Click **Edit Tag** in the upper left corner. In the displayed dialog box, modify the key and value of a tag, and click **OK**.

- Delete

Click **Edit Tag** in the upper left corner. In the displayed dialog box, click **Delete** in the row that contains a tag, and click **OK**.



## Searching for Clusters by Tag

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the search box, and select a tag key and value to search for the target cluster.

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.

## 2.13.4 Configuring Default Parameters in the .yaml Configuration File of an Elasticsearch Cluster

You can modify an Elasticsearch cluster's configuration file for purposes like accelerating queries, modifying cross-domain access configuration, adjusting the internal cache size, and managing the task queue size. The core configuration information of an Elasticsearch cluster is stored in a file named **elasticsearch.yaml**. You can modify specific parameters in this file on the CSS management console.

### Configuring the .yaml File

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Cluster Settings > Parameter Settings**.
5. Click **Edit** to modify module parameters as required.

**Table 2-151** Module parameters

| Module Name         | Parameter                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cross-domain Access | http.cors.allow-credentials | <p>Whether to carry authentication information in cross-domain requests, that is, whether to contain the Access-Control-Allow-Credentials field in the response header.</p> <p>The value can be:</p> <ul style="list-style-type: none"><li>• <b>true</b>: The response header contains the Access-Control-Allow-Credentials field.</li><li>• <b>false</b> (default value): The response header does not contain the Access-Control-Allow-Credentials field.</li></ul> |

| Module Name | Parameter                | Description                                                                                                                                                                                                                                                                                                                                          |
|-------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|             | http.cors.allow-origin   | Source IP addresses or domain names allowed for cross-domain access. When there are multiple values, separate them with commas (,).<br>Example: <b>192.168.122.122:9200</b> or <b>192.168.1.1:9200, 192.168.1.2:9200</b>                                                                                                                             |
|             | http.cors.allow-headers  | Request header fields for cross-domain access. When there are multiple values, separate them using a comma (,) and a space.<br>Example: <b>X-Requested-With, Content-Type, Content-Length</b><br>The value can contain a maximum of 1024 characters. It can only contain letters, digits, hyphens (-), underscores (_), colons (:), and slashes (/). |
|             | http.cors.max-age        | Maximum retention duration of pre-check responses in the browser cache.<br>Unit: s<br>Value range: 0 to 31,536,000 (0 to 1 year)<br>Default value: 1,728,000 (20 days)                                                                                                                                                                               |
|             | http.cors.enabled        | Whether to enable cross-domain access.<br>The value can be: <ul style="list-style-type: none"><li>• <b>true</b>: Enable cross-domain access.</li><li>• <b>false</b> (default): Disable cross-domain access.</li></ul>                                                                                                                                |
|             | http.cors.allow-methods  | HTTP methods allowed for cross-domain access. When there are multiple values, separate them using a comma (,) and a space.<br>Example: <b>OPTIONS, GET, POST</b><br>Value range: OPTIONS, HEAD, GET, POST, PUT, and DELETE                                                                                                                           |
| Reindexing  | reindex.remote.whitelist | Which remote hosts are allowed for remote reindexing operations. Use commas (,) to separate multiple values.<br>Example: <b>192.168.122.122:9200</b> or <b>192.168.1.1:9200, 192.168.1.2:9200</b>                                                                                                                                                    |

| Module Name                 | Parameter                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Custom Cache                | indices.queries.cache.size   | Maximum heap space allocated to the query cache.<br>Unit: percentage (%)<br>Value range: 1 to 100<br>Default value: 10%                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Queue Size in a Thread Pool | thread_pool.force_merge.size | Size of the thread pool used for force merge operations.<br>Value range: a positive integer<br>Default value: 1                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Custom                      | Custom parameters            | You can add any custom parameters supported by Elasticsearch.<br>The parameter values must meet the following requirements: <ul style="list-style-type: none"><li>• If a custom parameter has multiple values, use a comma (,) and a space to separate them. The input format is [value1, value1, value1...].</li><li>• Colons (:) are not allowed.</li></ul> <b>WARNING</b><br>Set the values of any custom parameters to those supported by Elasticsearch. Otherwise, the cluster may fail to restart. Exercise caution when performing this operation. |

6. After the change is complete, click **Submit**. In the displayed dialog box, confirm the settings, 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.
7. 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 not updated** will be displayed in the **Task Status** column in the cluster list.
  - If you restart the cluster after the modification, and **Task Status** displays **Configuration error** in the cluster list, the parameter configuration file has failed to be modified.

## 2.13.5 Associating an Elasticsearch Cluster with an Enterprise Project

In CSS, you can associate your Elasticsearch cluster with 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.

You can associate your cluster with an enterprise project either when creating the cluster or afterward.

Prerequisites

To use the enterprise project function, you need to assign permissions to the corresponding account. You can [submit a service ticket](#) to apply for the permissions.

You have [created an enterprise project](#).

Managing Enterprise Projects


- 1. Log in to the [CSS management console](#).
- 2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
- 3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
- 4. On the **Overview** tab, find **Enterprise Project** in the **Configuration** area, and click the enterprise project name next to it. The **Project Management** page is displayed.

Figure 2-53 Enterprise Project




- 5. On the **Resources** tab, select the region of the current cluster, and select **CloudSearchService** as **Service**. Then click **Search**. The CSS clusters are displayed in the resource list below.

Figure 2-54 Filtering CSS clusters



- 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**.

Figure 2-55 Removing resources



8. After the change is complete, return to the Elasticsearch cluster list on the CSS console and check the cluster's new enterprise project.

Figure 2-56 Enterprise project

| Name/ID                | Clust...<br>▼ | Task ...<br>▼ | Version<br>▼ | Private IP... | Secur...<br>▼ | Enterprise Project<br>▼ |
|------------------------|---------------|---------------|--------------|---------------|---------------|-------------------------|
| CSS-...<br>60a7b7a3... | NEW           | ○ Availa...   | --           | 10...<br>1... | Disabled      | d...                    |

### 2.13.6 Changing the Node Specifications and Storage Type of an Elasticsearch Cluster

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

**Table 2-152** Change scenarios


| Change Type                  | Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Change Process                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Changing node specifications | <p>Typically, you increase node specifications instead of decreasing them. Common scenarios include:</p> <ul style="list-style-type: none"><li>• If the allocation of new indexes or shards takes too long or the coordination and scheduling of nodes are inefficient, increase the master node specifications.</li><li>• If too many requests need to be handled or too many results need to be aggregated, increase the client node specifications.</li><li>• If data nodes are becoming slower in responding to data writing and query requests, increase the specifications of the data nodes.</li><li>• If cold data query becomes slow, increase the cold data node specifications.</li><li>• If the CPU or JVM metrics of cluster nodes show performance bottlenecks, increase node specifications to improve cluster performance.</li></ul> <p>Alternatively, you may also decrease node specifications, but doing so will decrease the cluster's data processing and storage capacities. Exercise caution.</p> | <ol style="list-style-type: none"><li>1. Bring a node offline.</li><li>2. Change the node specifications.</li><li>3. Restart the node and restore its data.</li><li>4. After the node recovers, proceed to another node and repeat the steps above. This goes on until all nodes are changed.</li></ol> <p>The node specifications are changed one node at a time. This is to ensure that there are sufficient resources to keep services running.</p> |

| Change Type                                | Scenario                                                                                                                 | Change Process                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Changing the node storage type (disk type) | Change the node storage type if disk I/O has become a performance bottleneck, which impacts query and write performance. | <ol style="list-style-type: none"><li>1. Select a random node and migrate its data to other nodes.</li><li>2. Rebuild the node using the new storage type.</li><li>3. Add the node back into the cluster. The system automatically triggers a shard reallocation, migrating some of the shards to the new node.</li><li>4. After the node recovers, proceed to another node and repeat the steps above.</li></ol> <p>The nodes are changed one at a time to prevent service interruptions.</p> |

## Prerequisites

- The cluster status is **Available**, and there are no ongoing tasks.
- Your CSS resource quotas are sufficient for the capacity expansion you are about to perform. You can check available resources on the **Modify Configuration** page.

**Figure 2-57** Checking available resources



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

## Impact on Billing

For a pay-per-use cluster, you can see its new price when confirming the node specifications or storage type change on the console. After the change is complete, the new price will apply. For pricing details, see [Cloud Search Service Price Calculator](#).

For a yearly/monthly cluster, you can also see the new price when you confirm the change. In the case of an upgrade, a new order is triggered when you confirm the

change; and in the case of a downgrade, a refund request is triggered when you do so.

## Constraints

- The node specifications and storage type cannot be changed for nodes that use local disks.
- The node specifications and storage type cannot be changed at the same time.
- The node storage type can be changed only for data nodes and cold data nodes.
- When you change the node storage type, data needs to be migrated between different nodes. The timeout for data migration per node is 48 hours. Upgrade will fail if this timeout expires. When the cluster has large quantities of data, you are advised to manually adjust the data migration rate and avoid performing the migration during peak hours.
- For a cluster without master nodes, the node storage type can be changed only if the number of data nodes plus cold data nodes is at least three.
- For a cluster with master nodes, this operation is allowed only if the cluster has at least two data nodes.
- During a node storage type change, there is always one node that is unavailable. To ensure service continuity, make sure the **total number of data nodes and cold data nodes is greater than the maximum number of index replicas plus 1**. For a single-AZ or dual-AZ cluster, also make sure there are at least two nodes of each type for the cluster in each of its AZs.
- During a node specifications change, nodes are brought offline in order to make the changes. To ensure service continuity, make sure all shards have replicas.
- Make sure the disk usage is always less than 80% during the change.

## Change Impact

Before changing a cluster's node specifications or storage type, it is essential to assess the potential impacts and review operational recommendations. This enables proper scheduling of the change, minimizing service interruptions.

- Performance impact (only when the node storage type is changed)

Changing the node storage type does not interrupt services. However, data migration that occurs during this process consumes I/O performance, and taking individual nodes offline still has some impact on the overall cluster performance.

To minimize this impact, it is advisable to adjust the data migration rate based on the cluster's traffic cycle: increase the data migration rate during off-peak hours to shorten the task duration, and decrease it **before** peak hours arrive to ensure optimal cluster performance. The data migration rate is determined by the **indices.recovery.max\_bytes\_per\_sec** parameter. The default value of this parameter is the number of vCPUs multiplied by 32 MB. For example, for four vCPUs, the data migration rate is 128 MB. Set this parameter to a value between 40 MB and 1000 MB based on your service requirements.



```
PUT /_cluster/settings
{
 "transient": {
 "indices.recovery.max_bytes_per_sec": "1000MB"
 }
}
```

- Impact on request handling  
Taking nodes offline one at a time usually does not interrupt services. However, requests sent to offline nodes may fail. To mitigate this impact, the following measures may be taken:
  - Use a VPC endpoint or a dedicated load balancer to handle access requests to your cluster, which makes sure that requests are automatically routed to available nodes.
  - Enable an exponential backoff & retry mechanism on the client (configure three retries).
  - Perform this operation during off-peak hours.
- Impact on index replicas  
Shards that have no replicas will become unavailable when the nodes that store them are taken offline, causing service interruptions. You are advised to add replicas for all important indexes before making the change described in this topic.
- Impact on Kibana and Cerebro  
Changing the node storage type for a cluster will cause Kibana and Cerebro to be rebuilt. During this period, Kibana and Cerebro are temporarily unavailable. During a node specifications change, if Kibana and Cerebro become unavailable because the node that runs them is taken offline, refresh the web page or try to log in again, and the system will reschedule Kibana and Cerebro to an available node.
- Characteristics of this process  
Once started, a change task cannot be stopped until it succeeds or fails. A change task failure only impacts a single node, and does not interrupt services if there are data replicas, but the failed node still needs to be restored promptly.

If a change of the node specifications or storage type is urgently needed, submit a [service ticket](#) to contact technical support and perform a pre-change evaluation and the necessary checks.

## Change Durations

- The following formulas can be used to estimate how long the node specifications change will take:  
**Change duration (min) = 10 (min) x Total number of nodes to change + Data recovery duration (min)**  
where,
  - 10 minutes indicates how long non-data recovery operations (e.g., initialization) typically take per node. It is an empirical value.
  - The total number of nodes is the sum of the number of data nodes, master nodes, client nodes, and cold data nodes in the cluster.**Data recovery duration (min) = Total data size (MB)/[Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s)]**

where,

- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.
- The formulas above use estimates under ideal conditions. The actual data recovery speed depends on cluster load.
- The following formulas can be used to estimate how long the node storage type change will take:

**Change duration (min) = 15 (min) x Total number of nodes to change + Data migration duration (min)**

where,

- 15 minutes indicates how long non-data migration operations (e.g., initialization) typically take per node. It is an empirical value.
- The total number of nodes is the sum of the number of data nodes, master nodes, client nodes, and cold data nodes in the cluster.

**Data migration duration (min) = Total data size (MB)/[Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s)]**

where,


- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.
- The formulas above use estimates under ideal conditions. The actual migration speed depends on cluster load.

## Changing the Node Specifications and Storage Type

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. Check that all service data has replicas so that services will not be interrupted during the change.
  - a. In the displayed cluster list, find the target cluster, and click **Access Kibana** in the **Operation** column to log in to the Kibana console.
  - b. In the left navigation pane, choose **Dev Tools**.
  - c. Run the **GET \_cat/indices?v** command in Kibana.
    - If the returned **rep** value is greater than **0**, data replicas exist. Go to the next step.
    - If the returned **rep** value is **0**, there are no data replicas. You are advised to manually create a snapshot for the cluster before moving on to the next step. For details, see [Creating Snapshots to Back Up the Data of an Elasticsearch Cluster](#).
4. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
5. Click the **Scale Cluster** tab.
6. Click **Change specifications** to set parameters.

**Table 2-153** Changing the node specifications and storage type

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                               |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Action    | Select <b>Change specifications</b> .                                                                                                                                                                                                                                                                                                                                     |
| Resources | Shows the change of resources for this operation.                                                                                                                                                                                                                                                                                                                         |
| Nodes     | Configure the changes you want to make.<br><ol style="list-style-type: none"><li>1. Select a node type in the <b>Node Type</b> column.</li><li>2. Select a new flavor in the <b>Node Specifications</b> column, or select a new storage type in the <b>Node Storage</b> column.</li></ol><br>The node specifications and storage type cannot be changed at the same time. |

**Figure 2-58** Changing specifications

7. Click **Next**.
8. Confirm the information and click **Submit**.
9. In the displayed dialog box, confirm check items, and click **OK** to start the specifications change.
  - Check items for node specifications change: **Verify index copies** and **Cluster status check**.
  - Check items for node storage type change: **Check cluster load**.

**Table 2-154** Check item description

| Item                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Verify index copies | <p>By default, CSS checks for indexes that do not have any replicas created for them. You can skip this step, but the lack of index replicas may impact service availability during a node specifications change.</p> <ul style="list-style-type: none"><li>• If you select <b>Verify index copies</b> and the cluster has no master nodes, each index is required to have at least one replica and the cluster must have at least three nodes.</li><li>• If you select <b>Verify index copies</b> and the cluster has a master node, each index is required to have at least one replica, but there is no requirement on how many nodes the cluster must have.</li></ul> |

| Item                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster status check | By default, the cluster status is checked before specifications change. The specifications of nodes are changed one by one to ensure service availability and data security. If a cluster is overloaded and services are faulty, the request for a specifications change may fail to be delivered. In this case, you can skip cluster status check, but doing so may cause the cluster to become faulty and services become interrupted. Exercise caution when performing this operation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Check cluster load   | <p>During a node storage type change, data migration between nodes and the stopping and restarting of nodes will consume cluster resources, causing the cluster load to increase. A cluster load check can identify possible overload risks for a cluster and reduce the likelihood of an overload condition causing the node storage type change to fail.</p> <p>The cluster load check items are as follows:</p> <ul style="list-style-type: none"><li>• <b>nodes.thread_pool.search.queue &lt; 1000</b>: Check that the maximum number of requests in the search queue is less than 1000.</li><li>• <b>nodes.thread_pool.write.queue &lt; 200</b>: Check that the maximum number of requests in the write queue is less than 200.</li><li>• <b>nodes.process.cpu.percent &lt; 90</b>: Check that the maximum CPU usage is less than 90%.</li><li>• <b>nodes.os.cpu.load_average/Number of vCPUs &lt; 80%</b>: Check that the number of running processes plus the number of processes waiting for CPUs is less than 80% of the total number of vCPUs.</li></ul> |

 **NOTE**

If the change request fails to be submitted and a message is displayed indicating that the cluster needs to be upgraded, it means the current cluster version does not support a node storage type change. Upgrade the cluster to the latest image version and then try again. For a detailed upgrade guide, see [Upgrading the Version of an Elasticsearch Cluster](#).

10. Click **Back to Cluster List** to go back to the **Clusters** page. The **Cluster Status** is **Configuration modified**. When **Cluster Status** changes to **Available**, the cluster node specifications have been successfully changed.

## 2.13.7 Replacing a Specified Node for an Elasticsearch Cluster

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

The node replacement process is as follows:

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

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

## Prerequisites

- The cluster status is **Available**, and there are no ongoing tasks.
- All mission-critical data has been backed up.

## Constraints

- Only one node can be replaced at a time. Each new node is rebuilt using the ID, IP address, specifications, and AZ of the node it is replacing.
- The configurations you modified manually will not be retained after node replacement. For example, if you have manually added a return route for the original node, you need to add it again for the new node after the node replacement is complete.
- If the node you want to replace is a data node or cold data node, pay attention to the following precautions:
  - When a data node or cold data node is replaced, its data is first migrated to other data nodes. This means the **total number of data nodes and cold data nodes must be greater than the maximum number of index replicas plus 1**.
  - Elasticsearch clusters whose version is earlier than 7.6.2 cannot have closed indexes. Otherwise, their data nodes or cold data nodes cannot be replaced.
  - In the AZ that contains the data node or cold data node to be replaced, there has to be at least another data node or cold data node.
  - If the cluster has not master nodes, the total number of data nodes plus cold data nodes must be at least three.
  - The precautions above do not apply if you are replacing a faulty node, regardless of its type. This is because faulty nodes are not included in `_cat/nodes`.

## Change Impact

Before replacing a node, it is essential to assess the potential impacts and review operational recommendations. This enables proper scheduling of the node replacement, minimizing service interruptions.

- Impact on performance  
Replacing a node does not interrupt services. However, data migration that occurs during this process consumes I/O performance, and taking individual nodes offline still has some impact on the overall cluster performance.

To minimize this impact, it is advisable to adjust the data migration rate based on the cluster's traffic cycle: increase the data migration rate during off-peak hours to shorten the task duration, and decrease it **before** peak hours arrive to ensure optimal cluster performance. The data migration rate is determined by the **indices.recovery.max\_bytes\_per\_sec** parameter. The default value of this parameter is the number of vCPUs multiplied by 32 MB. For example, for four vCPUs, the data migration rate is 128 MB. Set this parameter to a value between 40 MB and 1000 MB based on your service requirements.

```
PUT /_cluster/settings
```

```
{
 "transient": {
 "indices.recovery.max_bytes_per_sec": "1000MB"
 }
}
```

- Impact on request handling

While a node is being replaced, requests sent to it may fail. To mitigate this impact, the following measures may be taken:

- Use a VPC endpoint or a dedicated load balancer to handle access requests to your cluster, which makes sure that requests are automatically routed to available nodes.
- Enable an exponential backoff & retry mechanism on the client (configure three retries).
- Perform this operation during off-peak hours.

- Characteristics of this process

Once started, a node replacement task cannot be stopped until it succeeds or fails. A task failure only impacts a single node, and does not interrupt services if there are data replicas, but the failed node still needs to be restored promptly.

## Node Replacement Duration

The following formula can be used to estimate how long it will take to replace a specified node of a cluster:

**Change duration (min) = 15 (min) + Data migration duration (min)**

where, 15 minutes indicates how long non-data migration operations (e.g., initialization) typically take per node. It is an empirical value.

**Data migration duration (min) = Total data size (MB)/[Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s)]**

where,

- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.
- The formulas above use estimates under ideal conditions. The actual migration speed depends on cluster load.

## Replacing a Specified Node

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.

3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. On the **Modify Configuration** page, click the **Replace Node** tab.
5. On the **Replace Node** tab, set the parameters as needed.

**Table 2-155** Replacing a specified node

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>After a node is replaced, NICs need to be reattached to the new node. This means you need to have VPC permissions. Select an IAM agency to grant the current account the permission to access and use VPC.</p> <ul style="list-style-type: none"><li>• This parameter is available only when the new IAM plane is connected.</li><li>• If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-upgrade-agency</b>.</li><li>• If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>VPC Administrator</b> role or the <b>VPC FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/>"vpc:subnets:get",<br/>"vpc:ports:*"</li><li>• To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</li><li>• To use an IAM agency, the following minimum permissions are required:<br/>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</li></ul> |
| Node Type | Select the node you want to replace. You can expand a node type to check all the nodes under it.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

6. Click **Submit**. In the data migration confirmation dialog box, choose to migrate data, and click **OK**.

During data migration, the system migrates all data from the to-be-replaced node to the remaining nodes, and replace the node upon completion of the data migration. If the data on the to-be-replaced nodes has replicas on other nodes, data migration can be skipped and the cluster change can be completed faster.

- Click **Back to Cluster List** to go back to the **Clusters** page. The **Task Status** is **Replacing nodes**. When **Cluster Status** changes to **Available**, the node has been successfully replaced.

## 2.13.8 Changing the Security Mode of an Elasticsearch Cluster

With CSS, when an Elasticsearch cluster's security needs change, you can change its security mode settings.

Configure the security mode based on the security needs of your cluster.

**Table 2-156** Cluster security modes

| Cluster Type              | Key Settings                   | Description                                                                                                     | Applicable Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|--------------------------------|-----------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Non-security mode cluster | <b>Security Mode:</b> disabled | Access to such a cluster requires no user authentication, and data will be transmitted in plaintext using HTTP. | <p>Use when creating a cluster for internal testing or workloads that have a low security standard.</p> <ul style="list-style-type: none"><li>• Advantage: easy to access the cluster.</li><li>• Disadvantage: poor security, as anyone can access the cluster. When the security mode is disabled, public network access and Kibana public network access cannot be enabled.</li></ul> <p>Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.</p> |



| Cluster Type          | Key Settings                                                                                                                                          | Description                                                                                                                                                        | Applicable Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Security-mode cluster | Security-mode cluster + HTTP: <ul style="list-style-type: none"><li>• <b>Security Mode:</b> enabled</li><li>• <b>HTTPS Access:</b> disabled</li></ul> | Such a cluster requires user authentication. It supports access control and data encryption, and it uses HTTP to transmit data in plaintext.                       | Use to balance security and performance. <ul style="list-style-type: none"><li>• Advantage: User authentication improves cluster security. HTTP-based access ensures high performance of the cluster. User permissions can be configured to ensure proper isolation.</li><li>• Disadvantage: Public network access is not supported.</li></ul> Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.                                                  |
|                       | Security-mode cluster + HTTPS: <ul style="list-style-type: none"><li>• <b>Security Mode:</b> enabled</li><li>• <b>HTTPS Access:</b> enabled</li></ul> | Such a cluster requires user authentication. It supports access control and data encryption, and it uses HTTPS to encrypt communication and enhance data security. | Use when security takes precedence over performance and public network access is required. <ul style="list-style-type: none"><li>• Advantage: User authentication improves cluster security. HTTPS enhances cluster security by encrypting all communication over the public network. User permissions can be configured to ensure proper isolation.</li><li>• Disadvantage: When HTTPS is used, data encryption and decryption introduce computational overhead and impact the cluster's read and write performance.</li></ul> |

**Table 2-157** lists the various types of security mode changes supported for CSS clusters.

**Table 2-157** Security mode change scenarios

| Action                                                | Scenario                                                                                                                                                                                                                          | Change Process                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switching from the non-security mode to security mode | Non-security mode → Security mode + HTTP: Change a cluster from non-security mode to security mode + HTTP.                                                                                                                        | <ol style="list-style-type: none"><li>1. Select a node and modify the Elasticsearch, Kibana, and Cerebro configuration files.</li><li>2. Restart the Elasticsearch, Kibana, and Cerebro processes and restore data.</li><li>3. After the node recovers, proceed to another node and repeat the steps above. This goes on until all nodes are changed.</li></ol> |
|                                                       | Non-security mode → Security mode + HTTPS: Change a cluster from non-security mode to security mode + HTTPS.                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                 |
| Switching from the security mode to non-security mode | Security mode → Non-security mode: <ul style="list-style-type: none"><li>• Change a cluster from security mode + HTTP to non-security mode.</li><li>• Change a cluster from security mode + HTTPS to non-security mode.</li></ul> |                                                                                                                                                                                                                                                                                                                                                                 |
| Switching between HTTP and HTTPS in security mode     | HTTP → HTTPS: Change a cluster from security mode + HTTP to security mode + HTTPS.                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                 |
|                                                       | HTTPS → HTTP: Change a cluster from security mode + HTTPS to security mode + HTTP.                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                 |

## Prerequisites

- Before changing the security mode for a cluster, you are advised to back up its data.
- The cluster status is **Available**, and there are no ongoing tasks.
- Check whether load balancing is enabled for the cluster. If yes, disable load balancing first. Enable load balancing again after the security mode is changed. This prevents errors in accessing the cluster through the load balancer during the change.

## Constraints

The security mode settings can be changed only for Elasticsearch clusters created after November 2022 and whose version is 6.5.4 or later.

## Change Impact

Before changing the security mode for a cluster, it is essential to assess the potential impacts and review operational recommendations. This enables proper scheduling of the change, minimizing service interruptions.

**Table 2-158** Impact overview

| Action                                    | Service Interruption | Authentication Mode         | Performance | Public Network Access  | Security Account    |
|-------------------------------------------|----------------------|-----------------------------|-------------|------------------------|---------------------|
| Non-security mode → Security mode + HTTP  | Yes                  | Authentication required     | N/A         | Disallowed             | N/A                 |
| Non-security mode → Security mode + HTTPS | Yes                  | Authentication required     | Degraded    | Allowed                | N/A                 |
| Security mode → Non-security mode         | Yes                  | Authentication not required | Enhanced    | Automatically disabled | Permanently deleted |
| HTTP → HTTPS                              | Yes                  | No change                   | Degraded    | Allowed                | N/A                 |
| HTTPS → HTTP                              | Yes                  | No change                   | Enhanced    | Automatically disabled | N/A                 |

Impact description:

- The security mode can be enabled or disabled for a cluster, and a security-mode cluster can use either HTTPS and HTTP.
- Service interruption: When a cluster's security mode is changed, the cluster automatically restarts, leading to temporary service unavailability.
- Authentication mode: Changing a cluster's security mode also changes its user authentication mechanism. The client authentication logic must be updated accordingly. If client adaptation cannot be completed in a timely manner, services may be interrupted.
- Performance impact: Given the same hardware configuration, switching from HTTP to HTTPS leads to a performance loss (such as throughput) of around 20% under high concurrency. Conversely, when a cluster switches from HTTPS to HTTP, cluster performance improves.

- Public network access: Only security-mode clusters that use HTTPS support public network access.
- Security account: When the security mode is disabled for a cluster, the system permanently deletes the cluster's security account.
- Tool impact: If the security mode is changed for a cluster that has an ongoing Kibana session, a session error message will be displayed on Kibana. To rectify the error, clear the cache and launch Kibana again.

Changing the security mode for a cluster changes its accessibility mode, possibly causing service interruptions. You should perform this operation before services are brought online or when service interruptions can be tolerated.

## Change Duration

The following formula can be used to estimate how long it will take to change the security mode for a cluster:

**Change duration (min) = 5 (min) x Total number of nodes to change + Data recovery duration (min)**

where,

- 5 minutes indicates how long non-data recovery operations (e.g., initialization) typically take per node. It is an empirical value.
- The total number of nodes is the sum of the number of data nodes, master nodes, client nodes, and cold data nodes in the cluster.

**Data recovery duration (min) = Total data size (MB)/[Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s)]**

where,

- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.
- The formulas above use estimates under ideal conditions. The actual data recovery speed depends on cluster load.

## Switching from the Non-Security Mode to Security Mode

You can change a non-security mode cluster to a security-mode cluster that uses HTTP or HTTPS. After a cluster's security mode is enabled, user 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**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. Choose the **Change Security Mode** tab.
5. Enable the security mode. Enter and confirm the administrator password of the cluster.

**Figure 2-59** Enabling the security mode

6. Enable or disable **HTTPS Access**.
  - If you enable **HTTPS Access**, HTTPS is used to encrypt cluster communication, and public network access can be enabled for the cluster.
  - If you disable **HTTPS Access**, HTTP is used, and public network access is disallowed for the cluster.
7. Click **Submit**, and confirm the information. The cluster list is displayed.  
The **Task Status** of the cluster is **Changing security mode**. When the cluster status changes to **Available**, the security mode has been successfully changed.

## Switching from Security Mode to Non-Security Mode

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

---


### WARNING

- Clusters with the security mode disabled do not require user authentication, and HTTP is used to transmit data in plaintext. Make sure such a cluster are deployed in a secure environment, and do not expose the cluster's network interface to the public network.
- When a cluster is switched from the security mode to non-security mode, indexes used by the security-mode cluster will be deleted. Back up data before changing the security mode.
- If a public IP address has been assigned to a security-mode cluster, unassign it before changing the security mode.
- If Kibana public network access is enabled for a security-mode cluster, 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 > Elasticsearch**.
  3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.

4. Choose the **Change Security Mode** tab.
5. Disable the security mode.

**Figure 2-60** Disabling the security mode



6. Click **Submit**. In the displayed dialog box, confirm the information. The cluster list page is displayed.

The **Task Status** of the cluster is **Changing security mode**. When the cluster status changes to **Available**, the security mode has been successfully changed.

## Switching between HTTP and HTTPS in Security Mode

You can change the protocol of a security cluster.




### WARNING

If a public IP address has been assigned to a security-mode cluster, unassign it before changing from HTTPS to HTTP.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. Choose the **Change Security Mode** tab.
5. Enable or disable **HTTPS Access**.

**Figure 2-61** Configuring the protocol



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

## 2.13.9 Switching AZs for an Elasticsearch Cluster

CSS supports cross-AZ deployment. By switching AZs for a cluster, you can make more resources available to the cluster or improve the cluster's availability.

You can **Add AZ** or **Migrate AZ**.

- **Add AZ:** Add one or two AZs for a single-AZ cluster, or add one AZ for a dual-AZ cluster to enhance cluster availability.
- **Migrate AZ:** Migrate data from the current AZ to another AZ that has sufficient resources for the cluster.

The AZ change process is as follows:

1. Migrate data from one node to other available nodes.
2. Rebuild the node for the target AZ. Keep the node ID, IP address, and specifications unchanged.
3. Add the new node into the cluster. The system automatically triggers a shard reallocation, migrating some of the shards to the new node. Then repeat these steps on all the other nodes.

The nodes are changed one at a time to prevent service interruptions.

### Prerequisites

- The cluster status is **Available**, and there are no ongoing tasks.
- Make sure that the cluster has not had any non-standard modifications, such as modifying return routes and system parameters. Such modifications, if ever made, will be lost after the AZ change, which may affect your services.
- All mission-critical data has been backed up before an AZ switchover. This is to prevent data loss.

### Constraints

- To ensure service continuity, the total number of data nodes plus cold data nodes in the cluster must be at least three.
- During the change, nodes are brought offline one by one and then new nodes are created. Ensure that the total disk capacity of the remaining nodes is sufficient for handling all the data of the cluster, and that the nodes' disk usage stays below 80%.
- To make sure that all index shards in the cluster can be allocated to the remaining data nodes and cold data nodes, the **total number of data nodes and cold data nodes must be greater than the maximum number of index replicas plus 1**.
- Before a change completes, some nodes may have already been moved to a new AZ. In this case, the AZs before and after the change are both displayed. After the change succeeds, the new AZs and their nodes will be displayed properly.
- When you add nodes for a cluster (**Add AZ**), the current AZ must be retained. When adding one or two AZs for a single-AZ cluster, you must change AZs for all nodes. When adding one AZ for a dual-AZ cluster, you can choose to change AZs for specific types of nodes or all nodes in the cluster. For example, for a dual-AZ cluster, you can use three AZs for its master nodes, while still

using two AZs for the rest of its nodes. To complete this AZ change, the system will try to move as few nodes as possible to rebuild the cluster. During this process, the YML configuration file of the nodes that are not modified is also updated. You need to restart the cluster to make the change take effect.

- When migrating an AZ (**Migrate AZ**), you can select only one target AZ. For this operation, you can choose to migrate only specific types of nodes or all nodes in the cluster. For example, for a dual-AZ cluster, you may move the master nodes from one AZ to the other, while keeping the rest of the nodes unchanged. If you migrate from one AZ to another, there is no need to restart the cluster. If multiple AZs are migrated, you must restart the cluster to make the change take effect.

## Change Impact

Before changing a cluster's AZs, it is essential to assess the potential impacts and review operational recommendations. This enables proper scheduling of the change, minimizing service interruptions.

- Performance

Changing AZs for a cluster does not interrupt services. However, data migration that occurs during this process consumes I/O performance, and taking individual nodes offline still has some impact on the overall cluster performance.

To minimize this impact, it is advisable to adjust the data migration rate based on the cluster's traffic cycle: increase the data migration rate during off-peak hours to shorten the task duration, and decrease it **before** peak hours arrive to ensure optimal cluster performance. The data migration rate is determined by the **indices.recovery.max\_bytes\_per\_sec** parameter. The default value of this parameter is the number of vCPUs multiplied by 32 MB. For example, for four vCPUs, the data migration rate is 128 MB. Set this parameter to a value between 40 MB and 1000 MB based on your service requirements.

```
PUT /_cluster/settings
{
 "transient": {
 "indices.recovery.max_bytes_per_sec": "1000MB"
 }
}
```

- Impact on request handling

While a node is offline, requests sent to it may fail. To mitigate this impact, the following measures may be taken:

- Use a VPC endpoint or a dedicated load balancer to handle access requests to your cluster, which makes sure that requests are automatically routed to available nodes.
- Enable an exponential backoff & retry mechanism on the client (configure three retries).
- Perform this operation during off-peak hours.

- Characteristics of this process

Once started, an AZ change task cannot be stopped until it succeeds or fails.



## Change Duration

The following formula can be used to estimate how long it will take to change AZs for a cluster:

**Change duration (min) = 15 (min) x Total number of nodes involved + Data migration duration (min)**

where, 15 minutes indicates how long non-data migration operations (e.g., initialization) typically take per node. It is an empirical value.

**Data migration duration (min) = Total data size of the nodes to be migrated (MB) ÷ [Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s)]**

where,

- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.
- The formulas above use estimates under ideal conditions. The actual migration speed depends on cluster load.

## Changing AZs

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. Click the **Change AZ** tab.
5. On the **Change AZ** page, set parameters.

---

### CAUTION

If the number of data nodes plus cold data nodes in a cluster is not divisible by the number of AZs, data in the cluster may be unevenly distributed after the AZ change is completed. For example, if the cluster has two AZs and three data nodes, one node may need to store twice as much data as the other two nodes. This uneven distribution of data may affect the cluster's query or write performance.

---

**Table 2-159** Parameters for changing AZs

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Operation Type | <ul style="list-style-type: none"><li>● <b>Add AZ:</b> Add one or two AZs for a single-AZ cluster, or add one AZ for a dual-AZ cluster.<br/>To complete this AZ change, the system will try to move as few nodes as possible to rebuild the cluster. During this process, the YML configuration file of the nodes that are not modified is also updated. You need to restart the cluster to make the change take effect.</li><li>● <b>Migrate AZ:</b> Migrate data from one AZ to another.<br/>After adding AZs, you need to restart the cluster to make the modification take effect.</li></ul> |
| Node Type      | Select a node type for the AZ change.<br>Select one node type or <b>All nodes</b> to change their AZ. When adding one or two AZs for a single-AZ cluster, you can only select <b>All nodes</b> .                                                                                                                                                                                                                                                                                                                                                                                                 |
| Current AZ     | Current AZ of a cluster                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Target AZ      | AZ to add or migrate to. <ul style="list-style-type: none"><li>● <b>Add AZ:</b> Select up to three AZs, which must include all your current AZs.</li><li>● <b>Migrate AZ:</b> Select only one target AZ, which cannot be your current AZ.</li></ul>                                                                                                                                                                                                                                                                                                                                              |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>After a node is rebuilt, NICs need to be reattached to the new node. This means you need to have VPC permissions. Select an IAM agency to grant the current account the permission to access and use VPC.</p> <ul style="list-style-type: none"><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-upgrade-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>VPC Administrator</b> role or the <b>VPC FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/><pre>"vpc:subnets:get",<br/>"vpc:ports:*"</pre></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</pre></li><li>To use an IAM agency, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</pre></li></ul> |

6. Click **Submit**. In the displayed dialog box, choose whether to check for the presence of a full index snapshot, and click **OK** to start the change.

 **NOTE**

You are advised to select **Check full index snapshot**. This ensures that all data has been backed up, so that in case the task fails, the data can be restored using this snapshot.

When **Task Status** in the task list below changes to **Successful**, the AZ change is completed.

7. Confirm the result.
- In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  - On the **Overview** tab, check the cluster's AZs in the **Configuration** area to determine whether the switchover is successful.

**Figure 2-62** Checking the cluster's AZs

### Configuration

Region

AZ

AZ2

VPC

vpc-

## Checking the Progress of an AZ Change Task

In the **Jobs** area, find the AZ change task.

Expand the task, and click **View Progress** to check its progress.

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

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

If by the time an AZ change task is terminated, the AZ has not been changed for some nodes, you can try to restore these nodes by performing [Replacing a Specified Node for an Elasticsearch Cluster](#).

---

#### CAUTION

If the AZ change is only partially successful, the cluster will have nodes spanning different AZs, and there is no convenient way to continue to change the AZ for the remaining nodes. To avoid this situation, do not terminate the task in this state.

---

## 2.13.10 Configuring Custom Word Dictionaries for an Elasticsearch Cluster

### 2.13.10.1 Introduction to Word Dictionaries for Elasticsearch Clusters

#### Context

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

CSS uses the IK and synonym analyzers. The IK analyzer uses a main word dictionary and a stop word dictionary. The synonym analyzer uses a synonym word dictionary. The IK analyzer uses the **ik\_max\_word** and **ik\_smart** word segmentation policies. The synonym analyzer uses the **ik\_synonym** word segmentation policy.

- **ik\_max\_word**: splits the text at a fine granularity.
- **ik\_smart**: splits the text at a coarse granularity.

CSS provides four preset word dictionaries: static main word dictionary, static stop word dictionary, extra main word dictionary, and extra stop word dictionary.

- If these preset word dictionaries can already support your application needs, there is no need to configure custom word dictionaries.
- If these preset word dictionaries are no longer sufficient, you can add a custom main word dictionary, stop word dictionary, or synonym dictionary, or modify the four preset word dictionaries to enable search by keywords or synonyms.

**Table 2-160** describes the different word dictionaries used by CSS.

**Table 2-160** Introduction to word dictionaries

| Word Dictionary Type |                      | Description                                                                                                                                         | File Requirements                                                                                                                                                                                     |
|----------------------|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Custom               | Main word dictionary | Main words are words that users want to use as basis for performing word segmentation. A main word dictionary is a custom collection of such words. | The word dictionary file must be a text file encoded using UTF-8 without BOM, with one word per row. The maximum size of a main word dictionary file is 100 MB. Letters must be in lowercase.         |
|                      | Stop word dictionary | Stop words are words that users wish to be ignored during word segmentation. The stop word dictionary is a custom collection of stop words.         | The word dictionary file must be a text file encoded using UTF-8 without BOM, with one word per row. The maximum size of a stop word dictionary file is 100 MB.                                       |
|                      | Synonym dictionary   | Synonyms are words with the same or similar meanings. A synonym dictionary is a custom collection of synonyms.                                      | The word dictionary file must be a text file encoded using UTF-8 without BOM. Each row contains a group of synonyms separated by commas (,). The maximum size of a synonym dictionary file is 100 MB. |

| Word Dictionary Type |                             | Description                                                                                                                                                                                                                                                                                                                 | File Requirements                                                                                                                                                      |
|----------------------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Built-in             | Static main word dictionary | The static main word dictionary is a collection of common main words preconfigured in CSS. To view the static main word dictionary, visit <a href="https://github.com/infinilabs/analysis-ik/blob/master/config/main.dic">https://github.com/infinilabs/analysis-ik/blob/master/config/main.dic</a> .                       | The word dictionary file must be a text file encoded using UTF-8 without BOM, with one word per row. The maximum size of a static main word dictionary file is 100 MB. |
|                      | Static stop word dictionary | The static stop word dictionary is a collection of common stop words preconfigured in CSS. To view the static stop word dictionary, visit <a href="https://github.com/infinilabs/analysis-ik/blob/master/config/stopword.dic">https://github.com/infinilabs/analysis-ik/blob/master/config/stopword.dic</a> .               | The word dictionary file must be a text file encoded using UTF-8 without BOM, with one word per row. The maximum size of a static stop word dictionary file is 100 MB. |
|                      | Extra main word dictionary  | The Extra main word dictionary is a collection of uncommon main words preconfigured in CSS. To access the Extra main word dictionary, visit <a href="https://github.com/infinilabs/analysis-ik/blob/master/config/extra_main.dic">https://github.com/infinilabs/analysis-ik/blob/master/config/extra_main.dic</a> .         | The word dictionary file must be a text file encoded using UTF-8 without BOM, with one word per row. The maximum size of an Extra main word dictionary file is 100 MB. |
|                      | Extra stop word dictionary  | The Extra stop word dictionary is a collection of uncommon stop words preconfigured in CSS. To access the Extra stop word dictionary, visit <a href="https://github.com/infinilabs/analysis-ik/blob/master/config/extra_stopword.dic">https://github.com/infinilabs/analysis-ik/blob/master/config/extra_stopword.dic</a> . | The word dictionary file must be a text file encoded using UTF-8 without BOM, with one word per row. The maximum size of an Extra stop word dictionary file is 100 MB. |

## Constraints

- You cannot use the custom word dictionary function for clusters created before the function was launched (March 10, 2018).
- Custom word dictionaries are used for Chinese word segmentation. They can also be used to segment English words based on special characters other than `#&+-.@_`.
- The word dictionary files must meet the requirements described in [Table 2-160](#).

- Make sure the account you are using has the following permissions (choose **Permissions > Policies/Roles** on the IAM console to check your permissions):
  - **css:IKThesaurus\***: the permission to configure custom word dictionaries.
  - Permission to view the OBS bucket list and read permissions on OBS buckets and the objects in them. For details, see [Permissions Management](#).

## 2.13.10.2 Configuring and Using Custom Word Dictionaries for an Elasticsearch Cluster

### Prerequisites

You have uploaded the custom word dictionary files to an OBS bucket. The OBS bucket must be in the same region as the target CSS cluster, and the storage class must be **Standard**.

### Configuring Custom Word Dictionaries

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Cluster Settings > Custom Word Dictionaries**.
5. On the **Word Dictionaries** page, configure custom word dictionaries for the cluster or modify preset ones.
  - a. To configure custom word dictionaries, see [Table 2-161](#).

**Table 2-161** Configuring custom word dictionaries

| Parameter            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBS Bucket           | Select the OBS location for storing the word dictionary file.<br><br>If no OBS buckets meet your requirements, click <b>Create Bucket</b> to go to the OBS console and create one. For details, see <a href="#">Creating a Bucket</a> .<br><br>To grant an IAM user access to an OBS bucket, you need to grant the <b>GetBucketStoragePolicy</b> , <b>GetBucketLocation</b> , <b>ListBucket</b> , and <b>ListAllMyBuckets</b> permissions to that user. |
| Main Word Dictionary | Main Word Dictionary is a custom word dictionary. Its initial state is empty. By default, <b>No Update</b> is selected, meaning not to configure this word dictionary. <ul style="list-style-type: none"><li>• To add a custom main word dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>• If you do not want to use a certain word dictionary, click <b>Do Not Use</b>.</li></ul>                                      |

| Parameter            | Description                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Stop Word Dictionary | <p>Stop Word Dictionary is a custom word dictionary. Its initial state is empty. By default, <b>No Update</b> is selected, meaning not to configure this word dictionary.</p> <ul style="list-style-type: none"><li>• To add a custom stop word dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>• If you do not want to use a certain word dictionary, click <b>Do Not Use</b>.</li></ul> |
| Synonym Dictionary   | <p>Synonym Dictionary is a custom word dictionary. Its initial state is empty. By default, <b>No Update</b> is selected, meaning not to configure this word dictionary.</p> <ul style="list-style-type: none"><li>• To add a custom synonym dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>• If you do not want to use a certain word dictionary, click <b>Do Not Use</b>.</li></ul>     |

- b. To modify a preset word dictionary, toggle on **Modify Preset Word Dictionary**, and then modify that word dictionary as needed.

 **NOTE**

If the four preset word dictionaries (static main word, static stop word, extra main word, and extra stop word) are not displayed, the current cluster version does not support the deletion or modification of them. To use this function, you are advised to upgrade the cluster version, or create a new cluster and migrate data to it.

**Table 2-162** Configuring preset word dictionaries

| Parameter                   | Description                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Static Main Word Dictionary | <p>Static Main Word Dictionary is a preset collection of common main words. By default, <b>No Update</b> is selected, meaning this word dictionary will be used.</p> <ul style="list-style-type: none"><li>• To modify Static Main Word Dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>• If you do not want to use Static Main Word Dictionary, click <b>Do Not Use</b>.</li></ul> |
| Static Stop Word Dictionary | <p>Static Stop Word Dictionary is a preset collection of common stop words. By default, <b>No Update</b> is selected, meaning this word dictionary will be used.</p> <ul style="list-style-type: none"><li>• To modify Static Stop Word Dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>• If you do not want to use Static Stop Word Dictionary, click <b>Do Not Use</b>.</li></ul> |



| Parameter                  | Description                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Extra Main Word Dictionary | Extra Main Word Dictionary is a preset collection of uncommon main words. By default, <b>No Update</b> is selected, meaning this word dictionary will be used. <ul style="list-style-type: none"><li>• To modify Extra Main Word Dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>• If you do not want to use Extra Main Word Dictionary, click <b>Do Not Use</b>.</li></ul> |
| Extra Stop Word Dictionary | Extra Stop Word Dictionary is a preset collection of uncommon stop words. By default, <b>No Update</b> is selected, meaning this word dictionary will be used. <ul style="list-style-type: none"><li>• To modify Extra Stop Word Dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>• If you do not want to use Extra Stop Word Dictionary, click <b>Do Not Use</b>.</li></ul> |

6. Click **Save**. In the dialog box that is displayed, click **OK**. The word dictionary information is displayed in the lower part of the page. The word dictionary status is **Updating**. In approximately 1 minute, the word dictionaries are configured, and **Word Dictionary Status** changes to **Successful**.
7. The deletion or update of the four preset dictionaries (static main word, static stop word, extra main word, and extra stop word) requires a cluster restart to take effect. The update of other word dictionaries happens dynamically, and there is no need for cluster restart. For details about how to restart a cluster, see [Restarting an Elasticsearch Cluster](#).

## Example

Configure a custom word dictionary for the cluster, set main words, stop words, and synonyms. Search for the target text by keyword and synonym and check the search results.

**Step 1** Configure a custom word dictionary and check the word segmentation result. If there is no need for a custom word dictionary because the preset word dictionaries are already sufficient, skip this step.

1. Prepare a word dictionary file (a text file encoded using UTF-8 without BOM) and upload it to the target OBS path.

Set the main word dictionary file, stop word dictionary file, and synonym word dictionary file.

### NOTE

The built-in static stop word dictionary contains common stop words such as **are** and **the**. If the built-in stop word dictionary was never deleted or updated, you do not need to upload such stop words.

2. Configure the word dictionary by referring to [Table 2-161](#).
3. After the word dictionary takes effect, return to the cluster list. Locate the target cluster and click **Kibana** in the **Operation** column to access the cluster.

4. On the Kibana page, click **Dev Tools** in the navigation tree on the left. The operation page is displayed.
5. Run the following commands to check the performance of the `ik_smart` word segmentation policy and `ik_max_word` word segmentation policy.
  - Use the `ik_smart` word segmentation policy to split a piece of target text.

Example code:

```
POST /_analyze
{
 "analyzer": "ik_smart",
 "text": "Text used for word segmentation"
}
```

After the operation is completed, view the word segmentation result.

```
{
 "tokens": [
 {
 "token": "word-1",
 "start_offset": 0,
 "end_offset": 4,
 "type": "CN_WORD",
 "position": 0
 },
 {
 "token": "word-2",
 "start_offset": 5,
 "end_offset": 8,
 "type": "CN_WORD",
 "position": 1
 }
]
}
```

- Use the `ik_max_word` word segmentation policy to split a piece of target text.

Example code:

```
POST /_analyze
{
 "analyzer": "ik_max_word",
 "text": "Text used for word segmentation"
}
```

After the operation is completed, view the word segmentation result.

```
{
 "tokens" : [
 {
 "token": "word-1",
 "start_offset" : 0,
 "end_offset" : 4,
 "type" : "CN_WORD",
 "position" : 0
 },
 {
 "token": "word-3",
 "start_offset" : 0,
 "end_offset" : 2,
 "type" : "CN_WORD",
 "position" : 1
 },
 {
 "token": "word-4",
 "start_offset" : 0,
 "end_offset" : 1,
 "type" : "CN_WORD",
 "position" : 2
 },
],
}
```

```
{
 "token" : "word-5",
 "start_offset" : 1,
 "end_offset" : 3,
 "type" : "CN_WORD",
 "position" : 3
},
{
 "token" : "word-6",
 "start_offset" : 2,
 "end_offset" : 4,
 "type" : "CN_WORD",
 "position" : 4
},
{
 "token" : "word-7",
 "start_offset" : 3,
 "end_offset" : 4,
 "type" : "CN_WORD",
 "position" : 5
},
{
 "token" : "word-2",
 "start_offset" : 5,
 "end_offset" : 8,
 "type" : "CN_WORD",
 "position" : 6
},
{
 "token" : "word-8",
 "start_offset" : 5,
 "end_offset" : 7,
 "type" : "CN_WORD",
 "position" : 7
},
{
 "token" : "word-9",
 "start_offset" : 6,
 "end_offset" : 8,
 "type" : "CN_WORD",
 "position" : 8
},
{
 "token" : "word-10",
 "start_offset" : 7,
 "end_offset" : 8,
 "type" : "CN_WORD",
 "position" : 9
}
}
```

**Step 2** Create an index and configure a word segmentation policy. After data is imported, search for data by keyword.

The commands used for Elasticsearch 7.x and later are different from those used for versions earlier than that.

Sample code for Elasticsearch earlier than 7.x:

1. Create the **book** index and configure a word segmentation policy.

In this example, both **analyzer** and **search\_analyzer** are set to **ik\_max\_word**. You can also use **ik\_smart**.

```
PUT /book
{
 "settings": {
 "number_of_shards": 2,
```

```

 "number_of_replicas": 1
 },
 "mappings": {
 "type1": {
 "properties": {
 "content": {
 "type": "text",
 "analyzer": "ik_max_word",
 "search_analyzer": "ik_max_word"
 }
 }
 }
 }
}

```

2. Import the text information to the **book** index.

```

PUT /book/type1/1
{
 "content": "Imported text"
}

```

3. Use a keyword to search for the text and view the search results.

```

GET /book/type1/_search
{
 "query": {
 "match": {
 "content": "Keyword"
 }
 }
}

```

#### Search result

```

{
 "took" : 20,
 "timed_out" : false,
 "_shards" : {
 "total" : 2,
 "successful" : 2,
 "skipped" : 0,
 "failed" : 0
 },
 "hits" : {
 "total" : 1,
 "max_score" : 1.1507283,
 "hits" : [
 {
 "_index" : "book",
 "_type" : "type1",
 "_id" : "1",
 "_score" : 1.1507283,
 "_source" : {
 "content" : "Imported text"
 }
 }
]
 }
}

```

Sample code for Elasticsearch 7.x or later:

1. Create the **book** index and configure the word segmentation policy.  
In this example, both **analyzer** and **search\_analyzer** are set to **ik\_max\_word**.  
You can also use **ik\_smart**.

```

PUT /book
{
 "settings": {
 "number_of_shards": 2,
 "number_of_replicas": 1
 },

```

```
"mappings": {
 "properties": {
 "content": {
 "type": "text",
 "analyzer": "ik_max_word",
 "search_analyzer": "ik_max_word"
 }
 }
}
```

2. Import the text information to the **book** index.

```
PUT /book/_doc/1
{
 "content": "Imported text"
}
```

3. Use a keyword to search for the text and view the search results.

```
GET /book/_doc/_search
{
 "query": {
 "match": {
 "content": "Keyword"
 }
 }
}
```

#### Search result

```
{
 "took" : 16,
 "timed_out" : false,
 "_shards" : {
 "total" : 2,
 "successful" : 2,
 "skipped" : 0,
 "failed" : 0
 },
 "hits" : {
 "total" : {
 "value" : 1,
 "relation" : "eq"
 },
 "max_score" : 1.7260926,
 "hits" : [
 {
 "_index" : "book",
 "_type" : "_doc",
 "_id" : "1",
 "_score" : 1.7260926,
 "_source" : {
 "content" : "Imported text"
 }
 }
]
 }
}
```

- Step 3** Create an index and configure a synonym policy. After data is imported, search for data by synonyms.

The commands used for Elasticsearch 7.x and later are different from those used for versions earlier than that.

Sample code for Elasticsearch earlier than 7.x:

1. Create the **myindex** index and configure the word segmentation policy.

```
PUT myindex
{
 "settings": {
```

```

"analysis": {
 "filter": {
 "my_synonym": {
 "type": "dynamic_synonym"
 }
 },
 "analyzer": {
 "ik_synonym": {
 "filter": [
 "my_synonym"
],
 "type": "custom",
 "tokenizer": "ik_smart"
 }
 }
},
"mappings": {
 "mytype": {
 "properties": {
 "desc": {
 "type": "text",
 "analyzer": "ik_synonym"
 }
 }
 }
}

```

2. Import the text information to the **myindex** index.

```

PUT /myindex/mytype/1
{
 "desc": "Imported text"
}

```

3. Conduct search based on the synonym and view the search results.

```

GET /myindex/_search
{
 "query": {
 "match": {
 "desc": "Keyword"
 }
 }
}

```

#### Search result

```

{
 "took" : 2,
 "timed_out" : false,
 "_shards" : {
 "total" : 5,
 "successful" : 5,
 "skipped" : 0,
 "failed" : 0
 },
 "hits" : {
 "total" : 1,
 "max_score" : 0.49445358,
 "hits" : [
 {
 "_index" : "myindex",
 "_type" : "mytype",
 "_id" : "1",
 "_score" : 0.49445358,
 "_source" : {
 "desc" : "Imported text"
 }
 }
]
 }
}

```

```
}
}
```

Sample code for Elasticsearch 7.x or later:

1. Create the **myindex** index and configure the word segmentation policy.

```
PUT myindex
{
 "settings": {
 "analysis": {
 "filter": {
 "my_synonym": {
 "type": "dynamic_synonym"
 }
 },
 "analyzer": {
 "ik_synonym": {
 "filter": [
 "my_synonym"
],
 "type": "custom",
 "tokenizer": "ik_smart"
 }
 }
 }
 },
 "mappings": {
 "properties": {
 "desc": {
 "type": "text",
 "analyzer": "ik_synonym"
 }
 }
 }
}
```

2. Import the text information to the **myindex** index.

```
PUT /myindex/_doc/1
{
 "desc": "Imported text"
}
```

3. Conduct search based on the synonym and view the search results.

```
GET /myindex/_search
{
 "query": {
 "match": {
 "desc": "Keyword"
 }
 }
}
```

Search result

```
{
 "took" : 1,
 "timed_out" : false,
 "_shards" : {
 "total" : 1,
 "successful" : 1,
 "skipped" : 0,
 "failed" : 0
 },
 "hits" : {
 "total" : {
 "value" : 1,
 "relation" : "eq"
 },
 "max_score" : 0.1519955,
 "hits" : [
 {

```

```
{
 "_index": "myindex",
 "_type": "_doc",
 "_id": "1",
 "_score": 0.1519955,
 "_source": {
 "desc": "Imported text"
 }
}
```

----End

## 2.13.11 Switching Between Simplified and Traditional Chinese for Data Search in an Elasticsearch Cluster

This topic describes how to switch between simplified and traditional Chinese for data search in an Elasticsearch cluster.

### Scenario

The simplified-traditional Chinese conversion plugin converts between simplified and traditional Chinese. With this plugin, you can search index data containing the corresponding simplified Chinese based on the traditional Chinese keyword, and vice versa.

This plugin is installed by default. You do not need to install it by yourself.

The simplified-traditional Chinese conversion plugin can be used as the analyzer, tokenizer, token-filter, or char-filter.

The simplified-traditional Chinese conversion plugin provides the following two conversion types:

- s2t: converts simplified Chinese to traditional Chinese.
- t2s: converts traditional Chinese to simplified Chinese.

### Switching Between Simplified and Traditional Chinese for Data Search

1. Log in to Kibana.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
  - c. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
  - d. In the left navigation pane, choose **Dev Tools**.
2. Run the following command to create the **stconvert** index and specify the user-defined mapping to define the data type:

Sample code for Elasticsearch clusters earlier than 7.x:

```
PUT /stconvert
{
 "settings": {
 "number_of_shards": 1,
 "number_of_replicas": 0,
 "analysis": {
 "analyzer": {
 "ts_ik": {
```



```
 "tokenizer": "ik_smart",
 "char_filter": [
 "tsconvert",
 "stconvert"
]
 },
 "char_filter": {
 "tsconvert": {
 "type": "stconvert",
 "convert_type": "t2s"
 },
 "stconvert": {
 "type": "stconvert",
 "convert_type": "s2t"
 }
 }
},
"mappings": {
 "type": {
 "properties": {
 "desc": {
 "type": "text",
 "analyzer": "ts_ik"
 }
 }
 }
}
```

Sample code for Elasticsearch 7.x or later:

```
PUT /stconvert
{
 "settings": {
 "number_of_shards": 1,
 "number_of_replicas": 0,
 "analysis": {
 "analyzer": {
 "ts_ik": {
 "tokenizer": "ik_smart",
 "char_filter": [
 "tsconvert",
 "stconvert"
]
 }
 },
 "char_filter": {
 "tsconvert": {
 "type": "stconvert",
 "convert_type": "t2s"
 },
 "stconvert": {
 "type": "stconvert",
 "convert_type": "s2t"
 }
 }
 }
 },
 "mappings": {
 "properties": {
 "desc": {
 "type": "text",
 "analyzer": "ts_ik"
 }
 }
 }
}
```

The command output is similar to the following:

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

3. Run the following command to import data to index **stconvert**:

Sample code for Elasticsearch clusters earlier than 7.x:

```
POST /stconvert/type/1
```

```
{
 "desc": "Text in traditional Chinese"
}
```

Sample code for Elasticsearch 7.x or later:

```
POST /stconvert/_doc/1
```

```
{
 "desc": "Text in traditional Chinese"
}
```

If the value of **failed** in the command output is **0**, the data is imported successfully.

4. Run the following command to search for the keyword and check the search result:

```
GET /stconvert/_search
```

```
{
 "query": {
 "match": {
 "desc": "Keyword"
 }
 }
}
```

The command output is similar to the following:

```
{
 "took" : 15,
 "timed_out" : false,
 "_shards" : {
 "total" : 1,
 "successful" : 1,
 "skipped" : 0,
 "failed" : 0
 },
 "hits" : {
 "total" : 1,
 "max_score" : 0.5753642,
 "hits" : [
 {
 "_index" : "stconvert",
 "_type" : "type",
 "_id" : "1",
 "_score" : 0.5753642,
 "_source" : {
 "desc": "Text in traditional Chinese"
 }
 }
]
 }
}
```

## 2.13.12 Restarting an Elasticsearch Cluster

If an Elasticsearch cluster has stopped working, you can try to recover it by restarting it. Two restart options are available: quick restart and rolling restart. You are advised to perform the restart during off-peak hours.

## Constraints

- Avoid restarting a cluster during peak hours, as doing so may impact data integrity and service availability.
- During a quick restart, the cluster is unavailable for the entire duration of the restart.
- Rolling restart is supported only for clusters with three or more nodes, including the master node, client node, and cold data nodes. During a rolling restart, only the nodes that are being restarted are unavailable. When the cluster stores a large amount of data, the restart may take a long time.

## Prerequisites

- The cluster you want to restart is not frozen and has no task in progress.
- The cluster is available and is not processing any service requests, such as importing and searching for data. Otherwise, data loss may occur during a cluster restart.

## Quick Restart

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and choose **More > Restart** in the **Operation** column.
4. On the **Restart Cluster** page, set parameters.

**Table 2-163** Configuring a quick restart

| Parameter       | Description                                                                                                                                                                                                                                                                                                                            |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Restart Mode    | Select <b>Quick Restart</b> .<br><b>WARNING</b><br>The cluster will be unavailable for the entire duration of a quick restart.                                                                                                                                                                                                         |
| Operation Role  | Select <b>Node type</b> or <b>Node name</b> . <ul style="list-style-type: none"><li>• <b>Node type</b>: You can select multiple node types to restart them at the same time.</li><li>• <b>Node name</b>: You can restart only one node at a time.</li></ul>                                                                            |
| Operation Value | Select one or more node types or a node name, depending on the setting of <b>Operation Role</b> .<br>Values for <b>Node type</b> : <ul style="list-style-type: none"><li>• <b>ess</b>: data node</li><li>• <b>ess-master</b>: master node</li><li>• <b>ess-client</b>: client node</li><li>• <b>ess-cold</b>: cold data node</li></ul> |

5. Click OK to start the restart task.

Refresh the page and check the cluster status. During the restart, the cluster status is **Processing**, and the task status is **Restarting**. When the cluster status changes to **Available**, the cluster has restarted successfully.

## Rolling Restart

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and choose **More > Restart** in the **Operation** column.
4. On the **Restart Cluster** page, set parameters.

**Table 2-164** Configuring a rolling restart

| Parameter       | Description                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Restart Mode    | Select <b>Rolling Restart</b> . <ul style="list-style-type: none"><li>• Rolling restart is supported only for clusters with three or more nodes, including the master node, client node, and cold data nodes.</li><li>• During a rolling restart, the nodes of a cluster are restarted one at a time. This may take a long time if the cluster has a large number of indexes.</li></ul> |
| Operation Role  | Only <b>Node type</b> is supported. You can select multiple node types to restart them at the same time.                                                                                                                                                                                                                                                                                |
| Operation Value | Select node types for the restart.<br>Values for <b>Node type</b> : <ul style="list-style-type: none"><li>• <b>ess</b>: data node</li><li>• <b>ess-master</b>: master node</li><li>• <b>ess-client</b>: client node</li><li>• <b>ess-cold</b>: cold data node</li></ul>                                                                                                                 |

5. Click OK to start the restart task.  
Refresh the page and check the cluster status. During the restart, the cluster status is **Processing**, and the task status is **Restarting**. When the cluster status changes to **Available**, the cluster has restarted successfully.

## 2.13.13 Deleting an Elasticsearch Cluster

You can delete clusters that you no longer need.

### Constraints

- Deleting a cluster will delete all its data. Any data that has not been backed up will be permanently lost. Exercise caution.
- A cluster that has no snapshots created for it cannot be restored once it is deleted. The snapshots of the cluster stored in OBS, if any, will not be deleted

with the cluster. You can restore a deleted cluster using its snapshots stored in the OBS bucket. For details, see [Can I Restore a Deleted Cluster?](#)

- For a cluster for which the VPCEP service has been enabled, deleting the cluster may not delete the VPC endpoint associated with the cluster if the current account does not have the necessary VPCEP permissions. For details about the VPCEP permissions, see [Permissions](#).

## Deleting a Pay-per-Use Cluster

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and choose **More > Delete** in the **Operation** column.
4. In the displayed dialog box, manually type in **DELETE**, and click **OK**.

## Deleting a Yearly/Monthly Cluster

To delete a yearly/monthly cluster, unsubscribe from it. After unsubscription, the cluster's resources and data will be permanently deleted. If you want to retain data, manually back up data before submitting the unsubscription request.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and choose **More > Unsubscribe/Release** in the **Operation** column.
4. In the displayed dialog box, manually type in **YES**, and click **OK**.

On the **Unsubscribe from Resource** page, confirm the resource information and refund amount.

5. Select a reason for unsubscription, select the acknowledgement check boxes, and click **Confirm**.

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

### NOTE

If the cluster is in **Available** state, an unsubscription order will be generated for refund, and the cluster will be deleted. If the cluster has expired or has been frozen, it will be directly deleted. For more information, see [Billing Center](#).

## 2.14 Managing Index Policies for an Elasticsearch Cluster

### 2.14.1 Creating and Managing Index Policies for an Elasticsearch Cluster

This topic describes how to create and manage index lifecycle policies for Elasticsearch clusters.

## Context

Index State Management (ISM) of Elasticsearch is a plugin that allows you to automate periodic, administrative operations on indexes by triggering them based on changes in the index age, index size, or number of documents. With ISM, you can define custom policies to automate index rollovers and deletion, thus optimizing cluster search performance or cutting storage costs. The procedure for using ISM is as follows:

1. **Creating an Index Lifecycle Policy:** Create an index lifecycle policy on OpenSearch Dashboards.
2. **Associating Indexes with a Lifecycle Policy:** Associate indexes with a lifecycle policy.
3. **Managing Index Policies:** Modify, retry, and change index lifecycle policies.

For more information about ISM, see [Index State Management](#).

## Constraints

- ISM is available only in Elasticsearch 7.6.2 clusters or later.
- You can customize policy names in Kibana. Chinese characters are not supported.
- The Kibana GUI varies depending on the Kibana version. Kibana 7.6.2 is used as an example here.

## Creating an Index Lifecycle Policy

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
4. On the Kibana console, expand the menu in the upper-left corner, and choose **IM** or **Index Management** to go to the index management page.
5. Click **Create policy** to configure an index lifecycle policy.
  - **Policy ID:** use a custom policy name.
  - **Define policy:** define a custom policy based on the reference example provided.

**Figure 2-63** Configuring a policy

Create policy

Name policy

Policy ID

hot\_cold\_workflow

Define policy

```
{
 "policy": {
 "description": "A simple default policy that changes the replica count between hot and cold states.",
 "default_state": "hot",
 "states": [
 {
 "name": "hot",
 "actions": [
 {
 "replica_count": {
 "number_of_replicas": 5
 }
 }
],
 "transitions": [
 {
 "state_name": "cold",
 "conditions": {
 "min_index_age": "360"
 }
 }
]
 },
 {
 "name": "cold",
 "actions": [
 {
 "replica_count": {
 "number_of_replicas": 2
 }
 }
]
 }
]
 }
}
```

6. Click **Create**.

## Associating Indexes with a Lifecycle Policy

You can attach a policy to one or more indexes, or associate a policy with an index template, so that the policy will be applied to all indexes created using this index template.

Use either of the following methods to associate a lifecycle policy with an index.

- **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**

- a. On the Kibana console, choose **IM** or **Index Management** from the menu bar to go to the index management page.
- b. Select **Indices** on the left.

Figure 2-64 Choosing Indices



- c. In the **Indices** list, select the target index to which you want to bind a policy.
- d. Click **Apply policy** in the upper right corner. In the displayed dialog box, select a policy ID.

Figure 2-65 Adding an index policy





Figure 2-66 Selecting a policy



- e. Click **Apply**.  
After you bind a policy to an index, ISM creates a job that runs every 5 minutes by default to execute the policy, check criteria, and change index states.



## Managing Index Policies

1. On the **Index Management** page of the Kibana console, choose **Managed Indices**.  
The displayed page shows index policies configured for the current cluster.
2. Manage policies in the index policy list.
  - For a policy whose status is abnormal, click **Retry policy**.
  - To update a policy for an index, select the associated index, and click **Change policy**. After the policy is updated, click **Change** to apply the new policy.
  - To remove a policy from an index, select the index, and click **Remove policy**. In the displayed dialog box, click **Remove**.

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


### 2.14.2 Automating Index Rollover in an Elasticsearch Cluster Through Index Lifecycle Management

As time series data is continuously written into Elasticsearch and the index size keeps growing, you can configure Index Lifecycle Management (ISM) to periodically roll over data to new indexes and delete old indexes.

In this example, a lifecycle policy is created for an Elasticsearch cluster. According to this policy, when the size of an index reaches 1 TB or the index is one day old, a new index will be automatically generated; seven days after the index is created, data duplicates will be disabled; 30 days after the index is created, the index will be deleted.

Assume that an index generates approximately 2.4 TB of data each day. The index alias is **log-alias**. The following figure illustrates the data format in the cluster. During reads, it points to all indexes whose name starts with **test**. During writes, it points to the latest index.

Figure 2-67 log-alias format



#### NOTE

The one day in the rollover time refers to 24 hours following the index creation time, not a calendar day.

## Prerequisites

The target Elasticsearch cluster is available.

## Automating Index Rollover Through Index Lifecycle Management

### NOTICE

The Elasticsearch cluster version must be 7.6.2 or later.

1. Log in to Kibana.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
  - c. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
  - d. In the left navigation pane, choose **Dev Tools**.
2. Create a rollover lifecycle policy named **rollover\_workflow**.

Policy description: When the size of an index reaches 1 TB or the index has been created for more than one day, the index rollover is performed. When the index has been created for seven days, the data copy is disabled. When the index has been created for 30 days, the index is deleted.

```
PUT _opendistro/_ism/policies/rollover_workflow
```

```
{
 "policy": {
 "description": "rollover test",
 "default_state": "hot",
 "states": [
 {
 "name": "hot",
 "actions": [
 {
 "rollover": {
 "min_size": "1tb",
 "min_index_age": "1d"
 }
 }
],
 "transitions": [
 {
 "state_name": "warm",
 "conditions": {
 "min_index_age": "7d"
 }
 }
]
 },
 {
 "name": "warm",
 "actions": [
 {
 "replica_count": {
 "number_of_replicas": 0
 }
 }
],
 "transitions": [
 {
 "state_name": "delete",
 "conditions": {
 "min_index_age": "30d"
 }
 }
]
 }
]
 }
}
```

```
 },
 {
 "name": "delete",
 "actions": [
 {
 "delete": {}
 }
]
 }
]
}
```

After a lifecycle policy is created, run the following command to query the policy details:

```
GET _opendistro/_ism/policies/rollover_workflow
```

3. Create the index template **template\_test**.

Template description: All the new indexes starting with **test** are automatically associated with the rollover lifecycle policy **rollover\_workflow**. The alias **log\_alias** is used during rollover.

```
PUT _template/template_test
```

```
{
 "index_patterns": "test-*",
 "settings": {
 "number_of_replicas": 1,
 "number_of_shards": 1,
 "opendistro.index_state_management.policy_id": "rollover_workflow",
 "index.opendistro.index_state_management.rollover_alias": "log_alias"
 },
 "mappings": {
 "properties": {
 "name": {
 "type": "text"
 }
 }
 }
}
```

**Table 2-165** Parameter description

| Parameter                                              | Description                    |
|--------------------------------------------------------|--------------------------------|
| number_of_shards                                       | Number of index shards         |
| number_of_replicas                                     | Number of index shard replicas |
| opendistro.index_state_management.policy_id            | Lifecycle policy name          |
| index.opendistro.index_state_management.rollover_alias | Index alias for rollover       |

After an index template is created, you can run the following command to query the template details:

```
GET _template/template_test
```

4. Create an index, specify **aliases**, and set **is\_write\_index** to **true**. The index template **template\_test** is automatically used for the index and is associated with the lifecycle policy **rollover\_workflow** based on the index template

configuration. In this way, when the index size reaches 1 TB or the index is created for more than one day, the rollover automatically starts. After an index is created for seven days, the data copy is disabled. After an index is created for 30 days, the index is deleted.

The following index is the URL code of **<test-{now/d}-000001>**. By default, an index name contains the creation date. For example, if an index is created on 2022-06-02, the index name is **test-2022.06.02-000001**.

```
PUT %3Ctest-%7Bnow%2Fd%7D-000001%3E
```

```
{
 "aliases": {
 "log_alias": {
 "is_write_index": true
 }
 }
}
```

5. The alias **log\_alias** is used to during data write, and **log\_alias** always points to the last index.

```
POST log_alias/_bulk
```

```
{ "index": {} }
{ "name": "name1" }
{ "index": {} }
{ "name": "name2" }
{ "index": {} }
{ "name": "name3" }
{ "index": {} }
{ "name": "name4" }
{ "index": {} }
{ "name": "name5" }
{ "index": {} }
{ "name": "name6" }
```

6. Query data and check whether the rollover takes effect.
- One day after the indexes are created, check the indexes starting with **test**.

```
GET _cat/indices/test*?s=i
```

There are supposed to be at least two indexes, for example:

```
green open test-<Date>-000001 r8ab5NX6T3Ox_hoGUanogQ 1 1 6 0 416b 208b
green open test-<Date>-000002 sfwkVgy8RSSEw7W-xYjM2Q 1 1 0 0 209b 209b
```

In the preceding information, **test-<Date>-000001** is the index created in [4](#), and **test-<Date>-000002** is the index generated through rollover.

- To query the index associated with the alias **log\_alias**, run the following command:

```
GET _cat/aliases/log_alias?v
```

The alias is supposed to point to multiple indexes, for example:

| alias     | index              | filter | routing | index | routing | search | is_write_index |
|-----------|--------------------|--------|---------|-------|---------|--------|----------------|
| log_alias | test-<Date>-000001 | -      | -       | -     | -       | -      | false          |
| log_alias | test-<Date>-000002 | -      | -       | -     | -       | -      | true           |

## FAQ: How Do I Skip Rollover for an Index?

### Scenarios

- A manual rollover is performed: A time-based ISM policy was configured to automate index rollover. Before the rollover criteria were met, however, a manual rollover was performed. When the auto rollover time configured by ISM was reached, an automatic rollover was attempted but failed because the index had already been rolled over manually. As a result, the ISM task stopped. To solve this problem, you need to skip the index rollover and then retry the ISM task to resume automatic index rollover.

- For an index that has already been rolled over automatically, a remove or add policy operation is performed: After the rollover policy is modified (through the remove policy or add policy operation), the ISM task restarts from the beginning. When rollover starts, the task fails. To solve this, you also need to skip index rollover and then retry the ISM task to resume automatic index rollover.

**WARNING**

After rollover is skipped for an index, ISM will no longer attempt rollover or generate rolled-over indexes. This means skipping index rollover may lead to data loss. Please exercise caution.

### Constraints

To skip index rollover, the Elasticsearch cluster version must be 7.6.2 or 7.10.2, and the cluster image version must not be earlier than 7.x.2\_25.1.0\_x.x.x.

### Procedure

1. If an ISM task stops due to an index rollover failure, run the following command to skip index rollover:

```
PUT index_name/_settings
{
 "index.plugins.index_state_management.rollover_skip": true
}
```

If **true** is returned, the configuration is successful.

2. After index rollover is skipped, run the following command to retry the ISM task:

```
POST _opendistro/_ism/retry/index_name
```

If the following information is returned, the retry is successful:


```
{
 "updated_indices": 1,
 "failures": false,
 "failed_indices": []
}
```

## 2.14.3 Decoupling Index Storage and Compute in an Elasticsearch Cluster Through Index Lifecycle Management

### Overview

CSS supports decoupled storage and compute. That is, indexes can be frozen and their data can be dumped to OBS to reduce the storage cost of cold data. This document describes how to use index lifecycle management to automatically freeze indexes at a specific time to decouple storage and compute.

In this example, a lifecycle policy is created for an Elasticsearch 7.10.2 cluster. According to this policy, a newly created index is automatically frozen in three days, with data dumped to OBS; and deleted in seven days.

**Figure 2-68** Storage-compute decoupling

## Constraints

Only Elasticsearch 7.6.2 and 7.10.2 clusters support decoupled storage and compute.

## Prerequisites

The Elasticsearch cluster is available.

## Decoupling Index Storage and Compute Through Index Lifecycle Management

1. Log in to Kibana.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
  - c. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
  - d. In the left navigation pane, choose **Dev Tools**.
2. Create a lifecycle policy named **hot\_warm\_policy**.

Policy description: Three days after an index is created, the API for freezing indexes is automatically called to dump data to OBS. Seven days after index creation, the index is deleted.

```
PUT _opendistro/_ism/policies/hot_warm_policy
{
 "policy": {
 "description": "hot warm archive delete workflow",
 "error_notification": null,
 "default_state": "hot",
 "states": [
 {
 "name": "hot",
 "actions": [],
 "transitions": [
 {
 "state_name": "warm",
 "conditions": {
 "min_index_age": "3d"
 }
 }
]
 },
 {
 "name": "warm",
 "actions": [
 {
 "freeze_low_cost": {}
 }
]
 }
]
 }
}
```

```
 }
],
 "transitions": [
 {
 "state_name": "delete",
 "conditions": {
 "min_index_age": "7d"
 }
 }
]
},
{
 "name": "delete",
 "actions": [
 {
 "delete": {}
 }
],
 "transitions": []
}
]
}
```

3. Create the index template **template\_hot\_warm**.  
Template description: All the new indexes starting with **data** are automatically associated with the lifecycle policy **hot\_warm\_policy**.

```
PUT _template/template_hot_warm
{
 "index_patterns": "data*",
 "settings": {
 "number_of_replicas": 5,
 "number_of_shards": 1,
 "opendistro.index_state_management.policy_id": "hot_warm_policy"
 },
 "mappings": {
 "properties": {
 "name": {
 "type": "text"
 }
 }
 }
}
```

**Table 2-166** Parameter description

| Parameter                                   | Description                    |
|---------------------------------------------|--------------------------------|
| number_of_shards                            | Number of index shards         |
| number_of_replicas                          | Number of index shard replicas |
| opendistro.index_state_management.policy_id | Lifecycle policy name          |

4. Create the **data-2022-06-06** index. The index automatically uses the **template\_hot\_warm** template and associates the index template with the lifecycle policy **hot\_warm\_policy**. As such, the index is automatically frozen in three days, and deleted in seven days after creation.

```
POST data-2022-06-06/_bulk
{"index":{}}
{"name":"name1"}
```

```
{"index":{}}
{"name":"name2"}
{"index":{}}
{"name":"name3"}
{"index":{}}
{"name":"name4"}
{"index":{}}
{"name":"name5"}
{"index":{}}
{"name":"name6"}
```

5. Query data and check whether storage and compute is automatically decoupled.

- Three days after the index is created, check the frozen index.

```
GET _cat/freeze_indices?s=i&v
```

The index generated three days ago is expected to be frozen.

| health | status | index           | uuid                   | pri | rep | docs.count | docs.deleted | store.size |
|--------|--------|-----------------|------------------------|-----|-----|------------|--------------|------------|
| green  | open   | data-2022-06-06 | x8ab5NX6T3Ox_xoGUanogQ | 1   | 1   | 6          | 0            |            |
| 7.6kb  | 3.8kb  |                 |                        |     |     |            |              |            |

- Seven days after index creation, check the index. The index is expected to be deleted.

## 2.15 Elasticsearch Cluster Monitoring and Log Management

### 2.15.1 Elasticsearch Cluster Monitoring Metrics

#### 2.15.1.1 Monitoring Metrics for Elasticsearch Clusters in Cloud Eye

You can use Cloud Eye to monitor CSS cluster metrics in real time and quickly handle exceptions.

##### Overview

You can use Cloud Eye to monitor CSS cluster metrics in real time and quickly handle exceptions. This topic describes CSS cluster and node metrics that can be monitored through Cloud Eye, as well as their namespaces and dimensions. You can query metrics and alerts generated by CSS on the [Cloud Eye console](#) or an [API](#)

##### Namespace

SYS.ES

##### Cluster Monitoring Metrics

###### NOTE

Accumulated value: The value is accumulated from when a node is started. After the node is restarted, the value is reset to zero and starts accumulating again.



Table 2-167 CSS cluster metrics

| Metric ID | Metric                | Description                                                                                                                                                                                                                                                                                                                                                                                                                     | Value Range | Unit | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------|---------------|------------------------------|---------------------------|
| status    | Cluster Health Status | Health status of the monitored object <ul style="list-style-type: none"><li>0: The cluster is 100% available.</li><li>1: All primary shards are allocated, but some replicas are 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</li></ul> | 0, 1, 2, 3  | N/A  | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID              | Metric                     | Description                                                                                                                                                                  | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|------------------------|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
|                        |                            | investigation . <ul style="list-style-type: none"><li>• 2: Data is missing and the cluster cannot work properly.</li><li>• 3: Failed to obtain the cluster status.</li></ul> |             |       |               |                              |                           |
| disk_util              | Max. Disk Usage            | Maximum disk usage of nodes in a CSS cluster                                                                                                                                 | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| max_jvm_heap_usage     | Max. JVM Heap Usage        | Maximum JVM heap usage of nodes in a CSS cluster.                                                                                                                            | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| max_jvm_young_gc_time  | Max. JVM Young GC Duration | Maximum accumulated JVM Young GC duration of nodes in a CSS cluster.                                                                                                         | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |
| max_jvm_young_gc_count | Max. JVM Young GC Count    | Maximum accumulated JVM Young GC count of nodes in a CSS cluster.                                                                                                            | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                   | Metric                         | Description                                                                    | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------|--------------------------------|--------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| max_jvm_old_gc_time         | Max. JVM Old GC Duration       | Maximum accumulated JVM Old GC duration of nodes in a CSS cluster.             | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |
| max_jvm_old_gc_count        | Max. JVM Old GC Count          | Maximum accumulated JVM Old GC count of nodes in a CSS cluster.                | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| total_fs_size               | Total Size of File Systems     | Total size of file systems in a CSS cluster.                                   | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| free_fs_size                | Available Size of File Systems | Available size of file systems in a CSS cluster.                               | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| max_cpu_usage               | Max. CPU Usage                 | Maximum node CPU usage in a CSS cluster.                                       | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| 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. | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                              | Metric                                  | Description                                                                            | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|----------------------------------------|-----------------------------------------|----------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| 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.                    | $\geq 0$    | byte  | 10 24         | CSS cluster                  | 1 minute                  |
| max_current_opened_http_count          | Current Max. Opened HTTP Connections    | Maximum number of HTTP connections that are currently open for nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_total_opened_http_count            | Total Max. Opened HTTP Connections      | Maximum number of HTTP connections that were open for nodes in a CSS cluster.          | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| indices_count                          | Indexes                                 | Number of indexes in a CSS cluster                                                     | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| total_shards_count                     | Shards                                  | Number of shards in a CSS cluster                                                      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| primary_shards_count                   | Primary Shards                          | Number of primary shards in a CSS cluster                                              | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                | Metric            | Description                                  | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|--------------------------|-------------------|----------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| docs_count               | Documents         | Number of documents in a CSS cluster         | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| docs_deleted_count       | Deleted Documents | Number of documents deleted in a CSS cluster | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| nodes_count              | Nodes             | Number of nodes in a CSS cluster             | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| data_nodes_count         | Data Nodes        | Number of data nodes in a CSS cluster        | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| coordinating_nodes_count | Coordinator Nodes | Number of coordinator nodes in a CSS cluster | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| master_nodes_count       | Master Nodes      | Number of master nodes in a CSS cluster      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| ingest_nodes_count       | Client Nodes      | Number of client nodes in a CSS cluster      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                 | Metric                     | Description                                                                       | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|---------------------------|----------------------------|-----------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| max_load_average          | Max. Node Load             | Maximum number of average queued tasks per minute on nodes in a cluster           | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_cpu_usage             | Avg. CPU Usage             | Average node CPU usage in a CSS cluster.                                          | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| avg_load_average          | Avg. Node Load             | Average number of queued tasks per minute on nodes in a CSS cluster               | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_jvm_heap_usage        | Avg. JVM Heap Usage        | Average node JVM heap usage in a CSS cluster.                                     | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| max_open_file_descriptors | Max. Open File Descriptors | Maximum number of node file descriptors that are currently open in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                         | Metric                        | Description                                                                               | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------------|-------------------------------|-------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_open_file_descriptors         | Avg. Open File Descriptors    | Average number of node file descriptors that are currently open in a CSS cluster.         | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_max_file_descriptors          | Max. Allowed File Descriptors | Maximum number of allowed node file descriptors in a CSS cluster.                         | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_open_file_descriptors         | Open File Descriptors         | Number of node file descriptors that are currently open in a cluster.                     | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_write_queue       | Tasks in Write Queue          | Number of tasks in queue for the write thread pool.                                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_search_queue      | Tasks in Search Queue         | Total number of tasks in queue for the search thread pool on nodes in a CSS cluster.      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_force_merge_queue | Tasks in ForceMerge Queue     | Total number of tasks in queue for the force merge thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                            | Metric                             | Description                                                                                 | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|--------------------------------------|------------------------------------|---------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_write_rejected       | Rejected Tasks in Write Queue      | Total number of tasks rejected by the write thread pool on nodes in a CSS cluster.          | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_search_rejected      | Rejected Tasks in Search Queue     | Total number of tasks rejected by the search thread pool on nodes in a CSS cluster.         | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_force_merge_rejected | Rejected Tasks in ForceMerge Queue | Total number of tasks rejected by the force merge thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_search_queue         | Max. Tasks in Search Queue         | Maximum number of tasks in queue for the search thread pool on nodes in a CSS cluster.      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_force_merge_queue    | Max. Tasks in ForceMerge Queue     | Maximum number of tasks in queue for the force merge thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |



| Metric ID                           | Metric                         | Description                                                                                 | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------------|--------------------------------|---------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_write_threads       | Size of Write Thread Pool      | Total size of the write thread pools on nodes in the CSS cluster.                           | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_search_threads      | Size of Search Thread Pool     | Total size of the search thread pools on nodes in the CSS cluster.                          | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_force_merge_threads | Size of ForceMerge Thread Pool | Total size of the force merge thread pools on nodes in the CSS cluster.                     | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_write_queue         | Avg. Tasks in Write Queue      | Average number of tasks in queue for the write thread pool on nodes in a CSS cluster.       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_search_queue        | Avg. Tasks in Search Queue     | Average number of tasks in queue for the search thread pool on nodes in a CSS cluster.      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_force_merge_queue   | Avg. Tasks in ForceMerge Queue | Average number of tasks in queue for the force merge thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                           | Metric                              | Description                                                                          | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------------|-------------------------------------|--------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_thread_pool_search_threads      | Avg. Size of Search Thread Pool     | Average size of the search thread pool on nodes in a CSS cluster.                    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_write_threads       | Avg. Size of Write Thread Pool      | Average size of the write thread pool on nodes in a CSS cluster.                     | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_force_merge_threads | Avg. Size of ForceMerge Thread Pool | Average size of the force merge thread pool on nodes in a CSS cluster.               | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_write_rejected      | Avg. Rejected Tasks in Write Queue  | Average number of tasks rejected by the write thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_flush_queue         | Tasks in Flush Queue                | Total number of tasks in queue for the flush thread pool on nodes in a CSS cluster.  | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_flush_rejected      | Rejected Tasks in Flush Queue       | Total number of tasks rejected by the flush thread pool on nodes in a CSS cluster.   | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                     | Metric                         | Description                                                                           | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------|--------------------------------|---------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| max_thread_pool_flush_queue   | Max. Tasks in Flush Queue      | Maximum number of tasks in queue for the flush thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_flush_threads | Size of Flush Thread Pool      | Total size of the flush thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_flush_queue   | Avg. Tasks in Flush Queue      | Average number of tasks in queue for the flush thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_flush_threads | Avg. Size of Flush Thread Pool | Average number of tasks in queue for the flush thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_generic_queue | Tasks in Generic Queue         | Total number of tasks in queue for the generic thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                        | Metric                           | Description                                                                             | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|----------------------------------|----------------------------------|-----------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_generic_rejected | Rejected Tasks in Generic Queue  | Total number of tasks rejected by the generic thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_generic_queue    | Max. Tasks in Generic Queue      | Maximum number of tasks in queue for the generic thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_generic_threads  | Size of Generic Thread Pool      | Total size of the generic thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_generic_queue    | Avg. Tasks in Generic Queue      | Average number of tasks in queue for the generic thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_generic_threads  | Avg. Size of Generic Thread Pool | Average number of tasks in queue for the generic thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                           | Metric                             | Description                                                                                | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------------|------------------------------------|--------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_management_queue    | Tasks in Management Queue          | Total number of tasks in queue for the management thread pool on nodes in a CSS cluster.   | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_management_rejected | Rejected Tasks in Management Queue | Total number of tasks rejected by the management thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_management_queue    | Max. Tasks in Management Queue     | Maximum number of tasks in queue for the management thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_management_threads  | Size of Management Thread Pool     | Total size of the management thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_management_queue    | Avg. Tasks in Management Queue     | Average number of tasks in queue for the management thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                          | Metric                              | Description                                                                                | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|------------------------------------|-------------------------------------|--------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_thread_pool_management_threads | Avg. Size of Management Thread Pool | Average number of tasks in queue for the management thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_refresh_queue      | Tasks in Refresh Queue              | Total number of tasks in queue for the refresh thread pool on nodes in a CSS cluster.      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_refresh_rejected   | Rejected Tasks in Refresh Queue     | Total number of tasks rejected by the refresh thread pool on nodes in a CSS cluster.       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_refresh_queue      | Max. Tasks in Refresh Queue         | Maximum number of tasks in queue for the refresh thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_refresh_threads    | Size of Refresh Thread Pool         | Total size of the refresh thread pools on nodes in a CSS cluster.                          | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                             | Metric                               | Description                                                                                  | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|---------------------------------------|--------------------------------------|----------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_thread_pool_refresh_queue         | Avg. Tasks in Refresh Queue          | Average number of tasks in queue for the refresh thread pool on nodes in a CSS cluster.      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_refresh_threads       | Avg. Size of Refresh Thread Pool     | Average number of tasks in queue for the refresh thread pool on nodes in a CSS cluster.      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_searcher_queue    | Tasks in OBS Searcher Queue          | Total number of tasks in queue for the OBS Searcher thread pool on nodes in a CSS cluster.   | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_searcher_rejected | Rejected Tasks in OBS Searcher Queue | Total number of tasks rejected by the OBS Searcher thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_obs_searcher_queue    | Max. Tasks in OBS Searcher Queue     | Maximum number of tasks in queue for the OBS Searcher thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                            | Metric                                | Description                                                                                  | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|--------------------------------------|---------------------------------------|----------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_obs_searcher_threads | Size of OBS Searcher Thread Pool      | Total size of the OBS Searcher thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_searcher_queue   | Avg. Tasks in OBS Searcher Queue      | Average number of tasks in queue for the OBS Searcher thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_searcher_threads | Avg. Size of OBS Searcher Thread Pool | Average number of tasks in queue for the OBS Searcher thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_queue            | Tasks in OBS Queue                    | Total number of tasks in queue for the OBS thread pool on nodes in a CSS cluster.            | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_rejected         | Rejected Tasks in OBS Queue           | Total number of tasks rejected by the OBS thread pool on nodes in a CSS cluster.             | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |



| Metric ID                        | Metric                       | Description                                                                              | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|----------------------------------|------------------------------|------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| max_thread_pool_obs_queue        | Max. Tasks in OBS Queue      | Maximum number of tasks in queue for the OBS thread pool on nodes in a CSS cluster.      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_threads      | Size of OBS Thread Pool      | Total size of the OBS thread pools on nodes in a CSS cluster.                            | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_queue        | Avg. Tasks in OBS Queue      | Average number of tasks in queue for the OBS thread pool on nodes in a CSS cluster.      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_threads      | Avg. Size of OBS Thread Pool | Average number of tasks in queue for the OBS thread pool on nodes in a CSS cluster.      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_upload_queue | Tasks in OBS Upload Queue    | Total number of tasks in queue for the OBS Upload thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                           | Metric                              | Description                                                                                | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------------|-------------------------------------|--------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_obs_upload_rejected | Rejected Tasks in OBS Upload Queue  | Total number of tasks rejected by the OBS Upload thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_obs_upload_queue    | Max. Tasks in OBS Upload Queue      | Maximum number of tasks in queue for the OBS Upload thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_upload_threads  | Size of OBS Upload Thread Pool      | Total size of the OBS Upload thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_upload_queue    | Avg. Tasks in OBS Upload Queue      | Average number of tasks in queue for the OBS Upload thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_upload_threads  | Avg. Size of OBS Upload Thread Pool | Average number of tasks in queue for the OBS Upload thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                             | Metric                               | Description                                                                                  | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|---------------------------------------|--------------------------------------|----------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_obs_download_queue    | Tasks in OBS Download Queue          | Total number of tasks in queue for the OBS Download thread pool on nodes in a CSS cluster.   | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_download_rejected | Rejected Tasks in OBS Download Queue | Total number of tasks rejected by the OBS Download thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_obs_download_queue    | Max. Tasks in OBS Download Queue     | Maximum number of tasks in queue for the OBS Download thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_download_threads  | Size of OBS Download Thread Pool     | Total size of the OBS Download thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_download_queue    | Avg. Tasks in OBS Download Queue     | Average number of tasks in queue for the OBS Download thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                            | Metric                                 | Description                                                                                  | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|--------------------------------------|----------------------------------------|----------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_thread_pool_obs_download_threads | Avg. Size of OBS Download Thread Pool  | Average number of tasks in queue for the OBS Download thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| min_free_fs_size                     | Min. Available Storage Space           | Minimum available storage space of nodes in a CSS cluster.                                   | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| avg_jvm_old_gc_count                 | Avg. GCs of Old-Generation JVM         | Average number of old-generation garbage collections of nodes in a CSS cluster.              | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_jvm_old_gc_time                  | Avg. GC Duration of Old-Generation JVM | Average old-generation garbage collection duration of nodes in a CSS cluster.                | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |
| avg_jvm_young_gc_count               | Avg. GCs of Young-Generation JVM       | Average number of young-generation garbage collections of nodes in a CSS cluster.            | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                | Metric                                   | Description                                                                                    | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|--------------------------|------------------------------------------|------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_jvm_young_gc_time    | Avg. GC Duration of Young-Generation JVM | Average young-generation garbage collection duration of nodes in a CSS cluster.                | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |
| 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. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_mem_free_in_bytes    | Avg. Available Memory                    | Average unused memory capacity of nodes in a CSS cluster.                                      | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| avg_mem_free_percent     | Avg. Available Memory Percentage         | Average percentage of unused memory of nodes in a CSS cluster.                                 | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| avg_mem_used_in_bytes    | Avg. Used Memory                         | Average used memory of nodes in a CSS cluster.                                                 | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| avg_mem_used_percent     | Avg. Used Memory Percentage              | Average percentage of used memory of nodes in a CSS cluster.                                   | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID             | Metric                                  | Description                                                                   | Value Range | Unit  | Number of Systems | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------|-----------------------------------------|-------------------------------------------------------------------------------|-------------|-------|-------------------|------------------------------|---------------------------|
| max_mem_free_in_bytes | Max. Available Memory                   | Maximum unused memory of nodes in a CSS cluster.                              | $\geq 0$    | byte  | 1024              | CSS cluster                  | 1 minute                  |
| max_mem_free_percent  | Max. Available Memory Percentage        | Maximum percentage of unused memory of nodes in a CSS cluster.                | 0~100       | %     | N/A               | CSS cluster                  | 1 minute                  |
| max_mem_used_in_bytes | Max. Used Memory                        | Maximum used memory of nodes in a CSS cluster.                                | $\geq 0$    | byte  | 1024              | CSS cluster                  | 1 minute                  |
| max_mem_used_percent  | Max. Used Memory Percentage             | Maximum percentage of used memory of nodes in a CSS cluster.                  | 0~100       | %     | N/A               | CSS cluster                  | 1 minute                  |
| sum_jvm_old_gc_count  | Total GCs of Old-Generation JVM         | Number of old-generation garbage collections of nodes in a CSS cluster.       | $\geq 0$    | Count | N/A               | CSS cluster                  | 1 minute                  |
| sum_jvm_old_gc_time   | Total GC Duration of Old-Generation JVM | Total old-generation garbage collection duration of nodes in the CSS cluster. | $\geq 0$    | ms    | N/A               | CSS cluster                  | 1 minute                  |

| Metric ID                     | Metric                                    | Description                                                                     | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------|-------------------------------------------|---------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_jvm_young_gc_count        | Total GCs of Young-Generation JVM         | Number of young-generation garbage collections of nodes in a CSS cluster.       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_jvm_young_gc_time         | Total GC Duration of Young-Generation JVM | Total young-generation garbage collection duration of nodes in the CSS cluster. | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |
| sum_current_opened_http_count | Currently Open HTTP Connections           | Number of HTTP connections that are open on nodes in a CSS cluster.             | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_total_opened_http_count   | Historical Open HTTP Connections          | Number of HTTP connections that were open on nodes in a CSS cluster.            | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| IndexingLatency               | Average Index Latency                     | Average time required for a shard to complete an indexing operation.            | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                   | Metric                           | Description                                                                    | Value Range | Unit      | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------|----------------------------------|--------------------------------------------------------------------------------|-------------|-----------|---------------|------------------------------|---------------------------|
| IndexingRate                | Average Index Rate               | Average number of index operations per second in a cluster.                    | $\geq 0$    | Count     | N/A           | CSS cluster                  | 1 minute                  |
| SearchLatency               | Average Search Latency           | Average time required for a shard to complete a search operation.              | $\geq 0$    | ms        | N/A           | CSS cluster                  | 1 minute                  |
| SearchRate                  | Average QPS                      | Average queries per second (QPS) in a cluster.                                 | $\geq 0$    | Count     | N/A           | CSS cluster                  | 1 minute                  |
| task_max_running_time       | Max. Task Runtime                | Duration of the most time-consuming search or write task in the cluster.       | $\geq 0$    | ms        | N/A           | CSS cluster                  | 1 minute                  |
| number_of_pending_tasks     | Number of Pending Tasks in Queue | Number of pending tasks to be processed by the master node in the CSS cluster. | $\geq 0$    | Count     | N/A           | CSS cluster                  | 1 minute                  |
| sum_disk_read_requests_rate | Total Disk Read IOPS             | Number of disk read requests to all nodes in the CSS cluster per second.       | $\geq 0$    | request/s | N/A           | CSS cluster                  | 1 minute                  |



| Metric ID                               | Metric                                          | Description                                                                                     | Value Range | Unit       | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------------------|-------------------------------------------------|-------------------------------------------------------------------------------------------------|-------------|------------|---------------|------------------------------|---------------------------|
| sum_disk_write_requests_rate            | Total Disk Write IOPS                           | Number of disk write requests to all nodes in the CSS cluster per second.                       | $\geq 0$    | requests/s | N/A           | CSS cluster                  | 1 minute                  |
| sum_disk_read_bytes_rate                | Total Disk Read Rate                            | Number of bytes of data read from the disks of all nodes in the CSS cluster per second.         | $\geq 0$    | Byte/s     | 1024          | CSS cluster                  | 1 minute                  |
| sum_disk_write_bytes_rate               | Total Disk Write Rate                           | Number of bytes of data written into the disks of all nodes in the CSS cluster per second.      | $\geq 0$    | Byte/s     | 1024          | CSS cluster                  | 1 minute                  |
| avg_vector_index_off_heap_used_in_bytes | Average Off-Heap Memory Usage of Vector Indexes | Average off-heap memory (bytes) used for querying vector indexes on each node in a CSS cluster. | $\geq 0$    | byte       | 1024          | CSS cluster                  | 1 minute                  |

| Metric ID                               | Metric                                                    | Description                                                                                                                                                           | Value Range | Unit | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------------------|-----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------|---------------|------------------------------|---------------------------|
| avg_vector_index_off_heap_usage         | Average Off-Heap Memory Utilization (%) of Vector Indexes | Average off-heap memory (in %) used for querying vector indexes on each node in a CSS cluster.<br><br>This metric is calculated only when vector index queries occur. | 0~100       | %    | N/A           | CSS cluster                  | 1 minute                  |
| max_vector_index_off_heap_used_in_bytes | Maximum Off-Heap Memory Usage of Vector Indexes           | Maximum off-heap memory (bytes) used for querying vector indexes on each node in a CSS cluster.                                                                       | $\geq 0$    | byte | 1024          | CSS cluster                  | 1 minute                  |
| max_vector_index_off_heap_usage         | Maximum Off-Heap Memory Utilization (%) of Vector Indexes | Maximum off-heap memory (in %) used for querying vector indexes on each node in a CSS cluster.<br><br>This metric is calculated only when vector index queries occur. | 0~100       | %    | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                           | Metric                              | Description                                                                                                                                                                                             | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------------|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| vector_index_circuit_breaker_status | Vector Index Circuit Breaker Status | Circuit breaker status for vector indexes. <ul style="list-style-type: none"><li>0: Not enabled.</li><li>1: Enabled. Check whether the cluster's off-heap memory usage exceeds the threshold.</li></ul> | 0, 1        | N/A   | N/A           | CSS cluster                  | 1 minute                  |
| request_count                       | Total Requests                      | Total number of requests sent to all nodes in a CSS cluster.                                                                                                                                            | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| successfully_request_count          | Successful requests                 | Number of requests successfully handled by all nodes in a CSS cluster.<br>Status code: 200.                                                                                                             | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| failed_request_count                | Failed requests                     | Number of failed requests for all nodes in a CSS cluster.<br>Status code: non-200.                                                                                                                      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                             | Metric                                                               | Description                                                                                                                                                                                                                | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|---------------------------------------|----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| limited_request_count                 | Throttled requests                                                   | Number of throttled requests for all nodes in a CSS cluster.<br><br>Status code: 429.                                                                                                                                      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| cold_data_storage                     | The size of cold data stored                                         | The total size of cold data stored in a CSS cluster, which is the total size of the data of frozen indexes stored in an OBS bucket. This metric is measured hourly. Users cannot check this OBS bucket on the OBS console. | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| number_of_index_creation_failures     | Number of index creation failures                                    | Number of index creation failures in a CSS cluster.                                                                                                                                                                        | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| shard_document_exceed_threshold_count | Number of shards where the number of documents exceeds the threshold | Number of shards where the number of documents exceeds the threshold in a CSS cluster.                                                                                                                                     | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                         | Metric                                              | Description                                                                                                                                            | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------------|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| write_fail_alias_error            | Number of Write Failures Due to Alias Exceptions    | Number of write failures caused by abnormal index aliases in a CSS cluster.<br><br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric.   | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| write_fail_documents_count_exceed | Number of Write Failures Due to Excessive Documents | Number of write failures caused by too many index documents in a CSS cluster.<br><br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| write_fail_fields_count_exceed    | Number of Write Failures Due to Excessive Fields    | Number of write failures caused by too many index fields in a CSS cluster.<br><br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                      | Metric                                                  | Description                                                                                                                                                                           | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|--------------------------------|---------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| write_fail_shards_count_exceed | Number of Write Failures Due to Excessive Shards        | Number of write failures caused by too many index shards in a CSS cluster.<br><br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric.                                   | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| write_fail_index_blocks_write  | Number of Write Failures Due to Index Write Prohibition | Number of write failures caused by the setting of write prohibition in index configuration in a CSS cluster.<br><br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| write_fail_mapping_malformed   | Number of Write Failures Due to Incorrect Field Types   | Number of write failures caused by incorrect index field types in a CSS cluster.<br><br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric.                             | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                                | Metric                                                                                        | Description                                                                                                                                                                                        | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|------------------------------------------|-----------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| remote_sync_task_count                   | Number of Primary/Secondary Synchronization Jobs                                              | Number of primary/secondary synchronization jobs for a CSS cluster.<br><br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric.                                                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| remote_sync_failed_task_count            | Number of Failed Primary/Secondary Synchronization Jobs                                       | Number of failed primary/secondary synchronization jobs for a CSS cluster.<br><br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric.                                                | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| remote_sync_diff_bytes_of_primary_shards | Size of Data to Be Synchronized via Primary/Secondary Synchronization Jobs for Primary Shards | Total size of data to be synchronized via primary/secondary synchronization jobs for all primary shards in a CSS cluster.<br><br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric. | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |

| Metric ID                                    | Metric                                                                                               | Description                                                                                                                                                                                           | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|----------------------------------------------|------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| remote_sync_finished_bytes_of_primary_shards | Size of Data Synchronized via Primary/Secondary Synchronization Jobs for Primary Shards              | Total size of data already synchronized via primary/secondary synchronization jobs for all primary shards in a CSS cluster.<br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric.      | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| remote_sync_diff_docs_of_primary_shards      | Number of Documents to Be Synchronized via Primary/Secondary Synchronization Jobs for Primary Shards | Total number of documents to be synchronized via primary/secondary synchronization jobs for all primary shards in a CSS cluster.<br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |



| Metric ID                  | Metric                                                                     | Description                                                                                                                                                                                  | Value Range | Unit | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|----------------------------|----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------|---------------|------------------------------|---------------------------|
| remote_sync_diff_bytes     | Size of Data to Be Synchronized via Primary/Secondary Synchronization Jobs | Total size of data to be synchronized via primary/secondary synchronization jobs for all shards in a CSS cluster.<br><br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric.   | $\geq 0$    | byte | 1024          | CSS cluster                  | 1 minute                  |
| remote_sync_finished_bytes | Size of Data Synchronized via Primary/Secondary Synchronization Jobs       | Total size of data already synchronized via primary/secondary synchronization jobs for all shards in a CSS cluster.<br><br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric. | $\geq 0$    | byte | 1024          | CSS cluster                  | 1 minute                  |

| Metric ID                | Metric                                                                            | Description                                                                                                                                                                                       | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|--------------------------|-----------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| remote_sync_diff_docs    | Number of Documents to Be Synchronized via Primary/Secondary Synchronization Jobs | Total number of documents to be synchronized via primary/secondary synchronization jobs for all shards in a CSS cluster.<br><br>Only Elasticsearch 7.6.2 and 7.10.2 clusters support this metric. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| index_storage            | Index Data Volume                                                                 | Total data size of all indexes in a CSS cluster.                                                                                                                                                  | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| backup_failed_task_count | Number of Failed Backup Tasks                                                     | Number of failed snapshot backup tasks in a CSS cluster.                                                                                                                                          | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| ism_fail_count           | Number of Failed Index Lifecycle Tasks                                            | Number of failed index lifecycle tasks in a CSS cluster.<br><br>Only Elasticsearch 7.10.2 clusters support this metric.                                                                           | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                   | Metric                            | Description                                                                          | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------|-----------------------------------|--------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| disk_total                  | Total Disk Space                  | Total data disk capacity of all nodes in a CSS cluster.                              | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| number_of_unassigned_shards | Number of Unallocated Shards      | Total number of unallocated shards in a CSS cluster.                                 | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| disk_iops_qos_num           | Disk IOPS Overload Times          | Total number of times the disk IOPS exceeds the predefined threshold.                | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| disk_iobw_qos_num           | Disk I/O Bandwidth Overload Times | Total number of times the disk I/O bandwidth usage exceeds the predefined threshold. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_disk_io_util            | Maximum Disk I/O Usage            | Maximum disk I/O usage among all disks.                                              | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| avg_disk_io_util            | Average Disk I/O Usage            | Average disk I/O usage among all disks.                                              | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID         | Metric            | Description                                                         | Value Range | Unit | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------|-------------------|---------------------------------------------------------------------|-------------|------|---------------|------------------------------|---------------------------|
| shard_quota_usage | Shard Quota Usage | Ratio of allocated shards to the upper limit of the shard quantity. | 0~100       | %    | N/A           | CS Cluster                   | 1 minute                  |

Node Monitoring Metrics

Table 2-168 CSS node monitoring metrics

| Metric ID      | Metric         | Description                      | Value Range | Unit | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|----------------|----------------|----------------------------------|-------------|------|---------------|---------------------------------|---------------------------|
| jvm_heap_usage | JVM Heap Usage | JVM heap memory usage of a node. | 0~100       | %    | N/A           | CS cluster - cloud service node | 1 minute                  |

| Metric ID    | Metric          | Description                                         | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|--------------|-----------------|-----------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| cpu_usage    | CPU Usage       | CPU usage.                                          | 0~100       | %     | N/A           | CS cluster - cloud service node | 1 minute                  |
| disk_util    | Max. Disk Usage | Disk usage of a node                                | 0~100       | %     | N/A           | CS cluster - cloud service node | 1 minute                  |
| load_average | Average Load    | Average number of queued tasks per minute on a node | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |

| Metric ID               | Metric                        | Description                                         | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|-------------------------|-------------------------------|-----------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| open_file_descriptors   | Open File Descriptors         | Number of opened file descriptors on a node         | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| max_file_descriptors    | Max. Allowed File Descriptors | Maximum number of allowed file descriptors          | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| thread_pool_write_queue | Tasks in Write Queue          | Number of tasks in queue for the write thread pool. | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |

| Metric ID                     | Metric                        | Description                                               | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|-------------------------------|-------------------------------|-----------------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| thread_pool_search_queue      | Tasks in Search Queue         | Number of tasks in queue for the search thread pool.      | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| thread_pool_force_merge_queue | Tasks in ForceMerge Queue     | Number of tasks in queue for the force merge thread pool. | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| thread_pool_write_rejected    | Rejected Tasks in Write Queue | Number of tasks rejected by the write thread pool.        | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |

| Metric ID                        | Metric                             | Description                                             | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|----------------------------------|------------------------------------|---------------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| thread_pool_search_rejected      | Rejected Tasks in Search Queue     | Number of tasks rejected by the search thread pool.     | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |
| thread_pool_force_merge_rejected | Rejected Tasks in ForceMerge Queue | Number of tasks rejected by the force merge thread pool | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |
| thread_pool_write_threads        | Size of Write Thread Pool          | Size of a write thread pool                             | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |



| Metric ID                       | Metric                         | Description                                         | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|---------------------------------|--------------------------------|-----------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| thread_pool_search_threads      | Size of Search Thread Pool     | Size of a search thread pool                        | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |
| thread_pool_force_merge_threads | Size of ForceMerge Thread Pool | Size of a force merge thread pool                   | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |
| thread_pool_flush_queue         | Tasks in Flush Queue           | Number of tasks in queue for the flush thread pool. | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |

| Metric ID                  | Metric                        | Description                                           | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|----------------------------|-------------------------------|-------------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| thread_pool_flush_rejected | Rejected Tasks in Flush Queue | Number of tasks rejected by the flush thread pool.    | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |
| thread_pool_flush_threads  | Size of Flush Thread Pool     | Size of the flush thread pool.                        | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |
| thread_pool_generic_queue  | Tasks in Generic Queue        | Number of tasks in queue for the generic thread pool. | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |

| Metric ID                    | Metric                          | Description                                              | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|------------------------------|---------------------------------|----------------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| thread_pool_generic_rejected | Rejected Tasks in Generic Queue | Number of tasks rejected by the generic thread pool.     | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |
| thread_pool_generic_threads  | Size of Generic Thread Pool     | Size of the generic thread pool.                         | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |
| thread_pool_management_queue | Tasks in Management Queue       | Number of tasks in queue for the management thread pool. | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |

| Metric ID                       | Metric                             | Description                                             | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|---------------------------------|------------------------------------|---------------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| thread_pool_management_rejected | Rejected Tasks in Management Queue | Number of tasks rejected by the management thread pool. | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| thread_pool_management_threads  | Size of Management Thread Pool     | Size of the management thread pool.                     | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| thread_pool_refresh_queue       | Tasks in Refresh Queue             | Number of tasks in queue for the refresh thread pool.   | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |

| Metric ID                      | Metric                          | Description                                                | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|--------------------------------|---------------------------------|------------------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| thread_pool_refresh_rejected   | Rejected Tasks in Refresh Queue | Number of tasks rejected by the refresh thread pool.       | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| thread_pool_refresh_threads    | Size of Refresh Thread Pool     | Size of the refresh thread pool.                           | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_searcher_queue | Tasks in OBS Searcher Queue     | Number of tasks in queue for the OBS Searcher thread pool. | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |

| Metric ID                         | Metric                               | Description                                               | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|-----------------------------------|--------------------------------------|-----------------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| thread_pool_obs_searcher_rejected | Rejected Tasks in OBS Searcher Queue | Number of tasks rejected by the OBS Searcher thread pool. | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |
| thread_pool_obs_searcher_threads  | Size of OBS Searcher Thread Pool     | Size of the OBS Searcher thread pool.                     | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |
| thread_pool_obs_queue             | Tasks in OBS Queue                   | Number of tasks in queue for the OBS thread pool.         | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |

| Metric ID                    | Metric                      | Description                                              | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|------------------------------|-----------------------------|----------------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| thread_pool_obs_rejected     | Rejected Tasks in OBS Queue | Number of tasks rejected by the OBS thread pool.         | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_threads      | Size of OBS Thread Pool     | Size of the OBS thread pool.                             | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_upload_queue | Tasks in OBS Upload Queue   | Number of tasks in queue for the OBS Upload thread pool. | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |

| Metric ID                       | Metric                             | Description                                                | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|---------------------------------|------------------------------------|------------------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| thread_pool_obs_upload_rejected | Rejected Tasks in OBS Upload Queue | Number of tasks rejected by the OBS Upload thread pool.    | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_upload_threads  | Size of OBS Upload Thread Pool     | Size of the OBS Upload thread pool.                        | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_download_queue  | Tasks in OBS Download Queue        | Number of tasks in queue for the OBS Download thread pool. | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |



| Metric ID                         | Metric                               | Description                                             | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|-----------------------------------|--------------------------------------|---------------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| thread_pool_obs_download_rejected | Rejected Tasks in OBS Download Queue | Number of tasks rejected by the OBS Upload thread pool. | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_download_threads  | Size of OBS Download Thread Pool     | Size of the OBS Download thread pool.                   | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| free_fs_size                      | Available Size of File Systems       | Available size of file systems in a CSS cluster.        | $\geq 0$    | byte  | 1024          | CSS cluster - cloud service node | 1 minute                  |

| Metric ID        | Metric                                  | Description                                       | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|------------------|-----------------------------------------|---------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| total_fs_size    | Total Size of File Systems              | Total size of file systems in a CSS cluster.      | $\geq 0$    | byte  | 1024          | CSS cluster - cloud service node | 1 minute                  |
| jvm_old_gc_count | Total GCs of Old-Generation JVM         | Number of old-generation garbage collection times | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| jvm_old_gc_time  | Total GC Duration of Old-Generation JVM | Time spent on old-generation garbage collection   | $\geq 0$    | ms    | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID          | Metric                              | Description                                         | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|--------------------|-------------------------------------|-----------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| jvm_young_gc_count | Total GCs of Young-Generation JVM   | Number of young-generation garbage collection times | $\geq 0$    | Count | N/A           | CS Cluster - cloud service node | 1 minute                  |
| jvm_young_gc_time  | GC Duration of Young-Generation JVM | Time spent on young-generation garbage collection   | $\geq 0$    | ms    | N/A           | CS Cluster - cloud service node | 1 minute                  |
| mem_free_in_bytes  | Available Memory                    | Unused memory space of a node.                      | $\geq 0$    | byte  | 1024          | CS Cluster - cloud service node | 1 minute                  |

| Metric ID                 | Metric                          | Description                                         | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|---------------------------|---------------------------------|-----------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| mem_free_percent          | Available Memory Percentage     | Percentage of unused memory space on a node.        | 0~100       | %     | N/A           | CS cluster - cloud service node | 1 minute                  |
| mem_used_in_bytes         | Used Memory                     | Used memory space of a node.                        | $\geq 0$    | byte  | 1024          | CS cluster - cloud service node | 1 minute                  |
| current_opened_http_count | Currently Open HTTP Connections | Number of HTTP connections currently open on a node | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |

| Metric ID                | Metric                      | Description                                                            | Value Range | Unit       | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|--------------------------|-----------------------------|------------------------------------------------------------------------|-------------|------------|---------------|----------------------------------|---------------------------|
| total_opend_http_count   | Total Open HTTP Connections | Total number of HTTP connections on a node                             | $\geq 0$    | Count      | N/A           | CSS cluster - cloud service node | 1 minute                  |
| disk_read_requests_rate  | Disk Read IOPS              | Number of disk read requests to a node in the CSS cluster per second.  | $\geq 0$    | requests/s | N/A           | CSS cluster - cloud service node | 1 minute                  |
| disk_write_requests_rate | Disk Write IOPS             | Number of disk write requests to a node in the CSS cluster per second. | $\geq 0$    | requests/s | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID             | Metric          | Description                                                                             | Value Range | Unit   | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|-----------------------|-----------------|-----------------------------------------------------------------------------------------|-------------|--------|---------------|----------------------------------|---------------------------|
| disk_read_bytes_rate  | Disk Read Rate  | Number of bytes of data read from the disks of a node in the CSS cluster per second.    | $\geq 0$    | Byte/s | 1024          | CSS cluster - cloud service node | 1 minute                  |
| disk_write_bytes_rate | Disk Write Rate | Number of bytes of data written into the disks of a node in the CSS cluster per second. | $\geq 0$    | Byte/s | 1024          | CSS cluster - cloud service node | 1 minute                  |
| shards_count          | Shards          | Number of shards per node in the CSS cluster.                                           | $\geq 0$    | Count  | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID                           | Metric                                            | Description                                                                         | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|-------------------------------------|---------------------------------------------------|-------------------------------------------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| vector_index_off_heap_used_in_bytes | Off-Heap Memory Usage of Vector Indexes           | Off-heap memory (bytes) used for querying vector indexes on nodes in a CSS cluster. | $\geq 0$    | byte  | 1024          | CSS cluster - cloud service node | 1 minute                  |
| vector_index_off_heap_usage         | Off-Heap Memory Utilization (%) of Vector Indexes | Off-heap memory (in %) used for querying vector indexes on nodes in a CSS cluster.  | 0~100       | %     | N/A           | CSS cluster - cloud service node | 1 minute                  |
| disk_iops_qos_num                   | Disk IOPS Overload Times                          | Total number of times the disk IOPS exceeds the predefined threshold.               | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID         | Metric                            | Description                                                                          | Value Range | Unit  | Number System | Monitored Object (Dimension)    | Monitoring Interval (Raw) |
|-------------------|-----------------------------------|--------------------------------------------------------------------------------------|-------------|-------|---------------|---------------------------------|---------------------------|
| disk_iobw_qos_num | Disk I/O Bandwidth Overload Times | Total number of times the disk I/O bandwidth usage exceeds the predefined threshold. | $\geq 0$    | Count | N/A           | CS cluster - cloud service node | 1 minute                  |
| max_disk_io_util  | Maximum Disk I/O Usage            | Maximum disk I/O usage among all disks.                                              | 0~100       | %     | N/A           | CS cluster - cloud service node | 1 minute                  |
| avg_disk_io_util  | Average Disk I/O Usage            | Average disk I/O usage among all disks.                                              | 0~100       | %     | N/A           | CS cluster - cloud service node | 1 minute                  |



| Metric ID         | Metric            | Description                                                         | Value Range | Unit | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|-------------------|-------------------|---------------------------------------------------------------------|-------------|------|---------------|----------------------------------|---------------------------|
| shard_quota_usage | Shard Quota Usage | Ratio of allocated shards to the upper limit of the shard quantity. | 0~100       | %    | N/A           | CSS cluster - cloud service node | 1 minute                  |

Dimensions

Table 2-169 Dimension description

| Key                    | Value                               |
|------------------------|-------------------------------------|
| cluster_id             | CSS cluster ID                      |
| cluster_id,instance_id | CSS cluster - cloud service node ID |

### 2.15.1.2 Elasticsearch Cluster Monitoring Metrics Supported by the Kernel

**Table 2-170** Elasticsearch cluster monitoring metrics supported by the kernel

| Metric                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                     | Supported Cluster Version                 | Details                                                                   |
|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|---------------------------------------------------------------------------|
| P99 latency for search requests   | Open-source Elasticsearch provides only the average latency metric when monitoring responses to search requests. This may not accurately reflect the actual search performance of a cluster. To improve on this, the P99 latency metric is added in CSS to monitor the 99th percentile latency of each cluster.                                                                                                                 | Elasticsearch 7.6.2, Elasticsearch 7.10.2 | <a href="#">Monitoring P99 Latency</a>                                    |
| HTTP status codes for the cluster | When you access Elasticsearch through HTTP, you receive HTTP status codes in response to your requests. The native open-source Elasticsearch does not collect statistics on these status codes. To improve on this, HTTP status code monitoring is added in CSS, allowing you to monitor HTTP status codes and get a sense of how the service is running.                                                                       | Elasticsearch 7.6.2, Elasticsearch 7.10.2 | <a href="#">Monitoring HTTP Status Codes</a>                              |
| Cluster indexes                   | Index monitoring helps to monitor the index usage and trends of a cluster, allowing users to handle potential risks in a timely manner to ensure cluster reliability. Index usage statistics are collected and stored in the cluster's monitoring index (index name: <b>monitoring-eye-css-[yyyy-mm-dd]</b> ). By default, these monitoring indexes are retained for one week before being permanently deleted upon expiration. | Elasticsearch 7.6.2, Elasticsearch 7.10.2 | <a href="#">Configuring Index Monitoring for an Elasticsearch Cluster</a> |

## 2.15.2 Configuring Elasticsearch Cluster Monitoring

### 2.15.2.1 Using Cloud Eye to Monitor Elasticsearch Clusters

With CSS, you can use Cloud Eye to monitor 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 Cluster Monitoring Information:** View the statistics of the monitoring metrics in specific periods.

#### Recommended Monitoring Metrics

- For routine monitoring of a cluster, you are advised to focus on disk usage and cluster health status.
- 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 search 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.

#### Prerequisites

- The cluster is in the **Available** or **Processing** status.
- The cluster has been running properly for more than 10 minutes.

#### 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 to filter alarm rules.

If no alarm rules are available, create one by referring to [Creating an Alarm Rule](#). [Table 2-171](#) describes the key parameters. For more parameters, customize them as needed.

**Table 2-171** Alarm rule configuration parameters

| Parameter      | Description                                                                                                                                                                                                                                                                   |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Alarm Type     | Select <b>Metric</b> .                                                                                                                                                                                                                                                        |
| Cloud Service  | Select <b>Cloud Search Service</b> .                                                                                                                                                                                                                                          |
| Resource Layer | Select the layer by which you are going to configure alarm rules. <ul style="list-style-type: none"><li>• <b>CSS Clusters</b>: Alarm rules are specified by cluster.</li><li>• <b>CSS Clusters - CSS Instances</b>: Alarm rules are specified by node in a cluster.</li></ul> |


## Configuring Monitoring Metrics

1. Create a monitoring dashboard by referring to [Creating a Dashboard](#). If an available monitoring dashboard has been created, skip this step.
2. Add CSS monitoring graphs by referring to [Adding a Graph](#).

## Viewing Cluster Monitoring Information

In the cluster list, check the monitoring information of clusters and cluster nodes.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and choose **More > Monitoring Metrics** in the **Operation** column.
4. On the **Monitoring Metrics** page, check the cluster and node metrics.

**Figure 2-69** Monitoring Metrics

- Click the **Instances** tab to check cluster metrics.
  - **Instances**: Select a cluster to check its metrics. The current cluster is selected by default. Multiple clusters can be selected so that you may compare between them.

- You can also set monitoring periods so that you can compare metrics between different periods for the same cluster.
  - Add Graph Group: You can add a custom group in addition to the default group. When you move the pointer over a group name on the left, the icons for changing the group name and deleting the group are displayed on the right. You can modify the group.
  - Select Metric: In the displayed **Select Metric** dialog box, set **All Metrics** and **Top N Metrics**.
- Click the **CSS Instances** tab to check the metrics of cluster nodes. Only the node monitoring information of the first cluster selected on the **Instances** tab is displayed.
- CSS Instances: Select the cluster nodes whose metrics and information you want to check. By default, the first node of the current cluster is selected. Multiple nodes can be selected so that you may compare between them.
  - You can also set monitoring periods so that you can compare metrics between different periods for the same cluster.
  - Add Graph Group: You can add a custom group in addition to the default group. When you move the pointer over a group name on the left, the icons for changing the group name and deleting the group are displayed on the right. You can modify the group.
  - Select Metric: In the displayed **Select Metric** dialog box, set **All Metrics**.

## 2.15.2.2 Configuring Kernel Monitoring for an Elasticsearch Cluster

### Scenario

Table 2-172 Introduction to cluster kernel monitoring

| Enhanced Monitoring Feature | Description                                                                                                                                                                                                                                                                                                     | Cluster Version                           | Details                       |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|-------------------------------|
| P99 latency                 | Open-source Elasticsearch provides only the average latency metric when monitoring responses to search requests. This may not accurately reflect the actual search performance of a cluster. To improve on this, the P99 latency metric is added in CSS to monitor the 99th percentile latency of each cluster. | Elasticsearch 7.6.2, Elasticsearch 7.10.2 | <b>Monitoring P99 Latency</b> |

| Enhanced Monitoring Feature | Description                                                                                                                                                                                                                                                                                                                                               | Cluster Version                           | Details                                      |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|----------------------------------------------|
| HTTP status codes           | When you access Elasticsearch through HTTP, you receive HTTP status codes in response to your requests. The native open-source Elasticsearch does not collect statistics on these status codes. To improve on this, HTTP status code monitoring is added in CSS, allowing you to monitor HTTP status codes and get a sense of how the service is running. | Elasticsearch 7.6.2, Elasticsearch 7.10.2 | <a href="#">Monitoring HTTP Status Codes</a> |

## Logging In to Kibana

Log in to Kibana and go to the command execution page. Elasticsearch clusters support multiple access methods. This topic uses Kibana as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
4. In the left navigation pane, choose **Dev Tools**.

## Monitoring P99 Latency

Run the following command to obtain the P99 latency of the current cluster:

```
GET /search/stats/percentile
```

An example output is as follows:

```
{
 "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
}
}
```

**Table 2-173** Response parameters

| Parameter    | Description                                                    |
|--------------|----------------------------------------------------------------|
| overall      | Statistics between cluster starting and the current time.      |
| last_one_day | Statistics for the most recent day.                            |
| latest       | Statistics from the most recent resetting to the current time. |

 **NOTE**

- The calculated P99 latency is an estimation, but it is more precise than the P50 latency.
- When a cluster is restarted, its P99 latency data is cleared, and is re-measured after the cluster restarts successfully.

The command used for monitoring the P99 latency of clusters can also be used to set other configuration items.

- You can customize the percentile of latency to be monitored.

For example, run the following command to show the P1, P50, and P90 latency values:

```
GET /search/stats/percentile
{
 "percents": [1, 50, 90]
}
```

- You can manually reset the **latest** statistics.

Run the following command to reset the **latest** statistics:

```
POST /search/stats/reset
```

If **ok** is returned, the reset is successful.

```
{
 "nodes" : {
 "css-c9c8-ess-esn-1-1" : "ok"
 }
}
```

## Monitoring HTTP Status Codes

The command used for monitoring HTTP status codes varies with cluster versions.

- In an Elasticsearch 7.6.2 cluster, run the following command to obtain statistics on HTTP status codes:

```
GET /_nodes/http_stats
```

Example response:

```
{
 "_nodes" : {
 "total" : 1,
 "successful" : 1,
 "failed" : 0
 },
 "cluster_name" : "css-8362",
 "nodes" : {
 "F9IFdQPARaOJI7oL7HOxtQ" : {
 "http_code" : {
 "200" : 114,
 "201" : 5,
 "429" : 0,
 "400" : 7,
 "404" : 0,
 "405" : 0
 }
 }
 }
}
```

- In an Elasticsearch 7.10.2 cluster, run the following command to obtain statistics on HTTP status codes:

```
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
 }
 }
 }
 }
}
```

### 2.15.2.3 Configuring Index Monitoring for an Elasticsearch Cluster

#### Scenario

Index monitoring helps to monitor the index usage and trends of a cluster, allowing users to handle potential risks in a timely manner to ensure cluster reliability. Index usage statistics are collected and stored in the cluster's monitoring index (index name: **monitoring-eye-css-[yyyy-mm-dd]**). By default, these monitoring indexes are retained for one week before being permanently deleted upon expiration.



## Constraints

- Only Elasticsearch 7.6.2 and 7.10.2 clusters support index monitoring.
- Indexes starting with **monitoring-eye-css-\*** are identified as monitoring indexes and will not be monitored. Do not use this prefix for regular indexes.
- The index pattern of **monitoring-eye-css-\*** must not be deleted while index monitoring is enabled. Otherwise, monitoring charts will become abnormal.

## Logging In to Kibana

Log in to Kibana and go to the command execution page. Elasticsearch clusters support multiple access methods. This topic uses Kibana as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
4. In the left navigation pane, choose **Dev Tools**.

## Enabling Index Monitoring

1. Run the following command to enable index monitoring:

```
PUT _cluster/settings
{
 "persistent": {
 "css.monitoring.index.enabled": "true"
 }
}
```

2. To monitor a single index, run the following command:

```
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 2-174** Configuration items

| Configuration Item            | Type    | Description                                                                                                                        |
|-------------------------------|---------|------------------------------------------------------------------------------------------------------------------------------------|
| css.monitoring.index.enabled  | Boolean | Whether to enable index monitoring. Setting this parameter to <b>true</b> enables index monitoring.<br>Default value: <b>false</b> |
| css.monitoring.index.interval | Time    | Index monitoring cycle.<br>Minimum value: <b>1s</b><br>Default value: <b>10s</b>                                                   |

| Configuration Item              | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| css.monitoring.index.indices    | String | <p>Name of an index to be monitored. By default, all indexes are monitored. You can configure specific indexes or a specific type of indexes to monitor.</p> <p>Example:</p> <ul style="list-style-type: none"><li>• <b>"css.monitoring.index.indices":</b> <b>["index_name"]</b> indicates only <i>index_name</i> is monitored.</li><li>• <b>"css.monitoring.index.indices": ["log_*"]</b> indicates that only indexes starting with <b>log_</b> are monitored.</li><li>• <b>"css.monitoring.index.indices": ["index1", "index2"]</b> indicates that <b>index1</b> and <b>index2</b> are monitored.</li></ul> <p>Default value: * (indicating that all indexes are monitored)</p> |
| css.monitoring.history.duration | Time   | <p>Retention period of index monitoring data. The default period is a week.</p> <p>Minimum value: <b>1d</b></p> <p>Default value: <b>7d</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## Checking the Read and Write Traffic of Indexes

After index monitoring is enabled for a cluster, you can check the read/write traffic of this cluster's indexes during specific periods of time.

- Run the following command to check the read and write traffic of all indexes:  
GET /\_cat/monitoring
- Run the following command to check the read and write traffic of a specified index:  
GET /\_cat/monitoring/{indexname}  
**{indexName}** indicates the name of the index whose read and write traffic you want to check.
- Run the following command to check the read and write traffic of an index during specified periods of time:  
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 2-175** Configuration items

| Configurat<br>ion Item | Mandator<br>y | Description                                                                                                                                                                           |
|------------------------|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| indexname              | No            | Index name.<br><b>NOTE</b><br>You cannot check the traffic of system indexes, whose names start with a period (.).                                                                    |
| begin                  | No            | Start time (UTC time) of the monitoring data you want to view.<br>Time format: strict_date_optional_time epoch_millis<br>The default start time is 5 minutes before the current time. |
| end                    | No            | End time (UTC time) of the monitoring data you want to view.<br>Time format: strict_date_optional_time epoch_millis<br>The default end time is the current time.                      |

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 2-176** Parameters in the returned information

| Parameter  | Description                                                               |
|------------|---------------------------------------------------------------------------|
| index      | Index name.                                                               |
| 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.           |


| Parameter      | Description                                                                  |
|----------------|------------------------------------------------------------------------------|
| 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. |

## Checking Index Monitoring Data in Kibana Charts

You can check preconfigured index monitoring charts on the **Dashboard** and **Visualizations** pages of Kibana. You can also customize tables and charts.

- Check index monitoring results in preconfigured dashboard charts.
  - In the navigation tree on the left, click **Dashboard**.
  - Click **[Monitoring] Index monitoring Dashboard** to check preconfigured dashboards.

**Figure 2-70** Preconfigured dashboard charts



The preconfigured dashboard displays the number of read and write operations per second in the cluster and the top 10 indexes that have the most read and write operations per second.

**Table 2-177** 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 space occupied by documents in a cluster.                     |
| [monitoring] indices store_size for top10                       | Top 10 indexes that occupy the largest storage space.                       |
| [monitoring] indices docs_count for top10                       | Top 10 indexes that store 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.        |

- Check index monitoring results in custom visualizations.

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 custom charts.

The following procedure describes how to check the trend of document quantities in a chart as an example.

- Click **Visualize** in the navigation tree on the left of the Kibana console.
- Click **Create visualization** and select **TSVB**.
- Set chart parameters and view the visualizations.

On the **Data** tab page, set the parameters as needed.

- Select **Max** for **Aggregation**, and select **index\_stats primaries.docs.count** in **Field**, indicating the number of documents in a primary shard.
- Select **Derivative** from **Aggregation** to indicate differences between aggregation buckets. Set **Units** to **1s** to visualize network rates as "per second".
- Set **Aggregation** to **Positive Only** to prevent negative numbers after resetting.
- To show statistics by index, set **Group by** to **Terms** and **By** to **index\_stats.index**. Statistics will be grouped by index name.
- 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 2-71 TSVB page





Figure 2-72 Setting the interval




If the index monitoring charts are not displayed, load them on the Kibana console again.

#### NOTE

If the preset dashboards and visualizations cannot be found for a security-mode Elasticsearch cluster, try switching to the Private or Global space. If the issue persists, import the charts and graphs again.


- Create the **monitoring-kibana.ndjson** file by referring to [kibana-monitor Configuration File](#).
- On the Kibana console, choose **Management > Stack Management > Saved objects**.

Figure 2-73 Selecting Saved Objects




3. Click **Import** and upload the **monitoring-kibana.ndjson** file created in 1.

Figure 2-74 Uploading a file



4. After the file is uploaded, click **done**. The index monitoring charts are imported successfully.

Figure 2-75 Index monitoring charts imported successfully



kibana-monitor Configuration File

The content of **kibana-monitor** configuration file 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": "metrics",
 "aggs": [],
 "params": {
 "id": "61ca57f0-469d-11e7-af02-69e470af7417",
 "type": "timeseries",
 "series": [
 {
 "id": "61ca57f1-469d-11e7-af02-69e470af7417",
 "color": "#6BC000",
 "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
\","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":{"aggs":[],"params":{"axis_formatter\":"number\","axis_position\":"left\","axis_scale\":"normal\","default_index_pattern
\":"monitoring-eye-css-*\","default_timefield\":"timestamp\","filter":{"language\":"kuery\","query
\":"\"},"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
\":"index_stats.total.segments.count\","id\":"61ca57f2-469d-11e7-af02-69e470af7417\","type\":"max
"}],"point_size\":"1\","separate_axis\":"0\","split_color_mode\":"kibana\","split_model\":"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\n\nThis dashboard contains default table for you to play with. You can view 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_model\":"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_model\":"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\":\"#68BC00\",\"split_model\":\"everything\",\"split_color_model\":\"kibana\",\"metrics\":{\"unit\":\"bytes\",\"id\":\"61ca57f2-469d-11e7-af02-69e470af7417\",\"type\":\"max\",\"field\":\"indices_stats_all.total.store.size_in_bytes\"},\"separate_axis\":0,\"axis_position\":\"right\",\"formatter\":\"number\",\"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_model\":\"show_all\",\"default_index_pattern\":\"monitoring-eye-css-*\",\"default_timefield\":\"timestamp\",\"isModelInvalid\":false,\"legend_position\":\"bottom\"}}},\"id\":\"c7f72ae0-664e-11e7-af02-69e470af7417\",\"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\":\"timeseries\",\"series\":{\"id\":\"61ca57f1-469d-11e7-af02-69e470af7417\",\"color\":\"#68BC00\",\"split_model\":\"everything\",\"split_color_model\":\"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_model\":\"show_all\",\"default_index_pattern\":\"monitoring-eye-css-*\",\"default_timefield\":\"timestamp\",\"isModelInvalid\":false,\"drop_last_bucket\":1,\"legend_position\":\"bottom\"}}},\"id\":\"c7f72ae0-664e-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\":\"timeseries\",\"series\":{\"id\":\"61ca57f1-469d-11e7-af02-69e470af7417\",\"color\":\"#68BC00\",\"split_model\":\"everything\",\"split_color_model\":\"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_model\":\"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\"}}},\"id\":\"c7f72ae0-664e-11e7-af02-69e470af7417\",\"references\":[]},\"type\":\"visualization\",\"updated_at\":\"2022-12-01T12:41:01.165Z\",\"version\":\"WzlxMywyXQ==\"},{\"attributes\":{\"description\":\"\",\"kibanaSavedObjectMeta\":{\"searchSourceJSON\":\"\"},\"title\":\"[monitoring] indexing rate of index for top10(/s)\",\"uiStateJSON\":\"\"},\"version\":1,\"visState\":{\"title\":\"[monitoring] indexing rate of index for top10(/s)\",\"type\":\"timeseries\",\"aggs\":{\"params\":{\"id\":\"61ca57f0-469d-11e7-af02-69e470af7417\",\"type\":\"timeseries\",\"series\":{\"id\":\"61ca57f1-469d-11e7-af02-69e470af7417\",\"color\":\"#68BC00\",\"split_model\":\"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\"},\"separate_axis\":0,\"axis_position\":\"right\",\"formatter\":\"indexing_rate\",\"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_model\":\"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_model\":\"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\":\"timeseries\",\"aggs\":{\"params\":{\"id\":\"61ca57f0-469d-11e7-af02-69e470af7417\",\"type\":\"timeseries\",\"series\":{\"id\":\"61ca57f1-469d-11e7-af02-69e470af7417\",\"color\":\"#68BC00\",\"split_model\":\"terms\",\"metrics\":{\"id\":\"61ca57f2-469d-11e7-af02-69e470af7417\",\"type\":\"max\",\"field\"
```

```

{"index_stats.total.search.query_total":{"unit":"1s","id":"fd9dfad0-a5ef-11ec-aa10-992297d21a2e","type":"derivative","field":"61ca57f2-469d-11e7-af02-69e470af7417","unit":"","id":"0aaa26a0-a5f0-11ec-aa10-992297d21a2e","type":"positive_only","field":"fd9dfad0-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":"","timeseries":"","terms_field":"index_stats.index","terms_order_by":"61ca57f2-469d-11e7-af02-69e470af7417","split_color_model":"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_model":"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_model":"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_model":"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_model":"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_model":"kibana","split_model":"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_model":"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\n"},\ "panelIndex":\ "80ece867-cf23-4935-bfbc-430afa51bcca",\ "embeddableConfig":{\ },\ "panelRefName\n":\ "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\n"},\ "embeddableConfig":{\ },\ "panelRefName":\ "panel_11"},{\ "gridData":{\ "x":0,\ "y":68,\ "w":24,\ "h\n":11,\ "i":\ "dd14182d-d8b9-47f2-bf36-6cba3b09586c"},\ "panelIndex":\ "dd14182d-d8b9-47f2-bf36-6cba3b09586c",\ "embeddableConfig":{\ },\ "panelRefName":\ "panel_12"},{\ "gridData":{\ "x":24,\ "y\n":68,\ "w":24,\ "h":11,\ "i":\ "a47f9333-52b7-49b7-8cac-f470cf405131"},\ "panelIndex\n":\ "a47f9333-52b7-49b7-8cac-f470cf405131",\ "embeddableConfig":{\ },\ "panelRefName\n":\ "panel_13"}},\ "timeRestore":false,\ "title":\ "[Monitoring] Index monitoring\nDashboard",\ "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":\ "dashboar\n",\ "updated_at":\ "2022-12-01T12:41:01.165Z",\ "version":\ "WzlxOCwyXQ=="}\ {"exportedCount":16,\ "missingRefCount":0,\ "missingReferences":[]}]
```

## 2.15.3 Setting Alarm Alerting via SMN for an Elasticsearch Cluster

This topic describes how to use the open-source Open Distro alarm plugin to configure alarm alerting via the Simple Message Notification (SMN) service for Elasticsearch clusters in Kibana.

### Scenario Description

By default, CSS installs the open-source Open Distro alarm plugin (opendistro\_alerting) for Elasticsearch clusters of versions 7.1.1, 7.6.2 and 7.10.2. This plugin has three components: Dashboard, Monitors, and Destinations. You can configure alarm alerting via SMN using the Destinations component. When using the Open Distro alarm plugin, an Elasticsearch cluster needs to send notifications via SMN. This requires service authorization, the purpose of which is to allow CSS to access other cloud resources, such as SMN, through agency, so that alarms generated for Elasticsearch clusters can be sent via SMN notifications. For details about the Open Distro alarm plugin, see [Open Distro Monitors](#).

### Constraints

By default, the open-source Open Distro alarm plugin is installed only in Elasticsearch 7.1.1, Elasticsearch 7.6.2 and Elasticsearch 7.10.2. Therefore, only these clusters support this feature.

### Prerequisites

- You have created a topic on the SMN console. For details, see [Creating a Topic](#).

- You have obtained CSS administrator account. The account has the permissions to access CSS and view the agency list, create agencies, and grant permissions to agencies.

## Authorizing Users to Use SMN

1. Log in to the [CSS management console](#).  
You must log in using a CSS administrator account.
2. In the navigation pane, choose **Service Authorization**.
3. On the **Service Authorization** page, click **Create Agency**. In the dialog box displayed, confirm that the agency is successfully created.
  - If an agency has been created, "css\_smn\_agency exist, no need to created." is displayed in the upper right corner.
  - If you do not have the permission to create an agency, an error message will be displayed in the upper right corner indicating "no permission", in which case, check that the administrator account has been assigned the IAM permission.

## Setting Alarm Notifications via SMN

1. Log in to the [CSS management console](#).  
Log in using an account with CSS permissions.
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.
4. On the Kibana page, choose **Open Distro for Elasticsearch > Alerting** in the navigation pane on the left.
5. Create an SMN destination to send alert messages.
  - a. On the **Alerting** page, click the **Destinations** tab and click **Add destination** to configure destination information.

**Table 2-178** Destinations parameter description

| Parameter | Description                                                                                                                                                                                                                                                                           |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name      | User-defined destination name                                                                                                                                                                                                                                                         |
| Type      | Retain the default value <b>SMN</b> .                                                                                                                                                                                                                                                 |
| Topic     | Select the SMN topic you have created in <a href="#">Prerequisites</a> for sending alarm messages.<br><b>CAUTION</b><br>For the Elasticsearch cluster of version 7.1.1, you need to manually enter the topic name. Ensure that the topic name is the same as that in the SMN service. |

Figure 2-76 Add destination

Alerting / Destinations / Create destination

Add destination

Destination

Name

test

Specify a name of the destination.

Type

SMN

Settings

Topic

test

- b. Click **Create** to return to the destination list. If the created SMN destination is displayed in the list, the creation is complete.

Figure 2-77 Destination list

|                                 |                  |             |
|---------------------------------|------------------|-------------|
| Dashboard Monitors Destinations |                  |             |
| Destinations                    |                  |             |
| Search                          |                  | All type 1  |
| Destination name                | Destination type | Actions     |
| test                            | SMN              | Edit Delete |
| Rows per page: 20               |                  | 1           |

6. Create a monitoring task and configure the alarm triggering condition and monitoring frequency.
- a. Click the **Monitors** tab on the **Alerting** page and click **Create monitors** to configure monitoring information.

Table 2-179 Monitor parameters

| Parameter    | Description               |
|--------------|---------------------------|
| Monitor name | User-defined monitor name |

| Parameter            | Description                                                                                                                                                                                                                                                                                         |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Monitor state        | Monitoring status. You are advised to keep this function enabled.                                                                                                                                                                                                                                   |
| Method of definition | Select a method to define monitoring. You are advised to use <b>Define using extraction query</b> . <ul style="list-style-type: none"><li>• <b>Define using visual graph</b>: use visualized query statement</li><li>• <b>Define using extraction query</b>: use specific query statement</li></ul> |
| Index                | Index to be monitored                                                                                                                                                                                                                                                                               |
| Time field           | When <b>Define using visual graph</b> is selected, select a time field and define counting parameters such as <b>count</b> .                                                                                                                                                                        |
| Frequency            | Select the monitoring frequency and set the monitoring interval. The options include: <ul style="list-style-type: none"><li>• <b>By interval</b></li><li>• <b>Daily</b></li><li>• <b>Weekly</b></li><li>• <b>Monthly</b></li><li>• <b>Custom cron expression</b></li></ul>                          |

- b. Click **Create**. The **Create trigger** page is displayed.
- c. On the **Create trigger** page, set the alarm triggering conditions and actions to be triggered.

**Table 2-180** Trigger parameters

| Parameter         | Description                                                                                                                                             |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| Trigger name      | User-defined trigger name                                                                                                                               |
| Severity level    | Sensitivity of a trigger, that is, the number of alarms that are triggered before an alarm message is sent. <b>1</b> indicates the highest sensitivity. |
| Trigger condition | Trigger condition. An alarm is triggered when the trigger condition is hit.                                                                             |
| Action name       | Name of a trigger action                                                                                                                                |
| Destination       | Select the SMN destination created in section <a href="#">5</a> .                                                                                       |
| Message subject   | Title of the alarm message. This parameter is required only when Elasticsearch clusters of version 7.10.2 is used.                                      |

| Parameter         | Description                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Message           | Body of an alarm message. By default, the subject and body are defined when the destination is an email address. For details, see <a href="#">Message Publishing</a> .                                                                                                                                                                                                                  |
| Action throttling | Message sending frequency. It limits the number of notification messages can be received in a specified period.<br><br>For example, if this parameter is set to 10 minutes, SMN sends only one alarm notification in the next 10 minutes even if the trigger condition is hit for multiple times. After 10 minutes, SMN sends another alarm notification if the alarm condition is met. |

**Figure 2-78** Setting the destination of a trigger action

## Configure actions

### ▼ SMN notification: test-action

Action name

test-action

Names can only contain letters, numbers, and special characters

Destination

test - (SMN)



- d. Click **Send test message**. If a subscriber receives an email, as shown in [Figure 2-80](#), the trigger is configured successfully.

**Figure 2-79** Sending test messages

Destination  
test - (SMN)

Message subject  
test

Message [Info](#)


```
{
 "subject": "subject",
 "message": "Monitor {{ctx.monitor.name}} just entered alert status. Please investigate the issue. - Trigger: {{ctx.trigger.name}} - Severity: {{ctx.trigger.severity}} - Period start: {{ctx.periodStart}} - Period end: {{ctx.periodEnd}}",
 "time_to_live": "3600"
}
```

Embed variables in your message using Mustache templates. [Learn more about Mustache](#)

[Send test message](#)

Message preview

```
{
 "subject": "subject",
 "message": "Monitor test-monitor just entered alert status. Please investigate the issue. - Trigger: test-trigger - Severity: 1 - Period start: 2023-04-14T02:16:36Z - Period end: 2023-04-14T02:17:36Z",
 "time_to_live": "3600"
}
```

**Figure 2-80** Email notification

- e. Click **Create** to return to the Monitor details page.

## 2.15.4 Synchronizing Elasticsearch Alerts to Prometheus

Make necessary configurations in Kibana to synchronize Elasticsearch alerts to Prometheus so you can use Prometheus to monitor and analyze key performance metrics for Elasticsearch clusters in real time.

Prometheus is an open-source monitoring system with a dimensional data model, flexible query language, efficient time series database and modern alerting approach.

### Constraints

Only Elasticsearch 7.10.2 clusters (with an image version no earlier than 7.10.2\_24.3.3\_x.x.x) support alert synchronization to Prometheus.

### Prerequisites

- The Prometheus monitor server is ready, and the Pushgateway address has been obtained. Prometheus and Elasticsearch must be connected. Otherwise, alerts cannot be sent.
- The Elasticsearch cluster is available.

### Configuring Alert Synchronization

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, find the target cluster, and click **Kibana** in the **Operation** column to log in to the Kibana console.



4. On the Kibana page, choose **Open Distro for Elasticsearch > Alerting** in the navigation pane on the left.
5. Create a Prometheus destination to send alert messages.
  - a. On the **Alerting** page, click the **Destinations** tab, and click **Add destination** to configure destination information.

**Table 2-181** Destinations parameters

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name      | User-defined destination name                                                                                                                                                                                                                                                                                                                                                                             |
| Type      | Select <b>PROMETHEUS</b> .                                                                                                                                                                                                                                                                                                                                                                                |
| Settings  | <p>Enter the Pushgateway address of the Prometheus monitor server.</p> <ul style="list-style-type: none"><li>• Currently, only Prometheus Gauge dashboards can be added or deleted. Metrics are queried using specific statements and numeric values are synchronized to Pushgateway for monitoring via Prometheus.</li><li>• Two types of Pushgateway addresses are supported: HTTP and HTTPS.</li></ul> |

Figure 2-81 Add destination

# Add destination

## Destination

**Name**

test

Specify a name of the destination.

**Type**

PROMETHEUS

## Settings

**Pushgateway endpoint, such as http://ip:9091**

http://\*.\*.\*:9091

- b. Click **Create** to return to the Destinations list. The created Prometheus destination is displayed in the list.

Figure 2-82 Destinations list

# Destinations

Q Search

| Destination name ↓ | Destination type |
|--------------------|------------------|
| test               | PROMETHEUS       |

6. Create a monitoring task to configure the alert triggering conditions and monitoring interval.
- a. Click the **Monitors** tab on the **Alerting** page, and click **Create monitors** to configure monitoring information.

**Table 2-182** Monitor parameters

| Parameter            | Description                                                                                                                                                                                                                                                                                         |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Monitor name         | User-defined monitor name                                                                                                                                                                                                                                                                           |
| Monitor state        | Monitoring status. You are advised to keep this function enabled.                                                                                                                                                                                                                                   |
| Method of definition | Select a method to define monitoring. You are advised to use <b>Define using extraction query</b> . <ul style="list-style-type: none"><li>• <b>Define using visual graph</b>: use visualized query statement</li><li>• <b>Define using extraction query</b>: use specific query statement</li></ul> |
| Index                | Index to be monitored                                                                                                                                                                                                                                                                               |
| Time field           | When <b>Define using visual graph</b> is selected, select a time field and define counting parameters such as <b>count</b> .                                                                                                                                                                        |
| Frequency            | Select the monitoring frequency and set the monitoring interval. The options include: <ul style="list-style-type: none"><li>• <b>By interval</b></li><li>• <b>Daily</b></li><li>• <b>Weekly</b></li><li>• <b>Monthly</b></li><li>• <b>Custom cron expression</b></li></ul>                          |

- b. Click **Create**. The **Create trigger** page is displayed.
- c. On the **Create trigger** page, set the alert triggering conditions and the actions to be triggered.

**Table 2-183** Trigger parameters

| Parameter      |                | Description                                                                                                                                             |
|----------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| Define trigger | Trigger name   | User-defined trigger name.                                                                                                                              |
|                | Severity level | Sensitivity of the trigger, that is, how many alerts need to be triggered before an alert is actually sent. <b>1</b> indicates the highest sensitivity. |

| Parameter         |                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                   | Trigger condition | Trigger condition. An alert is triggered when the trigger condition is met.<br><b>NOTE</b><br>You are advised to set a trigger condition that can almost always be triggered so that the queried metrics will always be synchronized to the Pushgateway.                                                                                                                                                                                                                                       |
| Configure actions | Action name       | Name of the triggered action.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                   | Destination       | Select the destination created in <a href="#">5</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|                   | Message           | Defines the body of the message to be published, which must use the JSON format. See the following for an example. <pre>{   "metricsName":"hits_total_value", //Prometheus metric name   "metricsLabel":   {"label_key1":"label_value1","label_key2":"label_value2"}, //Prometheus labels   "metricsValue":{{ctx.results.0.hits.total.value}}, //Prometheus metric values   "jobName":"job_name" //Prometheus monitor task name   "metricsHelp":"****" //Metric explanation. Optional. }</pre> |
|                   | Action throttling | Message sending interval. It limits the number of notification messages that can be sent over a specified period.<br><br>For example, if this parameter is set to 10 minutes, Prometheus sends only one alert notification in the next 10 minutes even if the trigger condition is met multiple times. After 10 minutes, Prometheus sends another alert notification if the trigger condition is met again.                                                                                    |

- d. Click **Send test message** to send a test message to Prometheus to check whether the trigger is set successfully.

Figure 2-83 Sending a test message

Configure actions

Add action

▼ PROMETHEUS notification: test

Delete

Action name

test

Names can only contain letters, numbers, and special characters

Destination

test - (PROMETHEUS)

Message

Info

```
{
 "metricsName": "hits_total_value",
 "metricsLabel": ("label_key1":"label_value1","label_key2":"label_value2"),
 "metricsValue": {{ctx.results.0.hits.total.value}},
 "jobName": "job_name"
}
```

Embed variables in your message using Mustache templates. [Learn more about Mustache](#)

Send test message

Message preview

```
{
 "metricsName": "hits_total_value",
 "metricsLabel": ("label_key1":"label_value1","label_key2":"label_value2"),
 "metricsValue": 1,
 "jobName": "job_name"
}
```

Action throttling

Info

☐ Enable action throttling

As shown in [Figure 2-84](#), Prometheus can receive a triggered message, meaning the trigger is set successfully.

Figure 2-84 Message received successfully

Pushgateway Metrics Status Help

job="job\_name" label\_key1="label\_value1" label\_key2="label\_value2"

Delete Group

hits\_total\_value

CSS PushGateway

GAUGE

last pushed: 2024-12-24T10:40:12+08:00

| Labels                                                                         | Value |
|--------------------------------------------------------------------------------|-------|
| instance="" job="job_name" label_key1="label_value1" label_key2="label_value2" | 1     |

e. Click **Create** to return to the **Monitor** details page.

## 2.15.5 Intelligent Risk Detection for Elasticsearch Clusters

CSS provides intelligent O&M to help detect potential risks for clusters, along with risk handling suggestions.

### Scenario

Intelligent O&M for clusters supports the following functions:

- **Creating a Scan Task:** Start a scan task to trigger an intelligent health check and diagnosis on the current cluster.
- **Checking the Risk Items of a Cluster:** After a scan task is completed, check the risks identified by the task, confirm and handle these risks in a timely manner based on risk handling suggestions.
- **Deleting a Scan Task:** Delete scan tasks that you no longer need. After a scan task is deleted, the system deletes all diagnoses generated by it.

The check items of intelligent O&M are as follows:

- Check the current health status of a cluster. Red: Some primary shards are not allocated. Yellow: Some replica shards are not allocated. Green: All shards have been allocated.
- Check the number of nodes in a cluster and the number of AZs to evaluate the high availability status of a distributed Elasticsearch cluster.
- Check whether index replicas are enabled. An index with no replicas may become unavailable in the case of a node failure. If local disks are used, this may even lead to data loss.
- Check for Kibana index conflicts in a cluster.
- Check disk usage. If the disk usage of a node is too high, new index shards may fail to be allocated to the node and the cluster performance may be affected.
- Check whether the storage usage of a cluster's data nodes or cold data nodes is balanced. Unbalanced storage distribution may result in unbalanced cluster loads and increase read/write latency.
- Monitor node disconnection and unavailability in a cluster for 5 consecutive minutes each time.
- Check for nodes with too many shards. An excessively large number of shards will consume excessive node resources, increasing read/write latency and slowing down metadata update.
- Check the sizes of all shards. A large shard may impact performance, occupy excessive node memory, and slow down shard restoration during cluster scaling or fault recovery.
- Check for any new version that is now available.
- Check for snapshot creation failures or the absence of any snapshot creation records in the last seven days.

## Granting Access to SMN

To send alarm notifications via SMN after a scan task is completed, you must first be granted access to the SMN service. Additionally, you must create a topic on the SMN console in advance. For details, see [Creating a Topic](#).

1. Log in to the [CSS management console](#).  
You must log in using a CSS administrator account.
2. In the navigation pane, choose **Service Authorization**.
3. On the **Service Authorization** page, click **Create Agency**. In the dialog box displayed, confirm that the agency is successfully created.
  - If an agency has been created, "css\_smn\_agency exist, no need to created." is displayed in the upper right corner.
  - If you do not have the permission to create an agency, an error message will be displayed in the upper right corner indicating "no permission", in which case, check that the administrator account has been assigned the IAM permission.

## Creating a Scan Task

1. Log in to the [CSS management console](#).

2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Monitoring > Intelligent O&M**.
5. On the **Intelligent O&M** tab, click **Scan** in the upper-left corner. In the displayed dialog box, configure the scan task.

**Table 2-184** Configuring a scan task

| Parameter                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                                       | Custom name of a scan task.<br>Enter 4 to 64 characters and start with a lowercase letter. Only lowercase letters, digits, hyphens (-), and underscores (_) are allowed.                                                                                                                                                                                                                                                                                |
| Description                                | Brief description of the task.                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Send SMN notification upon task completion | Whether to send SMN notifications upon completion of the scan. <ul style="list-style-type: none"><li>• When selected, you need to further configure <b>SMN Topic</b> and <b>Notification Level</b>. Upon completion of the scan, an SMN notification is sent if a risk higher than or equal to the notification level you set is detected.</li><li>• When deselected (default), no SMN notification will be sent upon completion of the scan.</li></ul> |
| SMN Topic                                  | If you select <b>Send SMN notification upon task completion</b> , you need to set the SMN topic.                                                                                                                                                                                                                                                                                                                                                        |
| Notification Level                         | If you select <b>Send SMN notification upon task completion</b> , you need to set the risk level.<br>If the scan result contains risks at this level or higher, SMN will send notifications that list all the risk items in the result.                                                                                                                                                                                                                 |

6. After the configuration is complete, click **OK**.
7. When the status of the scan task changes to **Completed**, you can check the cluster's risk items.

## Checking the Risk Items of a Cluster

1. On the **Intelligent O&M** page, select a completed scan task, and click  on the lower-right left to check the task details, including the task creation time, summary, ID, and the risk items found by it.

Figure 2-85 Checking scan task details

| Name/ID          | Description | Task Status | Risk Item                                              | Created              | SMN Notificati... | Operation          |
|------------------|-------------|-------------|--------------------------------------------------------|----------------------|-------------------|--------------------|
| aiops-8pa3       | --          | Completed   | 1 high-risk items<br>1 medium-risk items<br>1 warnings | Aug 05, 2025 15:3... | Disabled          | Delete Export Risk |
| Total Records: 1 |             |             |                                                        |                      |                   |                    |

Selected: aiops-8pa3

2. Click  on the left of a risk item to check its details, including the check item, risk description, and risk handling suggestion.
- You can handle cluster risks in a timely manner based on the suggestions provided by the system.

Figure 2-86 Risk items

Selected: aiops-8pa3

|         |                                                                           |    |            |
|---------|---------------------------------------------------------------------------|----|------------|
| Created | Aug 05, 2025 15:38:29 GMT+08:00                                           | ID | 7f16e1c0db |
| Summary | Analysis completed. High-risk items: 1; medium-risk items: 1; warnings: 1 |    |            |

High-Risk Items(1)

|                  |                                                                                                                                                                                                                                                                                                              |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Check Item       | Check the number of nodes in the cluster and the number of Azs(Available Zones) to evaluate the high availability status of the distributed aiops-8pa3 cluster.                                                                                                                                              |
| Risk Description | The current cluster has one or two nodes. If a node is faulty, the entire cluster may become unavailable. The service availability risk is high.                                                                                                                                                             |
| Suggestion       | You are advised to change the cluster to a multi-AZ cluster. Procedure: On the CSS cluster console, choose Clusters > aiops-8pa3. In the Operation column of a cluster, choose More > Modify Configuration. Click the Change AZ tab and add AZs. Click the Scale Cluster tab and change the number of nodes. |

Medium-Risk Items(1)

3. Select a scan task, and click **Export Risk** in the **Operation** column to download the scan result.

Deleting a Scan Task

You can delete scan tasks that are no longer needed. Deleting a scan task also deletes all its diagnoses and scan report.

1. On the **Intelligent O&M** page, locate the scan task you want to delete, and click **Delete** in the **Operation** column.
2. In the displayed dialog box, enter **DELETE**, and click **OK** to delete the task.

2.15.6 Managing Elasticsearch Cluster Logs

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

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



- **Log ingestion:** Cluster logs are transmitted in real time to the current cluster or another cluster on the same network. (Relevant version and network compatibility requirements must be met.) You can use Kibana or other visualization tools for log search and analysis. You can also set index prefixes and retention periods to flexibly manage the log lifecycle. Logs can be transferred across clusters, facilitating centralized monitoring in a multi-cluster environment.

## Impact on Billing

When log backup is enabled, the generated log backups are stored in OBS buckets, which will result in additional costs. For details, see [Object Storage Service Billing](#).

## Prerequisites

The OBS bucket used for storing log backups has been created. The OBS bucket must meet the following requirements:

- **Storage Class:** **Standard**.
- **Region:** the same as that of the cluster.

## Querying Logs

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Logs > Log Search**. The **Log Search** page is displayed.

You can search records by log type, node, log level, or keyword. For a detailed description of each type of logs, see [Log Introduction](#).

When a log file reaches 128 MB or when the time reaches 00:00 UTC, the system automatically compresses and archives it. Only unarchived logs appear on the log search page, while archived logs remain accessible through the log backup function.

## Backing Up Logs

Cluster logs can be backed up to OBS buckets, where you can download them for in-depth analysis at any time.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Logs > Log Backup**. The **Log Backup** page is displayed.
5. Enable log backup.

Perform the following steps to enable log backup. If it is already enabled, skip this step.

- a. Click **Enable Backup**. In the displayed dialog box, configure necessary settings.

**Table 2-185** Log backup settings

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                            |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBS Bucket  | Select an OBS bucket for storing log backups from the drop-down list box.<br>If no OBS buckets meet your requirements, click <b>Create Bucket</b> to go to the OBS console and create one. For details, see <a href="#">Creating a Bucket</a> .                                                                                                                        |
| Backup Path | Set the log storage location in the OBS bucket.<br>The backup path cannot: <ul style="list-style-type: none"><li>• Contain the following characters: \:*?"&lt;&gt; '{}.</li><li>• Start with a slash (/).</li><li>• Start or end with a period (.).</li><li>• Contain more than two consecutive slashes (/) or periods (.).</li><li>• Exceed 512 characters.</li></ul> |


| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IAM Agency | <p>To back up data to an OBS bucket, you must have the write permission to it. By configuring an IAM agency, you can authorize CSS to access its OBS resources through an associated account.</p> <ul style="list-style-type: none"><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-obs-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the OBS Administrator permissions deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/><pre>"obs:bucket:GetBucketLocation",<br/>"obs:object:GetObjectVersion",<br/>"obs:object:GetObject",<br/>"obs:object:DeleteObject",<br/>"obs:bucket:HeadBucket",<br/>"obs:bucket:GetBucketStoragePolicy",<br/>"obs:object:DeleteObjectVersion",<br/>"obs:bucket:ListBucketVersions",<br/>"obs:bucket:ListBucket",<br/>"obs:object:PutObject"</pre></li><li>When OBS buckets use SSE-KMS encryption, the IAM agency must be granted KMS permissions. You can click <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b> to have the following custom policies created automatically.<br/><pre>"kms:cmk:create",<br/>"kms:dek:create",<br/>"kms:cmk:get",<br/>"kms:dek:decrypt",<br/>"kms:cmk:list"</pre></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</pre></li><li>To use an IAM agency, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</pre></li></ul> |

| Parameter | Description                                                                                                                                             |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
|           | <b>WARNING</b><br>The agency name can contain only letters (case-sensitive), digits, underscores (_), and hyphens (-). Otherwise, the backup will fail. |

Determine whether to enable **Automatic backup** based on service needs.

- To enable automatic, periodic log backup, select **Automatic backup**, and set the backup policy. For more information, see [Table 2-186](#).
  - If manual log backup is sufficient, deselect **Automatic backup**.
- b. Click **OK** to enable log backup.
- The configuration information will be displayed on the **Log Backup** tab.
6. Back up logs. Two options are available: automatic or manual.
- Automatic log backup: Logs are backed up periodically based on the preset policy.
- a. On the **Log Backup** page, click **Modify Settings** on the right.

**Figure 2-87** Modifying Settings



- b. In the **Modify Settings** dialog box, select **Automatic backup**, and configure the automatic backup policy.

**Table 2-186** Automatic backup settings

| Parameter         | Description                                                                                                                                                                     |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time Zone         | Select a time zone for the backup start time.                                                                                                                                   |
| Backup Start Time | Specify the start time of auto backup.<br>Select a value from the drop-down list. The value range is from <b>00:00</b> to <b>23:00</b> . The backup always happens on the hour. |


- c. Click **OK** to enable automatic log backup.
- Automatic log backup tasks are displayed in the task list. When **Task Status** in the task list changes to **Succeeded**, the backup is successful.
- If log backup fails, click **Failed Tasks** learn the failure causes. A maximum of 10 failed tasks can be displayed. When log backup is disabled or the cluster is deleted, the failure records are also cleared.

- d. If automatic log backup is not required, click **Modify Settings**. In the displayed dialog box, deselect **Automatic backup**, and click **OK**. In the displayed confirmation dialog box, click **OK**.
- Disabling log backup does not automatically delete log backups automatically generated earlier. Instead, you need to manually delete them on the OBS console.

Manual log backup: Back up logs right away.

- a. On the **Log Backup** page, click **Back Up Manually** under **Log Backup Tasks**.

Figure 2-88 Manually backing up logs



- b. In the displayed dialog box, confirm the log backup path and click **OK**.
- The log backup task is displayed in the task list. When **Task Status** in the task list changes to **Succeeded**, the backup is successful.
- If log backup fails, click **Failed Tasks** learn the failure causes. A maximum of 10 failed tasks can be displayed. When log backup is disabled or the cluster is deleted, the failure records are also cleared.

7. Check the backed-up log files.
- Logs are backed up incrementally. After the backup is successful, you can access the target OBS bucket to obtain the full log files by clicking **Log Path**.

Figure 2-89 OBS bucket address




Table 2-187 lists the log types, where **clustername** indicates the cluster name.

Table 2-187 Log types

| Log Name                                       | Description            |
|------------------------------------------------|------------------------|
| <i>clustername</i> _deprecation.log            | Deprecation log file   |
| <i>clustername</i> _index_indexing_slowlog.log | Slow indexing log file |

| Log Name                                    | Description         |
|---------------------------------------------|---------------------|
| <i>clustername_index_search_slowlog.log</i> | Slow query log file |
| <i>clustername.log</i>                      | Run log file        |
| <i>clustername_access.log</i>               | Access log file     |

8. If the log backup function is no longer needed, you can disable it.  
On the **Log Backup** page, click **Disable Backup**. In the displayed dialog box, click **OK**. Disabling log backup does not automatically delete existing log backups. Instead, you need to manually delete them on the OBS console.

## Log Ingestion

With log ingestion, you can store the real-time logs of a cluster in itself or a different cluster on the same network, facilitating log search and analytics using Kibana.

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

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Logs > Log Ingestion**. The **Log Ingestion** page is displayed.

---

### CAUTION

If the **Log Ingestion** tab is not displayed, the cluster does not support log ingestion.

---

5. Enable log ingestion. If it is already enabled, skip this step.
  - a. Click **Enable Log Ingestion**. In the displayed dialog box, configure necessary settings.

**Table 2-188** Log ingestion settings

| Parameter         | Description                                                                                                                                                                                                                                                                                                               |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Index Prefix Name | If you set a prefix for the log file indexes, the log index names will use the format <b>index prefix name + log ingestion date</b> . The unit of log ingestion is days.<br><br>An index prefix name is a string of 1 to 128 characters. It can contain only digits, lowercase letters, underscores (_), and hyphens (-). |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Retention Period    | Set the log retention period. Ingested logs will be deleted upon expiration of this period.<br>Value range: 1 to 3650<br>Unit: days                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Log Storage Cluster | Select a cluster to store ingested logs. Options include <b>Current cluster</b> and <b>Other clusters</b> . <ul style="list-style-type: none"><li>• <b>Current cluster</b>: Store ingested logs in the current cluster.</li><li>• <b>Other clusters</b>: Store ingested logs in another cluster.<br/>Select a cluster from the <b>Cluster</b> drop-down list to use as the log storage destination. Make sure this cluster is in the same VPC and have the same version as the current cluster.<br/>When you select another cluster, click <b>Check</b> to verify network connectivity between the current cluster and the selected other cluster. Log ingestion can start only when <b>The current cluster is available</b> is displayed.</li></ul> |

- b. Click **OK** to enable log ingestion.

The configuration information will be displayed on the **Log Ingestion** tab.

6. View ingested logs.

If the task status changes to **Running**, log ingestion has started.

**Figure 2-90** Viewing ingested logs

|                                         |                                            |                                            |                             |
|-----------------------------------------|--------------------------------------------|--------------------------------------------|-----------------------------|
| Index Prefix Name<br>z                  | Retention Period<br>30days                 | Status<br>Running                          | Log Storage Cluster<br>css- |
| Kibana<br><a href="#">Access Kibana</a> | Created<br>Aug 07, 2025 15:35:58 GMT+08:00 | Updated<br>Aug 07, 2025 15:36:15 GMT+08:00 |                             |

- Click **Kibana** to log in to the target cluster and check the logs.
  - Click the cluster name at **Log Storage Cluster** to go to the cluster details page.
7. If you no longer need the log ingestion function, you can disable it to reclaim storage resources.

On the **Log Ingestion** page, click **Disable Log Ingestion**. In the displayed dialog box, click **OK**. Disabling log ingestion does not automatically remove ingested log data. Rather, it will be deleted by the system upon expiration of their retention period. You can also manually delete it from the target cluster before the retention period expires.

## Log Introduction

**Table 2-189** Introduction to different log types

| Log Type           | Description                                                                                                                                                                                                                                                                                | Purpose                                                                                                                                                                       |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Run logs           | Run logs, or main logs, record the cluster status and key information about write and query operations. For example, write logs record operations such as index creation, index mapping update, and write queue exhaustion; and query logs record query queue status and query exceptions. | Check the status and write and query operations of each cluster node, including inter-node connectivity, full GC, index creation or deletion, and cluster-level query errors. |
| Slow indexing logs | Slow indexing logs record indexing operations (such as bulk, index, update, and delete) that took a long time to complete, helping you identify performance bottlenecks.                                                                                                                   | In the case of slow write performance, you can query slow indexing logs to locate the cause.                                                                                  |
| Slow query logs    | Slow query logs record search requests that took a long time to complete. They help you monitor and analyze time-consuming search requests, so you can identify performance bottlenecks, optimize SQL queries, and improve overall system performance.                                     | In the case of slow query performance, you can query slow query logs to locate the cause.                                                                                     |
| Deprecation logs   | Deprecation logs record deprecation warnings. Deprecation warnings are written to this log when you use APIs, configurations, or functions that are marked for removal in future versions.                                                                                                 | Check for APIs or features that are about to expire in future versions.                                                                                                       |
| Access logs        | Access logs record cluster access requests, such as the request path and source address.<br><br>You cannot check access logs on the console. To check them, you need to back them up to an OBS bucket or transfer them to a target cluster first.                                          | If there is a surge in service requests, you can analyze the request sources and paths by checking the access logs.                                                           |



- **Run log description**

Run logs record the cluster status and key information about write and query operations. For example, the log record below indicates that an index named **test** was created and afterwards the cluster status changed from **YELLOW** to **GREEN**.

**Figure 2-91** A sample of run logs

```
[2025-07-04T07:15:55,976][INFO][o.e.c.m.MetadataCreateIndexService] [css-ess-esn-3-1] [test] creating index, cause [auto(bulk api)], templates [], shards [1]/[1]
[2025-07-04T07:15:56,351][INFO][o.e.c.m.MetadataMappingService] [css-ess-esn-3-1] [test/KTXEJL0cSYfC1wy3AVrqA] create_mapping [_doc]
[2025-07-04T07:15:56,720][INFO][o.e.c.r.a.AllocationService] [css-ess-esn-3-1] Cluster health status changed from [YELLOW] to [GREEN] (reason: [shards started [[test]
]([0])]).
```

Log content:

- 1. Log generation time
- 2. Log level, which can be DEBUG, INFO, WARN, or ERROR
- 3. Log-generating module
- 4. Name of the log-generating node
- 5. Log content

- **Slow indexing log description**

Slow indexing logs record indexing operations that took a long time to complete. For example, the log record below shows an indexing request that lasted longer than the preset threshold. The log contains the index name, duration, and request content.

**Figure 2-92** A sample of slow indexing logs

```
[2025-07-04T07:24:46,898][INFO][i.i.s.index] [css-ess-esn-3-1] [test/KTXEJL0cSYfC1wy3AVrqA] took[1.5ms], took_millis[1], type[_doc], id[1NFS1Jcd1WhAeIP
MOS7R], routing[], source[{"message":"[scheduler] end","micro_service":"esim-gmt","@timestamp":"2025-07-04T00:00:17.087Z"}]
```

Log content:

- 1. Log generation time
- 2. Log level, which can be DEBUG, INFO, WARN, or ERROR
- 3. Log-generating module
- 4. Name of the log-generating node
- 5. Index name and ID
- 6. Log content. In this example, the log recorded the request execution duration, index type, and index request body.

- **Slow query log description**

Slow query logs record search requests that took a long time to complete. For example, the log record below shows a search request that lasted longer than the preset threshold. The log contains the index name, duration, and request content.

**Figure 2-93** A sample of slow query logs

```
[2025-07-04T07:24:51,289][WARN][i.s.s.query] [css-ess-esn-2-1] [test][0] took[3.6ms], took_millis[3], total_hits[42 hits], types[], stats[], search_type
[QUERY_FETCH], total_shards[1], source[{"query":{"match":{"message":{"query":"[scheduler]","operator":"OR","prefix_length":0,"max_expansions":50,"fuzzy_transpositions":true,
lenient:false,"zero_terms_query":"NONE","auto_generate_synonyms_phrase_query":true,"boost":1.0}}}], id[]]
```

Log content:

- 1. Log generation time
  - 2. Log level, which can be DEBUG, INFO, WARN, or ERROR
  - 3. Log-generating module
  - 4. Name of the log-generating node
  - 5. Index name and shard ID
  - 6. Log content. In this example, the log recorded the query duration, number of hits, and query request body.
- **Deprecation log description**  
Deprecation logs record deprecation warnings. For example, the log record below indicates that **GET /\_cat/master** has been deprecated and should be replaced with **GET /\_cat/cluster\_manager**.

**Figure 2-94** A sample of deprecation logs

```
[2025-06-27T06:47:12.185][DEPRECATION][o.o.d.r.RestController] [css-ess-esn-1-1] [GET /_cat/master] is deprecated! Use [GET /_cat/cluster_manager] instead.
```

Log content:

- 1. Log generation time
  - 2. Log level, which can only be DEPRECATION.
  - 3. Log-generating module
  - 4. Name of the log-generating node
  - 5. Log content
- **Access log description**  
Access logs record cluster access requests and source addresses. For example, the log record below has recorded source information for the `/_snapshot/my_backup/my_snapshot/_restore?pretty=true` operation.

**Figure 2-95** A sample of access logs

```
2025-07-09 06:35:00 815[elasticsearch[css-f557-ess-esn-1-1]][business_worker][T#3][INFO][357535][null][POST][_snapshot/my_backup/my_snapshot/_restore?pretty=true] received request from [Netty4HttpChannel{localAddress=:10.10.1.251.9200, remoteAddress=:10.10.1.251.38257}]
```

Log content:

- 1. Log generation time
- 2. Name of the log-generating node
- 3. Name of the log-generating thread
- 4. Log level, which can be DEBUG, INFO, WARN, or ERROR.
- 5. Log request method
- 6. Request path
- 7. Source and destination addresses of the request

## Related Documents

[How Do I Set Slow Query Log Thresholds for an Elasticsearch Cluster of CSS?](#)

## 2.16 Viewing Elasticsearch Cluster Audit Logs

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 2-190** Key operations recorded by CTS

| Operation                                                                 | Resource Type | Event Name                     |
|---------------------------------------------------------------------------|---------------|--------------------------------|
| Creating a cluster                                                        | cluster       | createCluster                  |
| Changing the billing mode of a cluster from pay-per-use to yearly/monthly | cluster       | updateOndemandClusterTo-Period |
| Querying cluster details                                                  | cluster       | showClusterDetail              |
| Deleting a cluster                                                        | cluster       | deleteCluster                  |
| Changing a cluster name                                                   | cluster       | updateClusterName              |
| Querying the cluster list                                                 | cluster       | listClusters                   |
| Changing the password of a cluster                                        | cluster       | resetPassword                  |
| Restarting a cluster                                                      | cluster       | restartCluster                 |
| Expanding cluster capacity                                                | cluster       | updateExtendCluster            |
| Adding instances and expanding storage                                    | cluster       | updateExtendInstanceStorage    |
| Changing specifications                                                   | cluster       | updateFlavor                   |
| Obtaining the instance specifications list                                | cluster       | listFlavors                    |
| Querying all tags                                                         | cluster       | listClustersTags               |
| Querying the tags of a specified cluster                                  | cluster       | showClusterTag                 |
| Adding tags for a cluster                                                 | cluster       | createClustersTags             |

| Operation                                            | Resource Type | Event Name              |
|------------------------------------------------------|---------------|-------------------------|
| Deleting the tags of a cluster                       | cluster       | deleteClustersTags      |
| Adding or deleting cluster tags in batches           | cluster       | updateBatchClustersTags |
| Removing specified nodes                             | cluster       | updateShrinkNodes       |
| Changing the specifications of a specified node type | cluster       | updateFlavorByType      |
| Scaling in nodes of a specific type                  | cluster       | updateShrinkCluster     |
| Downloading a security certificate                   | cluster       | downloadCert            |
| Replacing a node                                     | cluster       | updateInstance          |
| Changing the security mode                           | cluster       | changeMode              |
| Adding a dedicated Master or Client node             | cluster       | addIndependentNode      |
| Upgrading the cluster kernel                         | cluster       | upgradeCore             |
| Obtaining a target image ID                          | cluster       | listImages              |
| Obtaining upgrade details                            | cluster       | upgradeDetail           |
| Retrying a failed upgrade task                       | cluster       | retryUpgradeTask        |
| Changing the security group                          | cluster       | changeSecurityGroup     |
| Creating a V2 cluster                                | cluster       | createClusterV2         |
| Restarting a V2 cluster                              | cluster       | restartCluster          |
| Rolling restart                                      | cluster       | rollingRestart          |
| Loading a custom word dictionary                     | cluster       | createLoadIkThesaurus   |
| Viewing custom word dictionary configurations        | cluster       | showIkThesaurus         |

| Operation                                      | Resource Type | Event Name                  |
|------------------------------------------------|---------------|-----------------------------|
| Deleting a custom word dictionary              | cluster       | deleteIcThesaurus           |
| Enabling Kibana public access                  | cluster       | startKibanaPublic           |
| Disabling Kibana public access                 | cluster       | updateCloseKibana           |
| Modifying the Kibana public network bandwidth  | cluster       | updateAlterKibana           |
| Modifying Kibana public network access control | cluster       | updatePublicKibanaWhitelist |
| Disabling Kibana public network access control | cluster       | stopPublicKibanaWhitelist   |
| Enabling logging                               | cluster       | startLogs                   |
| Disabling logging                              | cluster       | stopLogs                    |
| Querying the job list                          | cluster       | listLogsJob                 |
| Querying basic log configurations              | cluster       | showGetLogSetting           |
| Modifying basic log configurations             | cluster       | updateLogSetting            |
| Enabling automatic log backup                  | cluster       | startLogAutoBackupPolicy    |
| Disabling automatic log backup                 | cluster       | stopLogAutoBackupPolicy     |
| Backing up logs                                | cluster       | createLogBackup             |
| Querying logs                                  | cluster       | showLogBackup               |
| Enabling public network access                 | cluster       | createBindPublic            |
| Disabling public access                        | cluster       | updateUnbindPublic          |
| Modifying public network access bandwidth      | cluster       | updatePublicBandWidth       |
| Enabling the public network access whitelist   | cluster       | startPublicWhitelist        |
| Disabling the public network access whitelist  | cluster       | stopPublicWhitelist         |

| Operation                                                        | Resource Type | Event Name               |
|------------------------------------------------------------------|---------------|--------------------------|
| Automatically setting basic configurations of a cluster snapshot | cluster       | startAutoSetting         |
| Modifying basic configurations of a cluster snapshot             | cluster       | updateSnapshotSetting    |
| Manually creating a snapshot                                     | snapshot      | createSnapshot           |
| Restoring a snapshot                                             | snapshot      | restoreSnapshot          |
| Deleting a specified snapshot                                    | snapshot      | deleteSnapshot           |
| Setting an automatic snapshot creation policy                    | cluster       | createAutoCreatePolicy   |
| Querying automatic snapshot creation policies                    | cluster       | showAutoCreatePolicy     |
| Querying the snapshot list                                       | cluster       | listSnapshots            |
| Disabling snapshot function                                      | cluster       | stopSnapshot             |
| Enabling automatic snapshot creation                             | cluster       | startAutoCreateSnapshots |
| Disabling automatic snapshot creation                            | cluster       | stopAutoCreateSnapshots  |
| Enabling VPC Endpoint Service                                    | cluster       | startVpecp               |
| Disabling VPC Endpoint Service                                   | cluster       | stopVpecp                |
| Obtaining a VPCEP connection                                     | cluster       | showVpcepConnection      |
| Updating a VPCEP connection                                      | cluster       | updateVpcepConnection    |
| Modifying the VPCEP whitelist                                    | cluster       | updateVpcepWhitelist     |
| Modifying parameter settings                                     | cluster       | listYmls                 |
| Obtaining the task list of parameter settings                    | cluster       | listYmlsJob              |

| Operation                                                 | Resource Type | Event Name         |
|-----------------------------------------------------------|---------------|--------------------|
| Obtaining the parameter settings list                     | cluster       | updateYmls         |
| Querying ELB V3 load balancers supported by a cluster     | cluster       | listElbs           |
| Enabling or disabling the Elasticsearch Load Balancer     | cluster       | enableOrDisableElb |
| Configuring an Elasticsearch listener                     | cluster       | createElbListener  |
| Obtaining information about an Elasticsearch ELB          | cluster       | showElbDetail      |
| Updating an Elasticsearch listener                        | cluster       | updateEsListener   |
| Querying the certificate list                             | cluster       | listElbCerts       |
| Obtaining the intelligent O&M task list and details       | cluster       | listAiOps          |
| Creating a cluster detection task                         | cluster       | createAiOps        |
| Deleting a detection task record                          | cluster       | deleteAiOps        |
| Obtaining SMN topics available for intelligent O&M alarms | cluster       | listSmnTopics      |

## Querying Real-Time Traces

After a management tracker is created on the CTS console, the system starts recording operations performed on cloud service resources. After a data tracker is created, the system starts recording operations performed on data in OBS buckets. CTS retains operation records generated in the latest seven days.

To view or export operation records of the last seven days on the CTS console, see [Querying Real-Time Traces](#).

# 3 Using OpenSearch for Data Search

## 3.1 Procedure for Using OpenSearch

Table 3-1 Procedure for using an OpenSearch cluster

| Cate<br>gory | Operation                      |                                                                                                                                                                                                                                                                                                                                                                             | Details                                                      |
|--------------|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------|
| Use          | Plann<br>ing a<br>cluste<br>r  | Before creating an OpenSearch cluster, develop a plan for it, such as whether to deploy the cluster across multiple AZs to improve availability; the node quantity and specifications; the cluster version and security mode; and index sharding, in order to ensure the desired performance and reliability.                                                               | <a href="#">OpenSearch Cluster Planning Suggestions</a>      |
|              | Creati<br>ng a<br>cluste<br>r  | Create an OpenSearch cluster based on the plan.                                                                                                                                                                                                                                                                                                                             | <a href="#">Creating an OpenSearch Cluster (New Version)</a> |
|              | Acces<br>sing a<br>cluste<br>r | There are many ways to access an OpenSearch cluster, such as OpenSearch Dashboards, Cerebro, open-source APIs, Java, Python, and Go clients, as well as multiple network configurations over an internal network or the public network. You can select the most appropriate access method based on the programming language you prefer as well as your network environment. | <a href="#">OpenSearch Cluster Access Methods</a>            |



| Category | Operation                                 | Details                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|          | Importing data                            | There are many ways to import data to an OpenSearch cluster, including Logstash, open-source OpenSearch APIs, Cloud Data Migration (CDM), and Data Replication Service (DRS), with support for different data sources and formats, as well as real-time synchronization for relational databases. You can select the best way for yourself based on your use case and the characteristics of your data.                                                                                                                                                     |
|          | Searching for data                        | With CSS, you are advised to use DSL for data search in OpenSearch clusters. You may also use SQL.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|          | Enhancing the cluster's search capability | On top of the open-source version, CSS's OpenSearch clusters offer a range of enhanced features, including vector search, storage-compute decoupling, and switchover between hot and cold data storage classes. These features help you meet performance and cost optimization requirements for different use cases, while enhancing the service's cluster stability and search capability.                                                                                                                                                                 |
| O&M      | Backup and restoration                    | Snapshots can be created to back up the data of an OpenSearch cluster, so that data can be quickly restored in the case of accidental data loss or in case historical data is needed, improving cluster data availability.                                                                                                                                                                                                                                                                                                                                  |
|          | Scaling a cluster                         | CSS provides flexible scale-out and scale-in options, using which you can add or reduce nodes (either randomly or with specified nodes), add node types, and increase or reduce node specifications. This allows you to dynamically adjust cluster resources to meet changing demand and optimize costs.                                                                                                                                                                                                                                                    |
|          |                                           | <a href="#">Different Ways to Ingest Data into an OpenSearch Cluster</a><br><br><a href="#">Using DSL to Search for Data in OpenSearch</a><br><a href="#">Using SQL to Search for Data in OpenSearch</a><br><br><a href="#">Search Enhancement Features for OpenSearch Clusters</a><br><br><a href="#">Creating Snapshots to Back Up the Data of an OpenSearch Cluster</a><br><a href="#">Restoring a Snapshot of an OpenSearch Cluster</a><br><br><a href="#">Scaling Out/Up an OpenSearch Cluster</a><br><a href="#">Scaling In an OpenSearch Cluster</a> |

| Category | Operation                       |                                                                                                                                                                                                                                                                                                                                                                                                                                    | Details                                                         |
|----------|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|
|          | Upgrade                         | OpenSearch clusters support both same-version upgrade and cross-version upgrade. Same-version upgrade means to upgrade the kernel patches to fix problems or optimize performance. Cross-version upgrade means to upgrade the cluster version to enhance functionality or incorporate versions.                                                                                                                                    | <a href="#">Upgrading the Version of an OpenSearch Cluster</a>  |
|          | Managing clusters               | CSS provides comprehensive cluster management functions. Users can check cluster information, authorize cluster access, change the cluster's security mode, replace nodes, bind clusters with enterprise projects, switches AZs, and configure custom word dictionaries for OpenSearch clusters. They help users efficiently manage OpenSearch clusters and ensure cluster security, high availability, and optimized performance. | <a href="#">Managing OpenSearch Clusters</a>                    |
|          | Managing cluster index policies | The Index State Management (ISM) plugin of OpenSearch can be used to create and manage index lifecycle policies. These policies help automate index rollovers and deletions, helping optimize cluster performance and cut storage costs.                                                                                                                                                                                           | <a href="#">Creating and Managing OpenSearch Index Policies</a> |

| Category | Operation                     | Details                                                                                                                                                                                                                                                                                                                                                 |
|----------|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|          | Monitoring and log management | <p>CSS provides comprehensive monitoring and log management functions. Users can configure and check monitoring metrics for clusters and nodes, configure alarm rules, and back up and view logs. Intelligent O&amp;M tools help users efficiently monitor, analyze, and maintain OpenSearch clusters and ensure cluster stability and performance.</p> |
|          | Audit logs                    | <p>Cloud Trace Service (CTS) can be used to log mission-critical operations related to OpenSearch clusters. Used for auditing and accountability purposes, these log records are retained for seven days on the management console.</p>                                                                                                                 |

## 3.2 OpenSearch Cluster Planning Suggestions

### 3.2.1 Planning Cluster AZs and HA

This topic describes how to improve cluster availability by deploying a cluster across multiple AZs, including the rules for distributing cluster nodes among AZs, the recommended number of replicas, and policies for handling AZ failures.

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

Multi-AZ deployment is a high availability feature provided by CSS. Deploying a cluster across two or three AZs located in the same region can help prevent data loss and lower the possibility of service outages.

### Suggestions on Multi-AZ Deployment

If you select multi-AZ deployment when creating a cluster, CSS automatically enables cross-AZ HA to ensure that cluster nodes will be evenly distributed across the selected AZs (the difference between the number of nodes in each AZ cannot exceed 1).

You are advised to select three AZs, instead of two AZs, for multi-AZ deployment. If only two AZs are selected and one AZ becomes faulty, the cluster may not be able to elect a master node. As a result, the cluster may become unavailable.

### Node Distribution Rules

When a multi-AZ cluster is created, nodes of all types are evenly distributed across different AZs. A maximum of three AZs can be used. [Table 3-2](#) shows the node distribution when the number of AZs varies.

 **WARNING**

- When creating a multi-AZ cluster, ensure that the number of selected nodes of any type is greater than or equal to the number of AZs. Otherwise, multi-AZ cluster deployment will fail.
- If the number of data nodes or cold data nodes in a cluster is not divisible by the number of AZs, data in the cluster may be unevenly distributed, affecting data query or write performance.

Table 3-2 Node distribution when the numbers of nodes and AZs vary

| Node Quantity | 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   |
| ...           | ...       | ...           | ... | ...           | ... | ... |

### Replica Configuration Suggestions

For a multi-AZ cluster, configure the number of index replicas in a manner that can better capitalize on the high availability that comes with such as deployment.

- With a dual-AZ deployment, if one AZ becomes unavailable, the other AZ continues to provide services. In this case, configure at least one replica. If higher query performance is desired, you can increase the number of replicas.
- In the case of a three-AZ deployment, if one AZ becomes unavailable, the other AZs can continue to provide services. In this case, also configure at least

one replica. To enhance the cluster's query performance, increase the number of replicas.

For Elasticsearch and OpenSearch clusters, the default number of index replicas is 1. To enable more replicas, modify settings. The following are two examples.

- Adjust the number of replicas for an existing index:  
**curl -XPUT http://ip:9200/{index\_name}/\_settings -d '{"number\_of\_replicas":2}'**
- Set the number of replicas for new indexes using a template:  
**curl -XPUT http://ip:9200/\_template/templatename -d '{ "template": "\*", "settings": {"number\_of\_replicas": 2}}'**

where, **ip** indicates the private IP address of the cluster; **index\_name** indicates the index name; **templatename** indicates the template name; **template** indicates the index name matching rule (meaning the template will automatically apply to indexes that match this rule. The asterisk (\*) indicates that the template will apply to all new indexes); and **number\_of\_replicas** indicates the number of index replicas to change to. In this example, the number of index replicas is changed to 2.

## Service Outage Pattern Analysis

Master nodes manage cluster-wide operations, including metadata, indexes, and shard allocation. In a cluster with master nodes, the master nodes perform these tasks. In a cluster without them, data nodes and cold data nodes share the responsibilities of the master nodes.

In the case of a single AZ failure, if fewer than half (including half) of the nodes assuming master node responsibilities are still available, service availability will be affected. In this case, you need to restore services by referring to [Table 3-3](#).

For example, a cluster has three master nodes distributed in two AZs. If the AZ that contains one master node becomes unavailable while the other AZ remains available, services will still be available. However, if it is the other AZ, which contains two of the three master nodes, that fails, services will be interrupted, because fewer than half of the master nodes (or nodes that assume master node responsibilities) are available. In this case, restore services by referring to [Table 3-3](#).

**Table 3-3** Possible service outage patterns in the face of the failure of a single AZ

| Number of AZs | Number of Master Nodes                                                                       | Service Outage Patterns and Handling Suggestions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------|----------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2             | 0 (In this case, data nodes and cold data nodes share the responsibilities of master nodes.) | <ul style="list-style-type: none"><li>• If the number of data nodes plus cold data nodes is an even number:<br/>If half of the nodes become unavailable due to the failure of a single AZ, you must add one node to the still available AZ to ensure a master node can be elected for the cluster. To add nodes to your cluster, contact technical support.</li><li>• If the number of data nodes plus cold data nodes is an odd number:<ul style="list-style-type: none"><li>– If the failed AZ is the one that has one more node, you must add two nodes to the surviving AZ to ensure a master node can be elected for the cluster. To add nodes to your cluster, contact technical support.</li><li>– If the failed AZ is the other one—the one that has one less node, services will not be interrupted and a master node can be elected.</li></ul></li></ul> |
| 2             | 3                                                                                            | <p>There is a 50% chance of service interruption when one AZ fails. In this situation, one AZ has two master nodes, and the other has one:</p> <ul style="list-style-type: none"><li>• If the AZ that has one master node fails while the other AZ remains available, two master nodes are still available, so a master node can be elected from among them. In this case, services remain available.</li><li>• If the failed AZ is the other one, only one master node survives, in which case, no master node can be elected, and services will be interrupted. In this case, contact technical support.</li></ul>                                                                                                                                                                                                                                               |
| 3             | 0 (In this case, data nodes and cold data nodes share the responsibilities of master nodes.) | <p>There is a small chance of service interruption when a single node fails. However, there is one high-risk scenario: If the cluster spans three AZs and the total number of data nodes plus cold data nodes is 4, so the node distribution across the three AZs is 2-1-1, the failure of the AZ containing two nodes will cause a service interruption. To avoid this, avoid selecting four nodes for a three-AZ cluster.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| 3             | 3                                                                                            | <p>A single AZ failure will not result in service interruptions.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## 3.2.2 Planning the Cluster Version and Security Mode

This topic introduces different OpenSearch cluster versions, including features supported by each version and its application scenarios. It also describes the characteristics and application scenarios of different cluster security modes (including different application-layer protocols).

### Introduction to Different Cluster Versions

When selecting an OpenSearch cluster version, consider factors such as service requirements, available features, performance, security updates, and long-term support, ensuring that the selected version can meet both current and future needs and provide a stable, secure environment for your data.

- If you are deploying CSS OpenSearch clusters for the first time, you are advised to use the latest version.
- If you are migrating an in-house built or third-party OpenSearch cluster to CSS without altering the cluster, keep the version of the source cluster.
- If you are migrating an in-house built or third-party OpenSearch cluster to CSS while recoding it, choose OpenSearch 2.19.0.

**Table 3-4** Features supported by different versions

| Feature                       | OpenSearch 1.3.6 | OpenSearch 2.17.1 | OpenSearch 2.19.0 | Related Documents                                                                   |
|-------------------------------|------------------|-------------------|-------------------|-------------------------------------------------------------------------------------|
| Open-source vector search     | √                | √                 | √                 | <a href="#">k-NN query</a>                                                          |
| CSS vector search             | √                | x                 | √                 | <a href="#">Configuring Vector Search for OpenSearch Clusters</a>                   |
| Decoupled storage and compute | √                | x                 | √                 | <a href="#">Configuring Decoupled Storage and Compute for an OpenSearch Cluster</a> |
| Flow Control 2.0              | x                | x                 | √                 | <a href="#">Configuring Flow Control for an OpenSearch Cluster</a>                  |
| Large query isolation         | x                | x                 | √                 | <a href="#">Configuring Large Query Isolation for an OpenSearch Cluster</a>         |
| Enhanced aggregation          | x                | x                 | √                 | <a href="#">Configuring Enhanced Aggregation for an OpenSearch Cluster</a>          |

| Feature                                | OpenSearch 1.3.6 | OpenSearch 2.17.1 | OpenSearch 2.19.0 | Related Documents                                                                |
|----------------------------------------|------------------|-------------------|-------------------|----------------------------------------------------------------------------------|
| Read/write splitting                   | x                | x                 | √                 | <a href="#">Configuring Read/Write Splitting Between Two OpenSearch Clusters</a> |
| Switching between hot and cold storage | √                | √                 | √                 | <a href="#">Switching Between Hot and Cold Storage for an OpenSearch Cluster</a> |
| Index recycle bin                      | x                | x                 | √                 | <a href="#">Configuring an Index Recycle Bin for an OpenSearch Cluster</a>       |
| Enhanced data ingestion performance    | √                | x                 | √                 | <a href="#">Enhancing the Data Ingestion Performance of OpenSearch Clusters</a>  |
| Index monitoring                       | x                | x                 | √                 | <a href="#">Configuring Index Monitoring for an OpenSearch Cluster</a>           |

## Cluster Security Modes

Configure the security mode based on the security needs of your cluster.



### CAUTION

HTTPS access can be disabled only for OpenSearch 1.3.6 and 2.19.0 security-mode clusters. For other versions, HTTPS access is forcibly enabled and cannot be disabled.



**Table 3-5** Cluster security modes

| Cluster Type              | Key Settings                                                                                                                                          | Description                                                                                                                                  | Applicable Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Non-security mode cluster | <b>Security Mode:</b> disabled                                                                                                                        | Access to such a cluster requires no user authentication, and data will be transmitted in plaintext using HTTP.                              | <p>Use when creating a cluster for internal testing or workloads that have a low security standard.</p> <ul style="list-style-type: none"><li>• Advantage: easy to access the cluster.</li><li>• Disadvantage: poor security, as anyone can access the cluster. When the security mode is disabled, public network access and Kibana public network access cannot be enabled.</li></ul> <p>Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.</p> |
| Security-mode cluster     | Security-mode cluster + HTTP: <ul style="list-style-type: none"><li>• <b>Security Mode:</b> enabled</li><li>• <b>HTTPS Access:</b> disabled</li></ul> | Such a cluster requires user authentication. It supports access control and data encryption, and it uses HTTP to transmit data in plaintext. | <p>Use to balance security and performance.</p> <ul style="list-style-type: none"><li>• Advantage: User authentication improves cluster security. HTTP-based access ensures high performance of the cluster. User permissions can be configured to ensure proper isolation.</li><li>• Disadvantage: Public network access is not supported.</li></ul> <p>Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.</p>                                   |

| Cluster Type | Key Settings                                                                                                                                          | Description                                                                                                                                                        | Applicable Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|              | Security-mode cluster + HTTPS: <ul style="list-style-type: none"><li>• <b>Security Mode:</b> enabled</li><li>• <b>HTTPS Access:</b> enabled</li></ul> | Such a cluster requires user authentication. It supports access control and data encryption, and it uses HTTPS to encrypt communication and enhance data security. | Use when security takes precedence over performance and public network access is required. <ul style="list-style-type: none"><li>• Advantage: User authentication improves cluster security. HTTPS enhances cluster security by encrypting all communication over the public network. User permissions can be configured to ensure proper isolation.</li><li>• Disadvantage: When HTTPS is used, data encryption and decryption introduce computational overhead and impact the cluster's read and write performance.</li></ul> |

To access a security-mode cluster, a username and password need to be provided. CSS supports authentication for the following two types of users:

- Administrator: The default administrator username is **admin**, and the password is the one specified during cluster creation.
- Cluster user: created by the cluster administrator on OpenSearch Dashboards. User permissions can be configured to ensure proper isolation. For details, see [Creating Users for an OpenSearch Cluster and Granting Cluster Access](#).

### 3.2.3 Planning Node Types and Specifications

This topic describes the application scenarios and configuration suggestions for each type of cluster nodes, including data nodes, master nodes, client nodes, and cold data nodes. The aim is to help you properly select and configure nodes for your cluster based on service requirements. It also provides suggestions on selecting node specifications and configuring the node storage type, storage capacity, and node quantity, helping you properly plan the capacities of your cluster.

#### Planning Node Types

Before creating a cluster, determine the types of nodes to use based on service requirements, query load, data growth patterns, and performance goals.

For each cluster, data nodes are mandatory, while master, client, and cold data nodes are optional. Choose whether to enable the latter three based on your service needs and performance requirements.

**Table 3-6** Introduction to different node types

| Node Type       | Suggestions on Node Selection                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Data node (ESS) | <p>Data nodes are primarily used to store the cluster's data. The functions of the data nodes vary depending on the node types enabled for the cluster.</p> <ul style="list-style-type: none"><li>• If <b>Master node</b> and <b>Client node</b> are both selected, data nodes will be used for data storage only.</li><li>• If <b>Master node</b> and <b>Client node</b> are both unselected, data nodes will be used for all of the following purposes: cluster management, data storage, cluster access, and data analysis. To ensure reliability, a cluster should have a least three nodes.</li><li>• If <b>Master node</b> is selected but <b>Client node</b> is not, data nodes will be used for data storage, cluster access, and data analysis.</li><li>• If <b>Master node</b> is unselected but <b>Client node</b> is selected, data nodes will be used for data storage and cluster management.</li></ul> |

| Node Type                | Suggestions on Node Selection                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Master node (ess-master) | <p>Master nodes manage cluster-wide operations, including metadata, indexes, and shard allocation. For large-scale deployments, using dedicated master nodes enhances cluster stability, service availability, and centralized control.</p> <ul style="list-style-type: none"><li>• Large-scale cluster: For a cluster that has more than 10 nodes, you are advised to add dedicated master nodes to effectively manage the cluster status and metadata.</li><li>• Large quantities of indexes and index shards: If the number of indexes or shards exceeds 10,000, master nodes will have better performance handling complex cluster management tasks, avoiding impact on the performance of data nodes.</li><li>• Better management of cluster nodes: Master nodes maintain the cluster metadata, including index mapping, settings, and aliases. For a complex cluster structure, dedicated master nodes offer better management, including node joining, exiting, and fault detection. Master nodes play a critical role in cluster node management.</li><li>• Improved cluster stability and reliability: Dedicated master nodes improve cluster stability and reliability by taking over cluster management responsibilities from data storage and query nodes.</li><li>• Optimized performance for data nodes: By offloading cluster management tasks from data nodes to master nodes, you can allow data nodes to focus on data processing, which leads to improved performance.</li></ul> |

| Node Type                | Suggestions on Node Selection                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Client node (ess-client) | <p>Client nodes route and coordinate search and index requests, offloading processing from data nodes for enhanced query performance and cluster scalability when there are heavy loads.</p> <ul style="list-style-type: none"><li>• High queries per second (QPS): In the face of a high QPS, a dedicated client node can evenly distribute query requests, reducing the load of data nodes and improving the overall query performance.</li><li>• Complex aggregation queries: For complex, compute-intensive aggregation queries, a client node can dedicate to the handling of aggregation results, thus improving the efficiency and response speed of such queries.</li><li>• Large number of shards: In a cluster with a large number of shards, a client node can effectively coordinate and manage query requests to each shard, improving efficiency in request forwarding and processing.</li><li>• Reducing the load of data nodes: A client node parses search requests, determines the locations of index shards, and coordinates different nodes to execute searches. This reduces the load of data nodes by allowing them to focus on data storage and indexing.</li><li>• Improved cluster scalability: The use of client nodes allows for better cluster scalability and flexibility, enabling support for large datasets and more complex query requirements.</li></ul> |

| Node Type                 | Suggestions on Node Selection                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cold data node (ess-cold) | <p>Cold data nodes are used to store and query latency-insensitive data in large quantities. They offer an effective way to manage large datasets while cutting storage costs.</p> <ul style="list-style-type: none"><li>• Storage of historical data in large quantities: Cold data nodes offer a more cost-effective solution for storing large quantities of historical data that are infrequently accessed but useful for analytical purposes.</li><li>• Optimizing hot data performance: By migrating cold data to cold data nodes, you reduce the storage load of hot data nodes, thereby optimizing their query and write performance.</li><li>• Insensitivity to query latency: Cold data nodes are a better option for storing data that is insensitive to a high query latency.</li><li>• Cost-effectiveness: Cold data nodes usually use large-capacity disks that offer inexpensive storage.</li></ul> <p><b>CAUTION</b><br/>If no cold data nodes were enabled during cluster creation, they cannot be added later, so you have to determine whether to use cold data nodes while creating a cluster.</p> |

 **NOTE**

- If no master or client nodes were enabled when a cluster was created, you can add them if data nodes become overloaded later at some point. For details, see [Adding New Node Types](#).
- When cold data nodes are enabled, you can switch between cold and hot data storage. For details, see [Switching Between Hot and Cold Storage for an Elasticsearch Cluster](#). If there are no cold data nodes, we recommend that you use decoupled storage and compute, which can also cut storage costs. For details, see [Configuring Decoupled Storage and Compute for an Elasticsearch Cluster](#).

## Suggestions on Data Node Configuration

Table 3-7 Data node configuration

| Parameter                      | Configuration Suggestions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node Specifications            | <p>In the node flavor list, <b>vCPUs   Memory</b> indicate the number of vCPUs and memory capacity available for each flavor, and <b>Recommended Storage</b> indicates the supported storage capacity range. We recommend that you select node specifications based on service needs, such as the data volumes, performance requirements, and your spending budget.</p> <p><a href="#">Node Specifications</a> describes the application scenarios and core features of different node specifications. It can help you properly plan your cluster.</p> <p>For more information about different node specifications, see <a href="#">ECS Types</a>.</p>                                                                                                                                                                                                                                                                                                               |
| Node Storage Type and Capacity | <ul style="list-style-type: none"><li>• If the selected node flavor uses EVS disks, you need to further select <b>Node Storage Type</b> and <b>Capacity</b> based on service requirements.<ul style="list-style-type: none"><li>– For more on EVS disk performance, see <a href="#">Disk Types and Performance</a>.</li><li>– The value range of node storage capacity is determined by the node flavor you select. The value must be divisible by 20. For how to calculate the required node storage capacity, see <a href="#">Recommended Node Storage Capacity</a>.</li><li>– Node storage capacity cannot be reduced once the cluster is created. Evaluate your long-term data needs and select an appropriate size.</li></ul></li><li>• If the selected node flavor uses local disks, there is no need to select the node storage type, and the node storage capacity is a fixed value. Both of them are determined by the local disk specifications.</li></ul> |
| Nodes                          | <ul style="list-style-type: none"><li>• If master nodes are configured, the number of data nodes ranges from 1 to 200.</li><li>• If no master nodes are configured, the number of data nodes ranges from 1 to 32.</li><li>• To ensure cluster availability, you should configure at least three data nodes. For how to calculate the number of nodes needed, see <a href="#">Recommended Node Quantity</a>.</li></ul> <p>If the number of data nodes in your cluster is not evenly divisible by the number of AZs, data distribution may become unbalanced across nodes. This will negatively impact both query and write performance.</p>                                                                                                                                                                                                                                                                                                                           |

## Node Specifications

The choice of node specifications affects disk deployment. There are two types of disks: local disks and EVS disks. For example, high-performance nodes may use local disks, while standard nodes may use cost-effective EVS disks.

- Local disks refer to physical storage directly attached to the host machines running your ECSs. They deliver high I/O performance, low latency, and high throughput, making them ideal for performance-sensitive workloads that do not require long-term storage.
- EVS disks are a virtual block storage service that is independent of ECSs. They provide high reliability and fast elasticity, making them ideal for workloads that require high data reliability and highly scalable storage capacity.



### WARNING

- Local disks are less reliable than EVS disks—if the underlying host fails, data stored on local disks may be permanently lost. For this reason, avoid using local disks to store mission-critical data. Instead, use EVS disks.
- If you do use local disks, you are advised to enable automatic snapshot creation when creating the cluster, and also enable replicas when creating indexes. This helps lower the risk of data loss.

---

The table below describes the application scenarios and core features of different node specifications. It can help you properly plan your cluster.



**Table 3-8** Comparing different node specifications

| CPU Architecture | Node Flavor                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| x86              | Compute-intensive           | <p><b>Core advantages</b></p> <ul style="list-style-type: none"><li>• High-performance CPU: designed for high computational load and low latency, ideal for real-time search and complex queries.</li><li>• High reliability: Local NVMe disks enhance cluster stability.</li></ul> <p><b>Application scenarios</b></p> <ul style="list-style-type: none"><li>• E-commerce search and recommendations: real-time responses to search requests, complex filtering and ranking supported.</li><li>• App search service: quick search on mobile apps, such as social apps and content recommendations.</li><li>• Database acceleration: low latency for complex OLAP queries.</li></ul> <p><b>Precautions</b></p> <ul style="list-style-type: none"><li>• This flavor is suitable for performance-demanding workloads, but it comes at a higher cost. Be sure to balance your performance needs with budget considerations.</li><li>• Use ultra-high I/O disks with this node flavor, making sure the data read/write performance matches the CPU performance.</li></ul> |
|                  | General computing-plus - AC | <p><b>Core advantages</b></p> <ul style="list-style-type: none"><li>• Dedicated CPU: no resource contention between different instances, stable performance at a relatively low cost.</li><li>• Low latency: ideal for latency-sensitive workloads.</li></ul> <p><b>Application scenarios</b></p> <ul style="list-style-type: none"><li>• High-concurrency search: e-commerce search, social platform search, etc., where fast responses are required.</li><li>• Real-time recommendation system: recommendations are generated based on real-time user behavior.</li><li>• OLTP scenarios: transactional data processing and real-time queries.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                            |

| CPU Architecture | Node Flavor       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                  | General computing | <p><b>Core advantages</b></p> <ul style="list-style-type: none"><li>• Balanced configuration: default specifications, suitable for medium-scale workloads.</li><li>• Cost-effective: able to meet the needs of most general search and analytical applications.</li></ul> <p><b>Application scenarios</b></p> <p>This flavor meets the needs of general use cases through a standard deployment, without the need for special tuning.</p> <ul style="list-style-type: none"><li>• Content search: medium-scale search and analytical workloads. (The data volume per node ranges from 100 GB to 1,000 GB.)</li><li>• Log analytics: real-time query and analysis of medium-scale log data.</li></ul>                                                                                                                                                                                                                                                                                                        |
|                  | Memory-optimized  | <p><b>Core advantages</b></p> <ul style="list-style-type: none"><li>• Large memory capacity: ideal for memory-intensive tasks, such as complex aggregation and caching.</li><li>• High throughput: fast processing of large data volumes.</li></ul> <p><b>Application scenarios</b></p> <p>Use this flavor for search and analytical workloads where the per-node data volume ranges from 100 GB to 2,000 GB.</p> <ul style="list-style-type: none"><li>• Large-scale data analytics, such as user behavior analysis and statistics.</li><li>• OLAP scenarios: complex queries and multi-dimensional aggregation.</li><li>• Cache-intensive workloads, such as real-time reporting, where data needs to be frequently loaded to the memory.</li></ul> <p><b>Precautions</b></p> <p>Avoid using this flavor for latency-sensitive workloads.</p> <p>This flavor comes at a relatively high cost, primary due to high memory usage. Be sure to balance your performance needs with budget considerations.</p> |

| CPU Architecture | Node Flavor               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                  | Ultra-high I/O            | <p><b>Core advantages</b></p> <ul style="list-style-type: none"><li>Local NVMe SSDs: impressively high disk I/O performance, ideal for workloads with high-concurrency reads/writes.</li><li>Low latency: Compared with cloud-based SSDs, local NVMe SSDs have lower latency and better performance.</li></ul> <p><b>Application scenarios</b></p> <ul style="list-style-type: none"><li>Real-time public sentiment analysis: frequent writes and fast queries.</li><li>Patent search: quick search and matching of large volumes of text data.</li><li>Database acceleration: read/write splitting architecture for MySQL/PostgreSQL.</li></ul> <p><b>Precautions</b></p> <ul style="list-style-type: none"><li>In terms of data loss prevention, local disks are less reliable than cloud-based risks. When using local disks, be sure to enable data replicas.</li><li>NVMe SSDs are expensive. Be sure to balance performance and costs.</li></ul> |
| Kunpeng          | Kunpeng general computing | <p><b>Core advantages</b></p> <p>Cost-effective: In general, the Arm architecture is more cost-effective than x86.</p> <p><b>Application scenarios</b></p> <ul style="list-style-type: none"><li>Cost-sensitive workloads, such as search services for small and medium enterprises and testing environments.</li><li>Arm ecosystem compatibility: when your search workloads must be compatible with Arm.</li></ul> <p><b>Precautions</b></p> <p>Ecosystem compatibility: Make sure your applications (including Java virtual machine and third-party plugins) are compatible with the Arm architecture.</p>                                                                                                                                                                                                                                                                                                                                          |

## Recommended Node Storage Capacity

If you select EVS disks, you can select a node storage capacity within the range supported by the selected node flavor.

When planning the storage capacity of a cluster, consider the following factors: the original data size, number of data replicas, data bloat rate, and disk usage.

The following is a recommended formula for determining the needed cluster storage capacity:

**Storage capacity = Original data size x (1 + Number of replicas) x (1 + Data bloat rate) x (1 + Ratio of reserved space)**

where,

- Original data size: Determine the size of original data that needs to be stored.
- Number of replicas: The default value is 1.
- Data bloat rate: Extra data may be generated due to data indexing. Generally, you are advised to use a 25% data bloat rate.
- Disk usage: Considering the space occupied by the operating system and file system and the space reserved for optimized disk performance and redundancy, you are advised to keep the disk usage under 70%. That is, you need to reserve 30% of the total disk capacity.

Using the recommended values above, the formula should be: Storage capacity = Original data size x 2 x 1.25 x 1.3 = Original data size x 3.25

This formula is for quick reference only. You still need to adjust it based on the actual applications and projected data growth rate. Moreover, the storage capacity must be divisible by 20.

## Recommended Node Quantity

Data nodes and cold data nodes are used to store data. Their quantities largely determine the cluster performance baseline. The following formulas can be used to estimate the number of data nodes and cold data nodes needed:

**Number of data nodes + Number of cold data nodes = Number of write nodes + Number of query nodes**

**Number of write nodes = Traffic during peak hours (MB/s)/Number of vCPUs per node/Per-vCPU write performance baseline (MB/s) x Number of index replicas**

**Number of query nodes = Query QPS/(Number of vCPUs per node x 3/2/Average query latency (s)) x Number of index shards**

where,

- The number of vCPUs per node is determined by the selected node flavor.
- The per-vCPU write performance baseline is determined by the selected disk type. For a node that uses EVS disks, the per-vCPU write performance baseline is 1 MB/s. For a node that uses local disks, this baseline is 1.5 MB/s.
- QPS indicates the number of queries that can be processed per second.
- The average query latency is used as the query performance baseline.
- Where a high write performance is desired, we recommend setting the number of index shards to equal the total number of data nodes plus cold data nodes, thus evenly distributing the write load across all nodes.
- Where a high QPS is desired, we recommend setting the number of index shards to equal the total number of data nodes plus cold data nodes **minus 1**, thus evenly distributing the query load across all nodes.

Example:

If the peak traffic is 100 MB/s; the node specifications are 16 vCPUs and 64 GB RAM, with EVS disks; the query throughput is 1,000 QPS; the average query latency is 100 ms; and each index has three shards and two replicas, then:

Number of write nodes =  $100/16/1 \times 2 = 12$

Number of query nodes =  $1000/(16 \times 3/2/0.1) \times 3 = 12$

Number of data nodes + Number of cold data nodes = Number of write nodes + Number of query nodes =  $12 + 12 = 24$

## Suggestions on Master Node Configuration

Table 3-9 Master node configuration

| Parameter                      | Configuration Suggestions                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node Specifications            | <p>In the node flavor list, <b>vCPUs   Memory</b> indicate the number of vCPUs and memory capacity available for each flavor, and <b>Recommended Storage</b> indicates the supported storage capacity range.</p> <p>For a large-scale cluster with a large number of index shards, use large-capacity nodes as master nodes.</p> <p>Master nodes support EVS disks only. For more information about different node specifications, see <a href="#">ECS Types</a>.</p> |
| Node Storage Type and Capacity | <p>Select an appropriate storage type and capacity for master nodes.</p> <ul style="list-style-type: none"><li>For more on EVS disk performance, see <a href="#">Disk Types and Performance</a>.</li><li>The node storage capacity is fixed at 40 GB.</li></ul>                                                                                                                                                                                                       |
| Nodes                          | <p>The number of master nodes can be 3, 5, 7, or 9.</p> <p>For split-brain protection, the number of master nodes must be an odd number. For the recommended master node quantity, see <a href="#">Recommended Master Node Quantity</a>.</p>                                                                                                                                                                                                                          |

## Recommended Master Node Quantity

Increasing the number of master nodes can improve the cluster's fault tolerance performance, but it can also lead to issues like increased communication overhead, longer election time, and low resource utilization. You are advised to select an appropriate number of master nodes based on service requirements.

The following table provides the recommended number of master nodes based on the cluster size.

**Table 3-10** Recommended number of master nodes

| Cluster Size                       | Recommended Number of Master Nodes | Performance Expectation                            |
|------------------------------------|------------------------------------|----------------------------------------------------|
| Small (< 10 nodes)                 | 3                                  | Basic high availability                            |
| Medium (50 > number of nodes ≥ 10) | 3 or 5                             | Balanced performance and high availability         |
| Large (≤ 50 nodes)                 | 5 or 7                             | Enhanced fault tolerance and stability             |
| Extra-large (≥ 100 nodes)          | 5 or 7                             | Pay attention to excessive communication overhead. |

## Suggestions on Client Node Configuration

**Table 3-11** Client node configuration

| Parameter                      | Configuration Suggestions                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node Specifications            | <p>In the node flavor list, <b>vCPUs   Memory</b> indicate the number of vCPUs and memory capacity available for each flavor, and <b>Recommended Storage</b> indicates the supported storage capacity range.</p> <p>For a cluster with heavy read and write traffic, use large-capacity nodes as client nodes.</p> <p>Client nodes support EVS disks only. For more information about different node specifications, see <a href="#">ECS Types</a>.</p> |
| Node Storage Type and Capacity | <p>Select an appropriate storage type and capacity for client nodes.</p> <ul style="list-style-type: none"><li>For more on EVS disk performance, see <a href="#">Disk Types and Performance</a>.</li><li>The node storage capacity is fixed at 40 GB.</li></ul>                                                                                                                                                                                         |
| Nodes                          | <p>The number of client nodes ranges from 1 to 32.</p> <p>For the recommended client node quantity, see <a href="#">Recommended Client Node Quantity</a>.</p>                                                                                                                                                                                                                                                                                           |

## Recommended Client Node Quantity

Client nodes route and coordinate search and index requests. Its quantity should be determined based on the cluster traffic. For your reference, the following table provides the recommended number of client nodes based on the number of data nodes.

**Table 3-12** Recommended number of client nodes

| Number of Data Nodes | Recommended Number of Client Nodes                |
|----------------------|---------------------------------------------------|
| < 5                  | Do not use client nodes.                          |
| ≥ 5 and < 10         | 2                                                 |
| ≥ 10                 | Number of client nodes/Number of data nodes = 1:5 |

## Suggestions on Cold Data Node Configuration

**Table 3-13** Cold data node configuration

| Parameter                      | Configuration Suggestions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node Specifications            | <p>In the node flavor list, <b>vCPUs   Memory</b> indicate the number of vCPUs and memory capacity available for each flavor, and <b>Recommended Storage</b> indicates the supported storage capacity range. We recommend that you select node specifications based on service needs, such as the data volumes, performance requirements, and your spending budget.</p> <p><a href="#">Node Specifications</a> describes the application scenarios and core features of different node specifications. It can help you properly plan your cluster.</p> <p>For more information about different node specifications, see <a href="#">ECS Types</a>.</p>                                                                                                                                                                                                                                                                                                               |
| Node Storage Type and Capacity | <ul style="list-style-type: none"><li>• If the selected node flavor uses EVS disks, you need to further select <b>Node Storage Type</b> and <b>Capacity</b> based on service requirements.<ul style="list-style-type: none"><li>– For more on EVS disk performance, see <a href="#">Disk Types and Performance</a>.</li><li>– The value range of node storage capacity is determined by the node flavor you select. The value must be divisible by 20. For how to calculate the required node storage capacity, see <a href="#">Recommended Node Storage Capacity</a>.</li><li>– Node storage capacity cannot be reduced once the cluster is created. Evaluate your long-term data needs and select an appropriate size.</li></ul></li><li>• If the selected node flavor uses local disks, there is no need to select the node storage type, and the node storage capacity is a fixed value. Both of them are determined by the local disk specifications.</li></ul> |

| Parameter | Configuration Suggestions                                                                                                                                                                                                                                                                                                                                                               |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nodes     | <p>The value range for cold data nodes is 1 to 32.</p> <p>For how to calculate the number of nodes needed, see <a href="#">Recommended Node Quantity</a>.</p> <p>If the number of cold data nodes in your cluster is not evenly divisible by the number of AZs, data distribution may become unbalanced across nodes. This will negatively impact both query and write performance.</p> |

### 3.3 Creating an OpenSearch Cluster (New Version)

You can use OpenSearch clusters created in CSS for purposes like log analytics, data search, real-time data analytics, and more.

You can create an OpenSearch cluster in either of the following ways:

- Method 1: Create an OpenSearch cluster on the CSS management console. This topic provides the operation guide for the default UI (that is, the new UI). For the operation guide for the old UI, see [Creating an OpenSearch Cluster \(Old Version\)](#).

#### NOTE

To improve efficiency in creating clusters, CSS has made several enhancements on the cluster creation page in terms of ease-of-use. Being more intuitive, the new UI streamlines cluster creation procedures.

- Method 2: Create an OpenSearch cluster using CSS APIs. For details, see [Creating a Cluster](#).

#### Prerequisites

- You have planned the OpenSearch cluster configuration by following the guidelines in [OpenSearch Cluster Planning Suggestions](#), including the cluster size, node configuration, and storage capacity requirements.
- You have created a VPC, subnet, and security group beforehand, or you can create them when creating the cluster. For more information, see [Virtual Private Cloud User Guide](#).
- If you are going to enable automatic snapshot creation for the cluster, make sure you have created an OBS bucket for storing snapshots. The OBS bucket must be in the same region as the target OpenSearch cluster, and the storage class must be **Standard** or **Warm**.

#### Impact on Billing


- If you enable automatic snapshot creation, automatically created snapshots will be stored in standard OBS buckets, and you will need to pay extra for the storage. For more information, see [Billing Overview](#).
- VPC endpoints, if created for the cluster, will incur extra fees, depending on the resource usage. For details, see [VPC Endpoint Billing](#).



### Opening the Create Cluster Page

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the upper right corner, click **Create Cluster**. The new-version UI for creating a cluster is displayed by default.

Figure 3-1 Create Cluster (new version)



### Cluster Configuration

Select the cluster type and version.

Table 3-14 Cluster configuration

| Parameter | Description                                       |
|-----------|---------------------------------------------------|
| Type      | Choose <b>OpenSearch</b> .                        |
| Version   | Select a cluster version from the drop-down list. |

### Basic Settings

Select the region, billing mode, and AZ of the cluster.

Table 3-15 Basic settings

| Parameter | Description                                                                                                                                                                                                                                                         |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Region    | Select the region where the cluster is located. A region is the location of a physical data center. Regions are defined based on their geographical location and network latency. For lower network latency and quicker resource access, select the nearest region. |

| Parameter    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AZ           | <p>Select one or more AZs associated with the cluster region. An AZ is a physical region where resources use independent power supplies and networks. AZs are physically isolated but interconnected through an internal network.</p> <p>A maximum of three AZs can be configured. For details about the use of multiple AZs, see <a href="#">Suggestions on Multi-AZ Deployment</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Billing Mode | <p>Billing mode for the cluster, which can be <b>Yearly/Monthly</b> or <b>Pay-per-use</b>.</p> <ul style="list-style-type: none"><li>• <b>Yearly/Monthly</b>: You prepay for a yearly or monthly subscription.<br/>This mode is cost-effective when the duration of your service usage is predictable. It is recommended for long-term users. For more information, see <a href="#">Yearly/Monthly</a>.<br/>A yearly/monthly subscription cannot be deleted. If you no longer need it, unsubscribe from it.</li><li>• <b>Pay-per-use</b> (postpaid): You will be billed hourly by actual duration of use. Any partial hour of usage will be rounded up to one hour.<br/>This mode is ideal for short-term or unpredictable workloads. For more information, see <a href="#">Pay-per-Use</a>.<br/>Pay-per-use resources can be provisioned and deleted at any time.</li></ul> |

## Data Nodes

Configure the cluster's data nodes.

Data nodes are primarily used to store the cluster's data. The functions of the data nodes vary depending on the node types enabled for the cluster.

- If the cluster has no client nodes, the data nodes will need to handle the additional task of query request distribution and data analysis.
- If the cluster has no master nodes, the data nodes will also need to handle cluster management.

**Table 3-16** Configuring data nodes

| Parameter        | Description                                                                                                                                                                            |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CPU Architecture | Select the CPU architecture of the data nodes. <b>x86</b> and <b>Kunpeng</b> nodes are supported. The architectures actually supported may vary depending on the regional environment. |

| Parameter                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node Specifications            | <p>Select the specifications of the data nodes. Click <b>Available</b>. On the displayed page, select a flavor that suits your needs.</p> <p>The nodes may use either local or EVS disks. For more information about different node specifications, see <a href="#">ECS Types</a>.</p> <p>In the node flavor list, <b>vCPUs   Memory</b> indicate the number of vCPUs and memory capacity available for each flavor, and <b>Recommended Storage</b> indicates the supported storage capacity range.</p> <p>The node flavors available may vary depending on the region you select.</p> <p><b>WARNING</b></p> <ul style="list-style-type: none"><li>Local disks are less reliable than EVS disks—if the underlying host fails, data stored on local disks may be permanently lost. For this reason, avoid using local disks to store mission-critical data. Instead, use EVS disks.</li><li>If you do use local disks, you are advised to enable automatic snapshot creation when creating the cluster, and also enable replicas when creating indexes. This helps lower the risk of data loss.</li></ul> |
| Node Storage Type and Capacity | <p>Select the storage type and capacity of the data nodes.</p> <ul style="list-style-type: none"><li>If the selected node flavor uses EVS disks, you need to further select <b>Node Storage Type</b> and <b>Capacity</b> based on service requirements.<ul style="list-style-type: none"><li>Available EVS disk types vary depending on your region. For more on EVS disk performance, see <a href="#">Disk Types and Performance</a>.</li><li>The value range of node storage capacity is determined by the node flavor you select. The value must be divisible by 20.</li><li>Node storage capacity cannot be reduced once the cluster is created. Evaluate your long-term data needs and select an appropriate size.</li></ul></li><li>If the selected node flavor uses local disks, there is no need to select the node storage type, and the node storage capacity is a fixed value. Both of them are determined by the local disk specifications.</li></ul>                                                                                                                                        |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nodes     | <p>Set the number of data nodes for the cluster.</p> <ul style="list-style-type: none"><li>• If master nodes are configured, the number of data nodes ranges from 1 to 200.</li><li>• If no master nodes are configured, the number of data nodes ranges from 1 to 32.</li><li>• To ensure cluster availability, you should configure at least three data nodes.</li></ul> <p>If the number of data nodes in your cluster is not evenly divisible by the number of AZs, data distribution may become unbalanced across nodes. This will negatively impact both query and write performance.</p> |

## Master Nodes

Configure master nodes for the cluster. By default, master nodes are not enabled. To enable master nodes, select **Master Node**, and configure the node settings.

Master nodes manage cluster-wide operations, including metadata, indexes, and shard allocation. For large-scale deployments, using dedicated master nodes enhances cluster stability, service availability, and centralized control.

**Table 3-17** Configuring master nodes

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CPU Architecture    | Select the CPU architecture of the master nodes. <b>x86</b> and <b>Kunpeng</b> nodes are supported. The architectures actually supported may vary depending on the regional environment.                                                                                                                                                                                                                                                                                                          |
| Node Specifications | <p>Select the specifications of the master nodes. Click <b>Available</b>. On the displayed page, select a flavor that suits your needs.</p> <p>Master nodes support EVS disks only. For more information about different node specifications, see <a href="#">ECS Types</a>.</p> <p>In the node flavor list, <b>vCPUs   Memory</b> indicate the number of vCPUs and memory capacity available for each flavor.</p> <p>The node flavors available may vary depending on the region you select.</p> |

| Parameter                      | Description                                                                                                                                                                                                                                                                                              |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node Storage Type and Capacity | Select the storage type and capacity of the master nodes. <ul style="list-style-type: none"><li>Available EVS disk types vary depending on your region. For more on EVS disk performance, see <a href="#">Disk Types and Performance</a>.</li><li>The node storage capacity is fixed at 40 GB.</li></ul> |
| Nodes                          | Set the number of master nodes for the cluster.<br>The value can be <b>3</b> , <b>5</b> , <b>7</b> , or <b>9</b> .                                                                                                                                                                                       |

## Client Nodes

Configure client nodes for the cluster. By default, client nodes are not enabled. To enable client nodes, select **Client Node**, and configure the node settings.

Client nodes route and coordinate search and index requests, offloading processing from data nodes for enhanced query performance and cluster scalability when there are heavy loads.

**Table 3-18** Configuring client nodes

| Parameter                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CPU Architecture               | Select the CPU architecture of the client nodes. <b>x86</b> and <b>Kunpeng</b> nodes are supported. The architectures actually supported may vary depending on the regional environment.                                                                                                                                                                                                                                                                                                     |
| Node Specifications            | Select the specifications of the client nodes. Click <b>Available</b> . On the displayed page, select a flavor that suits your needs.<br><br>Client nodes support EVS disks only. For more information about different node specifications, see <a href="#">ECS Types</a> .<br><br>In the node flavor list, <b>vCPUs   Memory</b> indicate the number of vCPUs and memory capacity available for each flavor.<br><br>The node flavors available may vary depending on the region you select. |
| Node Storage Type and Capacity | Select the storage type and capacity of the client nodes. <ul style="list-style-type: none"><li>Available EVS disk types vary depending on your region. For more on EVS disk performance, see <a href="#">Disk Types and Performance</a>.</li><li>The node storage capacity is fixed at 40 GB.</li></ul>                                                                                                                                                                                     |

| Parameter | Description                                                             |
|-----------|-------------------------------------------------------------------------|
| Nodes     | Set the number of client nodes for the cluster.<br>Value range: 1 to 32 |

## Cold Data Nodes

Configure cold data nodes for the cluster. By default, it is disabled. To enable cold data nodes, select **Cold Data Node**, and configure the node settings. Cold data nodes must be configured at cluster creation and cannot be modified later.

Cold data nodes are used to store and query large quantities of latency-insensitive data. They offer an effective way to manage large datasets while cutting storage costs. When cold data nodes are enabled, you can switch between cold and hot data storage. For details, see [Switching Between Hot and Cold Storage for an OpenSearch Cluster](#). If there are no cold data nodes, we recommend that you use decoupled storage and compute, which can also cut storage costs. For details, see [Configuring Decoupled Storage and Compute for an OpenSearch Cluster](#).

**Table 3-19** Configuring cold data nodes


| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CPU Architecture    | Select the CPU architecture of the cold data nodes. <b>x86</b> and <b>Kunpeng</b> nodes are supported. The architectures actually supported may vary depending on the regional environment.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Node Specifications | <p>Select the specifications of the cold data nodes. Click <b>Available</b>. On the displayed page, select a flavor that suits your needs.</p> <p>The nodes may use either local or EVS disks. For more information about different node specifications, see <a href="#">ECS Types</a>.</p> <p>In the node flavor list, <b>vCPUs   Memory</b> indicate the number of vCPUs and memory capacity available for each flavor, and <b>Recommended Storage</b> indicates the supported storage capacity range.</p> <p>The node flavors available may vary depending on the region you select.</p> <p><b>WARNING</b></p> <ul style="list-style-type: none"><li>Local disks are less reliable than EVS disks—if the underlying host fails, data stored on local disks may be permanently lost. For this reason, avoid using local disks to store mission-critical data. Instead, use EVS disks.</li><li>If you do use local disks, you are advised to enable automatic snapshot creation when creating the cluster, and also enable replicas when creating indexes. This helps lower the risk of data loss.</li></ul> |

| Parameter                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node Storage Type and Capacity | <p>Select the storage type and capacity of the cold data nodes.</p> <ul style="list-style-type: none"><li>• If the selected node flavor uses EVS disks, you need to further select <b>Node Storage Type</b> and <b>Capacity</b> based on service requirements.<ul style="list-style-type: none"><li>– Available EVS disk types vary depending on your region. For more on EVS disk performance, see <a href="#">Disk Types and Performance</a>.</li><li>– The value range of node storage capacity is determined by the node flavor you select. The value must be divisible by 20.</li><li>– Node storage capacity cannot be reduced once the cluster is created. Evaluate your long-term data needs and select an appropriate size.</li></ul></li><li>• If the selected node flavor uses local disks, there is no need to select the node storage type, and the node storage capacity is a fixed value. Both of them are determined by the local disk specifications.</li></ul> |
| Nodes                          | <p>Set the number of cold data nodes for the cluster.<br/>Value range: 1 to 32</p> <p>If the number of cold data nodes in your cluster is not evenly divisible by the number of AZs, data distribution may become unbalanced across nodes. This will negatively impact both query and write performance.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Network Settings

Configure the VPC, subnet, IP addresses, and security group of the cluster.

**Table 3-20** Network settings

| Parameter    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VPC          | <p>Select a VPC for the cluster for proper network isolation.</p> <p>You can select a shared VPC, which is indicated by <b>Shared VPC</b> in the VPC name. To create a CSS cluster in a shared VPC, the <b>default vpc subnet statement</b> permission is required.</p> <p>VPC</p>  <p>VPC sharing allows you to centrally manage resources across multiple accounts, helping to improve resource management efficiency and reduce O&amp;M costs. For more information about shared VPCs, see <a href="#">VPC Sharing</a>.</p> <p>If none of the existing VPCs meets your requirements, click <b>Create VPC</b> to create one. (Make sure you have the permission to create VPCs.) For details, see <a href="#">Creating a VPC and Subnet</a>.</p>                                  |
| Subnet       | <p>Select a subnet for the cluster. A subnet improves network security by providing exclusive network resources that are isolated from other networks.</p> <p>Select a subnet in the current VPC. Or you can select a subnet under a shared VPC. The owner of a VPC can share the VPC to specified users.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| IPv4 Address | <p>Assign IPv4 addresses to cluster nodes.</p> <ul style="list-style-type: none"><li>• <b>Assign automatically:</b> Private IPv4 addresses will be automatically assigned to cluster nodes.</li><li>• <b>Assign manually:</b> Manually assign private IPv4 addresses to cluster nodes. Before assigning an IP address, click <b>View In-Use IP Address</b> to check IP addresses that are already in use.<ul style="list-style-type: none"><li>– If you manually specify fewer IP addresses than the total number of cluster nodes, the system will automatically assign random private IPv4 addresses to the remaining nodes.</li><li>– Manually specified IP addresses will be assigned in the following priority order: client nodes &gt; data nodes &gt; cold data nodes &gt; master nodes.</li></ul></li></ul> <p>Default value: <b>Assign automatically</b></p> |



| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Security Group | <p>Select a security group for the cluster. A security group serves as a virtual firewall that provides access control policies for clusters.</p> <p>To select a security group that meets your requirements, click <b>View Security Group</b> to go to the security group list, where you can check the details of each security group.</p> <p>The selected security group must allow port 9200 in the inbound direction. Otherwise, the cluster may be inaccessible to external services. Furthermore, to enable OpenSearch Dashboards private network access, port 5601 must also be allowed.</p> <p>If none of the existing security groups meets your requirements, click <b>Create Security Group</b> to create one. (Make sure you have the permission to create security groups.) For details, see <a href="#">Creating a Security Group</a>.</p> |

## Security Mode

Configure the cluster's security mode and access methods. The security mode is enabled by default when you create a cluster. To disable it, deselect **Security Mode**.

Configure the security mode based on the security needs of your cluster.

- **Disabling the security mode:** Use when creating a cluster for internal testing or workloads that have a low security standard. When the security mode is disabled, HTTP is used, data is transmitted in plaintext, and no user authentication is required for cluster access. Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.
- **Security mode + HTTP:** Use to balance security and performance. In this setting, user authentication is required for cluster access, yet data is transmitted via HTTP.
- **Security mode + HTTPS:** Use when security takes precedence over performance. In this setting, user authentication is required for cluster access, and data is transmitted via HTTPS in ciphertext.

---

### CAUTION

HTTPS, public network access, and OpenSearch Dashboards public network access can be enabled only when the security mode is enabled.

---

**Table 3-21** Configuring the security mode

| Parameter                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Administrator account and password | <p>When the security mode is enabled, cluster access requires user authentication. In this case, the administrator password must be configured.</p> <ul style="list-style-type: none"><li>• <b>Administrator:</b> The default value is <b>admin</b> and cannot be changed.</li><li>• <b>Administrator Password:</b> Set and confirm the administrator password. This password will be required when you access this cluster.</li></ul>                                                                                                                                                                                                                    |
| HTTPS Access                       | <p>Whether to enable HTTPS access for the cluster. Enabled by default, HTTPS improves security by encrypting all communication with the cluster.</p> <p>To use HTTP instead, deselect <b>HTTPS Access</b>. Data security cannot be guaranteed when HTTPS is disabled. Exercise caution.</p> <p>HTTPS access can be disabled only for OpenSearch 1.3.6 and 2.19.0 clusters for which the security mode is enabled. For other versions, HTTPS access is forcibly enabled and cannot be disabled.</p> <p>When HTTPS is used, data encryption and decryption introduce computational overhead, which may impact the cluster's read and write performance.</p> |

| Parameter             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Public Network Access | <p>Whether to enable public network access. By default, it is disabled. To enable it, select <b>Public Network Access</b>, and configure the necessary settings.</p> <p>Public network access can be enabled only if the security mode and HTTPS are both enabled for the cluster.</p> <p>When public network access is enabled, the cluster is automatically assigned a public IP address with dedicated dynamic BGP bandwidth.</p> <p>When <b>Public Network Access</b> is selected, further configure the following parameters:</p> <ul style="list-style-type: none"><li>• <b>Bandwidth</b>: Set the bandwidth for public network access. The value ranges from 1 Mbit/s to 200 Mbit/s.</li><li>• <b>Configure Whitelist</b><br/>Control public network access to the cluster using a whitelist.<ul style="list-style-type: none"><li>– If a whitelist is configured, only IP addresses that are on this whitelist can access the cluster over the public network.<br/>Click <b>+Add</b>. In the displayed text box, enter IP addresses or CIDR blocks that are allowed to access the cluster from the public network. Separate them using commas (,). Each value must be unique. An example of valid values: <b>192.168.1.1,10.0.0.0/24</b>. Examples of invalid values: <b>0.0.0.0</b>, <b>xx.xx.xx.x0</b>, <b>172.16.0.0-172.16.255.255</b>, non-standard formats (e.g., <b>192.168.1</b>), and duplicate values.</li><li>– If no whitelist is configured, all public IP addresses can access the cluster. However, this can be a security risk and should be avoided.</li></ul></li></ul> |

## Cluster Management

Configure the cluster name, enterprise project, and more.

**Table 3-22** Cluster management information

| Parameter    | Description                                                                                                                                                                                |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Name | <p>User-defined cluster name.</p> <p>The cluster name must start with a letter and can contain 4 to 32 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.</p> |

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Add Description    | Add a description for the cluster for easy recognition.<br>Click + <b>Add Description</b> and then enter a description of the cluster. The value can contain 0 to 128 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Enterprise Project | Associate the cluster with an enterprise project.<br>An enterprise project groups cloud resources, so you can manage resources and members by project. The default project is <b>default</b> .<br>If enterprise projects are enabled, you can select an enterprise project from the drop-down list. You can click <b>View Enterprise Projects</b> to go to the Enterprise Project Management Service console and check enterprise project details. For more information, see <a href="#">Accessing the Enterprise Center</a> .                                                                                                                                                                                                |
| Tags               | Adding tags to clusters helps you identify and manage your cluster resources.<br>Click + <b>Add</b> . You can select existing tag keys and values from the drop-down lists or enter new ones.<br>If you want to use the same tag to identify multiple cloud resources for better resource grouping, we recommend that you predefine tags in Tag Management Service (TMS). For details, see <a href="#">Predefined Tags</a> .<br>Each cluster can have a maximum of 20 tags.<br>If your organization has configured tag policies for CSS, add tags to clusters based on these policies. If a tag does not comply with the tag policies, cluster creation may fail. Contact the administrator to learn more about tag policies. |

## Automatic Snapshot Creation

Click  to expand **More Settings** and configure automatic snapshot creation. Automatic snapshot creation is disabled by default. To enable it, select **Automatic Snapshot Creation**, and configure the necessary settings. Automatic snapshot creation provides continuous data protection by periodically backing up cluster data.

---

### CAUTION

If you enable automatic snapshot creation, automatically created snapshots will be stored in standard OBS buckets, and you will need to pay extra for the storage. For more information, see [Billing Overview](#).

---

**Table 3-23** Configuring automatic snapshot creation

| Parameter                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBS Bucket                       | <p>From the drop-down list, select an OBS bucket for storing snapshots.</p> <p>If no OBS buckets meet your requirements, click <b>Create Bucket</b> to go to the OBS console and create one. For details, see <a href="#">Creating a Bucket</a>.</p> <p>To grant an IAM user access to an OBS bucket, you need to grant the <b>GetBucketStoragePolicy</b>, <b>GetBucketLocation</b>, <b>ListBucket</b>, and <b>ListAllMyBuckets</b> permissions to that user.</p>                                                                                                                                                                                                                                                                                                    |
| Backup Path                      | <p>Snapshot storage path in the OBS bucket.</p> <p>The backup path cannot:</p> <ul style="list-style-type: none"><li>• Contain the following characters: \:*?"&lt;&gt; '{}.</li><li>• Start with a slash (/).</li><li>• Start or end with a period (.).</li><li>• Contain more than two consecutive slashes (/) or periods (.).</li><li>• Exceed 512 characters.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                           |
| Maximum Backup Rate (per Second) | <p>The parameter sets the maximum backup rate per node. When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual backup rate may not reach the configured value, as it depends on factors such as OBS performance and disk I/O.</p> <p><b>Value format:</b> number + unit</p> <ul style="list-style-type: none"><li>• Number range: 0-9999</li><li>• Unit: KB, MB, GB, TB, PB, or B</li></ul> <p><b>Default value:</b> 40 MB</p> <p>The value 0MB means there is no limit on how fast data is backed up to snapshots. An overly high backup rate may lead to excessive resource usage, which may impact cluster stability. Configure this parameter carefully to maintain optimal performance.</p> |

| Parameter                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Maximum Recovery Rate (per Second) | <p>The parameter sets the maximum recovery rate per node. When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual recovery rate may not reach the configured value, as it depends on factors such as OBS performance and disk I/O.</p> <p><b>Value format:</b> number + unit</p> <ul style="list-style-type: none"><li>• Number range: 0-9999</li><li>• Unit: KB, MB, GB, TB, PB, or B</li></ul> <p>The default value is 0MB, indicating no limit.</p> <p>An overly high recovery rate may lead to excessive resource usage, which may impact cluster stability. Configure this parameter carefully to maintain optimal performance.</p> <p>For OpenSearch clusters, the recovery rate is also limited by the <b>indices.recovery.max_bytes_per_sec</b> parameter.</p> <ul style="list-style-type: none"><li>• If <b>Maximum Recovery Rate (per Second)</b> is less than <b>indices.recovery.max_bytes_per_sec</b>, the former takes effect.</li><li>• If <b>Maximum Recovery Rate (per Second)</b> is greater than <b>indices.recovery.max_bytes_per_sec</b>, the latter takes effect.</li></ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>• To check the value of <b>indices.recovery.max_bytes_per_sec</b>, run the following command:<br/>GET _cluster/settings</li><li>• To modify <b>indices.recovery.max_bytes_per_sec</b>, run the following command:<br/>PUT _cluster/settings<br/>{<br/>  "transient": {<br/>    "indices.recovery.max_bytes_per_sec": "100mb"<br/>  }<br/>}</li></ul> |

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IAM Agency | <p>Select an IAM agency to grant the current account the permission to access and use OBS. To store snapshots to an OBS bucket, you must have the required OBS access permissions.</p> <ul style="list-style-type: none"><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-obs-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the OBS Administrator permissions deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/>"obs:bucket:GetBucketLocation",<br/>"obs:object:GetObjectVersion",<br/>"obs:object:GetObject",<br/>"obs:object:DeleteObject",<br/>"obs:bucket:HeadBucket",<br/>"obs:bucket:GetBucketStoragePolicy",<br/>"obs:object:DeleteObjectVersion",<br/>"obs:bucket:ListBucketVersions",<br/>"obs:bucket:ListBucket",<br/>"obs:object:PutObject"</li><li>When OBS buckets use SSE-KMS encryption, the IAM agency must be granted KMS permissions. You can click <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b> to have the following custom policies created automatically.<br/>"kms:cmk:create",<br/>"kms:dek:create",<br/>"kms:cmk:get",<br/>"kms:dek:decrypt",<br/>"kms:cmk:list"</li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</li><li>To use an IAM agency, the following minimum permissions are required:<br/>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</li></ul> <p><b>WARNING</b><br/>The agency name can contain only letters (case-sensitive), digits, underscores (_), and hyphens (-). Otherwise, the backup will fail.</p> |

| Parameter            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Snapshot Name Prefix | A snapshot name consists of the snapshot name prefix and timestamp, for example, <i>snapshot-1566921603720</i> .<br>The snapshot name prefix can contain 1 to 32 characters. It must start with a lowercase letter, and it can only contain lowercase letters, digits, hyphens (-), and underscores (_).<br>The default value is <b>snapshot</b> .                                                                                                                                                                                     |
| Time Zone            | Select a time zone for the backup start time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Backup Start Time    | Specify the start time of auto backup.<br>Select a value from the drop-down list. The interval can be <b>Daily</b> , <b>Hourly</b> , or weekly (by selecting a specific day of the week), and the backup time can be set to any hour from 00:00 to 23:00 (full hours only).                                                                                                                                                                                                                                                            |
| Retained Snapshots   | Number of automatically created snapshots to be retained. The value ranges from 1 to 90. The system automatically deletes excess snapshots every half hour. (The expiration deletion policy applies only to the snapshots that were automatically created at the same frequency as the current automatic snapshot creation policy.)<br><b>NOTE</b><br>If the snapshot creation interval is short or if the data size of indexes is large, the number of automatic snapshots retained may not reach the value set using this parameter. |

## VPC Endpoint

Expand **More Settings** and configure a VPC endpoint for the cluster. VPC endpoint is disabled by default. To enable it, select **VPC Endpoint**, and configure the necessary settings.

VPC Endpoint enables you to access resources across VPCs using a dedicated gateway without exposing the network information of servers. The address of a VPC endpoint consists of an IPv4 address and internal domain name.

---

### CAUTION

- If a shared VPC and a subnet within this shared VPC were selected earlier for the cluster, **VPC Endpoint** cannot be enabled.
  - VPC endpoints, if created for the cluster, will incur extra fees, depending on the resource usage. For details, see [VPC Endpoint Billing](#).
  - You need specific permissions to create VPC endpoints. For details, see [VPCEP Permissions](#).
-



**Table 3-24** Configuring VPC Endpoint

| Parameter                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create Private Domain Name | <p>Whether to create a private domain name for the VPC endpoint.</p> <ul style="list-style-type: none"><li>• Enable: The system automatically assigns a private domain name to the VPC endpoint. After cluster creation, you can check this private domain name on the <b>VPC Endpoint</b> tab of the cluster details page.</li><li>• Disable: No private domain name will be configured for the VPC endpoint. The cluster can only be accessed through an IP address assigned to the VPC endpoint.</li></ul>                   |
| Add Account                | <p>Add accounts that are allowed to access the cluster through the VPC endpoint service.</p> <p>To add an account, click <b>Add Account</b> and enter an account ID.</p> <ul style="list-style-type: none"><li>• To obtain your authorized account ID, point to your username in the upper right corner, and choose <b>My Credentials</b>. Copy the value of <b>Account ID</b>.</li><li>• If no account is added or the account ID is set to *, all users are allowed to access the cluster through the VPC endpoint.</li></ul> |

## OpenSearch Dashboards Public Network Access

Expand **More Settings** and configure OpenSearch Dashboards public network access for the cluster. OpenSearch Dashboards public network access is disabled by default. To enable it, select **OpenSearch Dashboards Public Network Access**, and configure the necessary settings.

When OpenSearch Dashboards public network access is enabled, a public IP address will be automatically assigned for OpenSearch Dashboards access. Users will be able to access this Elasticsearch cluster's OpenSearch Dashboards console from the public network through this address.



### CAUTION

This parameter is available only when the security mode is enabled for the cluster.

**Table 3-25** Configuring public network access for OpenSearch Dashboards

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bandwidth           | Bandwidth for accessing OpenSearch Dashboards from the public network.<br>Value range: 1 Mbit/s to 200 Mbit/s                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Configure Whitelist | Control OpenSearch Dashboards public network access using a whitelist. <ul style="list-style-type: none"><li>• If a whitelist is configured, only IP addresses that are on this whitelist can access the cluster's OpenSearch Dashboards over the public network. Click <b>+ Add</b>. In the displayed text box, enter IP addresses or CIDR blocks that are allowed to access the cluster's OpenSearch Dashboards console from the public network. Separate them using commas (,). Each value must be unique. An example of valid values: <b>192.168.1.1,10.0.0.0/24</b>. Examples of invalid values: <b>0.0.0.0, xx.xx.xx.x0, 172.16.0.0-172.16.255.255</b>, non-standard formats (e.g., <b>192.168.1</b>), and duplicate values.</li><li>• If no whitelist is configured, all public IP addresses can access the cluster's OpenSearch Dashboards console. However, this can be a security risk and should be avoided.</li></ul> |

## Quantity

Configure the cluster's subscription duration and auto-renewal option. This parameter is available only when **Billing Mode** is set to **Yearly/Monthly**.

**Table 3-26** Configuring the subscription duration

| Parameter         | Description                                                                                                                                                                                     |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Required Duration | Set the duration of your subscription in the range of 1 month to 3 years. If you plan to use the cluster for more than nine months, you should choose a yearly subscription for a better price. |

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Auto-renew | <p>Whether to automatically renew the subscription. By default, auto-renewal is disabled. To enable it, select <b>Auto-renew</b>.</p> <p>When <b>Auto-renew</b> is selected, a yearly/monthly subscription is automatically renewed upon expiry.</p> <ul style="list-style-type: none"><li>Monthly subscription: auto-renews for 1 month each time.</li><li>Yearly subscription: auto-renews for 1 year each time.</li></ul> <p>For more information, see <a href="#">Auto-Renewal Rules</a>.</p> |

## Submit

1. In the **Configuration Summary** panel on the right, review the cluster configuration details.  
Any mandatory fields that are missing are displayed in red, and you have to set them in the configuration area.
2. Click **Create Now**.  
Cluster creation time depends on the number of nodes. Typically, this process takes less than 60 minutes, though clusters with a large number of nodes may require additional time.
3. Return to the cluster list and check the newly created cluster.  
If the cluster is created successfully, **Cluster Status** changes to **Available**.

## Related Documents

- For details about how to check basic cluster information, such as the cluster status, private IP address, and node information, see [Viewing OpenSearch Cluster Information](#).
- For details about the methods you can use to access an OpenSearch cluster, see [OpenSearch Cluster Access Methods](#).
- For how to optimize the write and query performance of a cluster and improve cluster utilization, see [Optimizing the Write Performance of a Cluster](#) and [Optimizing the Query Performance of a Cluster](#).
- You can enhance cluster performance by properly planning the number of index shards. For details, see [How Do I Plan the Quantity of Index Shards for a Cluster?](#)
- For a security-mode cluster, you need to specify the administrator account and password for cluster access. You are advised to create roles and users with the corresponding permissions on the OpenSearch Dashboards console. For details, see [Creating Users for an OpenSearch Cluster and Granting Cluster Access](#).
- If cluster creation fails, follow the guidelines in [Why Does Cluster Creation Fail?](#) for troubleshooting.

## 3.4 Creating an OpenSearch Cluster (Old Version)

This topic describes how to create an OpenSearch cluster. You can create an OpenSearch cluster in either of the following ways:

- Method 1: Create an OpenSearch cluster on the CSS management console. This topic provides the operation guide for the old UI. For the operation guide for the new UI, see [Creating an OpenSearch Cluster \(New Version\)](#).

### NOTE

To improve efficiency in creating clusters, CSS has made several enhancements on the cluster creation page in terms of ease-of-use. Being more intuitive, the new UI streamlines cluster creation procedures. You are advised to use the new UI to create OpenSearch clusters.

- Method 2: Create an OpenSearch cluster using CSS APIs. For details, see [Creating a Cluster](#).

## Scenarios

[Table 3-27](#) lists key parameters that differentiate between different types of clusters.

**Table 3-27** Parameters that differentiate between different types of clusters

| Cluster Type                            | Security Mode | HTTPS Access | Internet Access   | OpenSearch Dashboards Public Network Access |
|-----------------------------------------|---------------|--------------|-------------------|---------------------------------------------|
| Cluster with the security mode disabled | Disabled      | N/A          | Cannot be enabled | Cannot be enabled                           |
| Cluster in security mode + HTTP         | Enabled       | Disabled     | Cannot be enabled | Can be enabled                              |
| Cluster in security mode + HTTPS        | Enabled       | Enabled      | Can be enabled    | Can be enabled                              |

## Prerequisites

You have planned the OpenSearch cluster configuration by following the instructions in [OpenSearch Cluster Planning Suggestions](#).


## Creating a Cluster

- Log in to the [CSS management console](#).
- On the **Dashboard** page, click **Create Cluster** in the upper right corner. The **Create Cluster** page is displayed.

Alternatively, choose **Clusters > OpenSearch** in the navigation tree on the left. Click **Create Cluster** in the upper right corner. The **Create Cluster** page is displayed.

- Click **Back to Old Version** on the right to go to the old UI for cluster creation.

**Figure 3-2** Create Cluster (old version)



- On the **Basic Configuration** page, configure basic information and resources for the OpenSearch cluster.

**Table 3-28** Basic settings of the OpenSearch cluster

| Parameter         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Billing Mode      | Select <b>Yearly/Monthly</b> or <b>Pay-per-use</b> . <ul style="list-style-type: none"><li>• <b>Yearly/monthly</b>: You pay for the cluster by year or month, in advance. The service duration ranges from one month to three years. If you plan to use a cluster for more than nine months, you are advised to purchase a yearly package for a better price.</li><li>• <b>Pay-per-use</b>: You will be billed hourly by actual duration of use. Any partial hour of usage will be rounded up to one hour.</li></ul> |
| Required Duration | The duration for which the purchased EIP will be used. The duration must be specified if the <b>Billing Mode</b> is set to <b>Yearly/Monthly</b> .<br>Configure automatic renewal if necessary.                                                                                                                                                                                                                                                                                                                      |
| Region            | Select the region where the cluster is located.<br>ECSs in different regions cannot communicate with each other over an intranet. For lower network latency and quicker resource access, select the nearest region.                                                                                                                                                                                                                                                                                                  |
| AZ                | Select one or more AZs associated with the cluster region.<br>A maximum of three AZs can be configured. For details about the use of multiple AZs, see <a href="#">Suggestions on Multi-AZ Deployment</a> .                                                                                                                                                                                                                                                                                                          |
| Type              | Choose <b>OpenSearch</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Version           | Select a cluster version from the drop-down list.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

| Parameter             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                  | User-defined cluster name.<br>The cluster name must start with a letter and can contain 4 to 32 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Cluster Description   | Add a description for the cluster for easy recognition.<br>The value can contain 0 to 128 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Nodes                 | Number of nodes in the cluster. Select a number from 1 to 32. You are advised to configure at least three nodes to ensure cluster availability. <ul style="list-style-type: none"><li>• If the cluster has no client nodes, the data nodes will need to handle the additional task of query request distribution and data analysis.</li><li>• If the cluster has no master nodes, the data nodes will also need to handle cluster management.</li></ul> <b>NOTE</b><br>If the number of data nodes in a cluster is not evenly divisible by the number of AZs, data in the cluster may be unevenly distributed, affecting data query or write performance. |
| CPU Architecture      | <b>x86</b> and <b>Kunpeng</b> are supported. The supported types depend on the actual regional environment.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Node Specifications   | Data node flavor. You can select a specified specification based on your needs. Each cluster supports only one specification. For details, see <a href="#">ECS Types</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Node Storage Type     | If you select EVS for node storage, you need to further select the EVS disk type for data nodes of the cluster. Options include Common I/O, High I/O, Ultra-high I/O, and Extreme SSD.<br><b>NOTE</b><br>If the type of storage in use is not supported, the storage type is not displayed.                                                                                                                                                                                                                                                                                                                                                               |
| Node Storage Capacity | Data node storage capacity. Its value range varies with node specifications.<br>The node storage capacity must be a multiple of 20.<br>The node storage capacity cannot be reduced once the cluster is created. Choose an appropriate capacity based on service needs.                                                                                                                                                                                                                                                                                                                                                                                    |

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Master node    | <p>The master node is responsible for important cluster management tasks, such as metadata management, index creation and deletion, and shard allocation. It plays a critical role in metadata management, node management, stability guarantee, and cluster operation control for large-scale clusters.</p> <p>After enabling the master node, specify <b>Node Specifications</b>, <b>Nodes</b>, and <b>Node Storage Type</b>. The value of <b>Nodes</b> must be an odd number greater than or equal to 3. Up to nine nodes are supported. The value of <b>Node Storage Capacity</b> is fixed. You can select a storage type based on your needs.</p>                                                                                                                                                                                                                     |
| Client node    | <p>Client nodes receive and coordinate external requests, such as search and write requests. They play an important role in handling high-load queries, complex aggregations, managing a large number of shards, and improving cluster scalability.</p> <p>After enabling the client node, specify <b>Node Specifications</b>, <b>Nodes</b> and <b>Node Storage Type</b>. The value of <b>Nodes</b> ranges from 1 to 32. The value of <b>Node Storage Capacity</b> is fixed. You can select a storage type based on your needs.</p>                                                                                                                                                                                                                                                                                                                                        |
| Cold data node | <p>Cold data nodes are used to store query latency-insensitive data in large quantities. They offer an effective way to manage large datasets and cut storage costs.</p> <p>After enabling cold data nodes, configure <b>Node Specifications</b>, <b>Nodes</b>, <b>Node Storage Type</b>, and <b>Node Storage Capacity</b>. The value of <b>Nodes</b> ranges from 1 to 32. Select <b>Node Storage Type</b> and <b>Node Storage Capacity</b> as required.</p> <p>When cold data nodes are enabled, you can switch between cold and hot data nodes. For details, see <a href="#">Switching Between Hot and Cold Storage for an OpenSearch Cluster</a>.</p> <p><b>NOTE</b></p> <p>If the number of cold data nodes in a cluster is not evenly divisible by the number of AZs, data in the cluster may be unevenly distributed, affecting data query or write performance.</p> |

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enterprise Project | <p>When creating a CSS cluster, you can bind an enterprise project to the cluster if you have enabled the enterprise project function.</p> <p>Select an enterprise project from the <b>Enterprise Project</b> drop-down list, or click <b>View Enterprise Project</b> to go to the <b>Enterprise Project Management Service</b> page and check existing enterprise projects.</p> |

5. Click **Next: Network**.
6. On the **Network** page, configure the network settings and security mode for the OpenSearch cluster.

**Table 3-29** Network settings for the OpenSearch cluster

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VPC       | <p>Specify a VPC to isolate the cluster's network.</p> <p>Click <b>View VPC</b> to go to the VPC management console and check the created VPCs or VPCs shared with the current account. You can select a shared VPC when creating a cluster, in which case, the <b>default vpc subnet statement</b> permission is required.</p> <p>VPC sharing allows you to centrally manage resources across multiple accounts, helping to improve resource management efficiency and reduce O&amp;M costs. For more information about shared VPCs, see <a href="#">VPC Sharing</a>.</p> <p>If no VPC is available, contact the CSS administrator to create a new VPC. For details, see <a href="#">Creating a VPC and Subnet</a>.</p> <p><b>NOTE</b></p> <p>The VPC must contain CIDRs. Otherwise, cluster creation will fail. By default, a created VPC contains CIDRs.</p> |
| Subnet    | <p>A subnet provides dedicated network resources that are isolated from other networks, improving network security.</p> <p>Select a subnet needed by the cluster in the current VPC. You may select a subnet in a shared VPC.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |



| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Security Group | <p>A security group serves as a virtual firewall that provides access control policies for clusters.</p> <p>Select a security group for the cluster. Click <b>View Security Group</b> to go to the security group list, where you can view details about security groups.</p> <p><b>NOTE</b><br/>Ensure that <b>Port Range/ICMP Type</b> is set to a port range that includes port <b>9200</b> for the selected security group.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Security Mode  | <p>Whether to enable the security mode for the cluster.</p> <ul style="list-style-type: none"> <li>The security mode is enabled by default. In security mode, a cluster's communication is encrypted and access to the cluster requires user authentication. This is why <b>Administrator Username</b> and <b>Administrator Password</b> must be configured for the cluster. <ul style="list-style-type: none"> <li>The default administrator username is <b>admin</b>.</li> <li>Set and confirm the <b>Administrator Password</b>. This password will be required when you access this cluster.</li> </ul> </li> <li>If <b>Security Mode</b> is disabled, a cluster in non-security mode will be created. With such a cluster, access to the cluster will require no user authentication, and data will be transmitted in plaintext using HTTP. Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.</li> </ul> |
| HTTPS Access   | <p>HTTPS access can be enabled only when security mode is enabled for the cluster. With HTTPS access enabled, communication will be encrypted when you access the cluster.</p> <p>If the OpenSearch cluster version is 2.x or later, HTTPS is enabled by default and cannot be disabled when the security mode is enabled.</p> <p><b>NOTE</b><br/>Compared with a non-security mode cluster that uses HTTP, a security-mode cluster that uses HTTPS has lower read performance. The performance loss is estimated at around 20% under high concurrency. If you want fast read performance as well as the isolation and permission control (such as indexes, documents, and fields) enabled by the security mode, you can disable <b>HTTPS Access</b>. After <b>HTTPS Access</b> 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>                                                       |

| Parameter         | Description                                                                                                                                                                                                                                                                                                                                                 |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Public IP Address | This parameter is available only when <b>Security Mode</b> and <b>HTTPS Access</b> are enabled. When <b>Public IP Address</b> is enabled, a public IP address is automatically assigned, which will enable access to the security cluster from the Internet. For details, see <a href="#">Configuring Public Network Access for an OpenSearch Cluster</a> . |

7. Click **Next: Advanced Settings**.
8. On the **Advanced Settings** page, configure a snapshot policy and other advanced settings for the OpenSearch cluster.
  - a. Set a cluster snapshot policy.

Cluster snapshots are disabled by default. To enable them, toggle on **Cluster Snapshots**. To store snapshots automatically created in OBS, an agency will need to be created in order to access OBS. Fees will be incurred for using standard OBS storage.

**Table 3-30** Basic configuration for a cluster snapshot policy

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBS Bucket  | <p>From the drop-down list, select an OBS bucket for storing snapshots. You can also click <b>Create Bucket</b> on the right to create a new OBS bucket. For details, see <a href="#">Creating a Bucket</a>.</p> <p>The newly created or existing OBS bucket must meet the following requirements:</p> <ul style="list-style-type: none"><li>• <b>Storage Class: Standard.</b></li><li>• <b>Region</b> must be the same as that of the created cluster.</li></ul> |
| Backup Path | <p>Storage path of the snapshot in the OBS bucket.</p> <p>The backup path cannot:</p> <ul style="list-style-type: none"><li>• Contain the following characters: \:*?"&lt;&gt; '{}.</li><li>• Start with a slash (/).</li><li>• Start or end with a period (.).</li><li>• Contain more than two consecutive slashes (/) or periods (.).</li><li>• Exceed 512 characters.</li></ul>                                                                                 |

| Parameter                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Maximum Backup Rate (per Second) | <p>The parameter sets the maximum backup rate per node. When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual backup rate may not reach the configured value, as it depends on factors such as OBS performance and disk I/O.</p> <p><b>Value format:</b> number + unit</p> <ul style="list-style-type: none"><li>• Number range: 0–9999</li><li>• Unit: KB, MB, GB, TB, PB, or B</li></ul> <p><b>Default value:</b> 40 MB</p> <p>The value 0MB means there is no limit on how fast data is backed up to snapshots. An overly high backup rate may lead to excessive resource usage, which may impact cluster stability. Configure this parameter carefully to maintain optimal performance.</p> |

| Parameter                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Maximum Recovery Rate (per Second) | <p>The parameter sets the maximum recovery rate per node. When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual recovery rate may not reach the configured value, as it depends on factors such as OBS performance and disk I/O.</p> <p><b>Value format:</b> number + unit</p> <ul style="list-style-type: none"><li>• Number range: 0–9999</li><li>• Unit: KB, MB, GB, TB, PB, or B</li></ul> <p>The default value is 0MB, indicating no limit.</p> <p>An overly high recovery rate may lead to excessive resource usage, which may impact cluster stability. Configure this parameter carefully to maintain optimal performance.</p> <p>For OpenSearch clusters, the recovery rate is also limited by the <b>indices.recovery.max_bytes_per_sec</b> parameter.</p> <ul style="list-style-type: none"><li>• If <b>Maximum Recovery Rate (per Second)</b> is less than <b>indices.recovery.max_bytes_per_sec</b>, the former takes effect.</li><li>• If <b>Maximum Recovery Rate (per Second)</b> is greater than <b>indices.recovery.max_bytes_per_sec</b>, the latter takes effect.</li></ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>• To check the value of <b>indices.recovery.max_bytes_per_sec</b>, run the following command:<br/>GET _cluster/settings</li><li>• To modify <b>indices.recovery.max_bytes_per_sec</b>, run the following command:<br/>PUT _cluster/settings<br/>{<br/>  "transient": {<br/>    "indices.recovery.max_bytes_per_sec": "100mb"<br/>  }<br/>}</li></ul> |

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IAM Agency | <p>To store snapshots to an OBS bucket, you must have the required OBS access permissions. Select an IAM agency to grant the current account the permission to access and use OBS.</p> <ul style="list-style-type: none"><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-obs-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the OBS Administrator permissions deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/>"obs:bucket:GetBucketLocation",<br/>"obs:object:GetObjectVersion",<br/>"obs:object:GetObject",<br/>"obs:object:DeleteObject",<br/>"obs:bucket:HeadBucket",<br/>"obs:bucket:GetBucketStoragePolicy",<br/>"obs:object:DeleteObjectVersion",<br/>"obs:bucket:ListBucketVersions",<br/>"obs:bucket:ListBucket",<br/>"obs:object:PutObject"</li><li>When OBS buckets use SSE-KMS encryption, the IAM agency must be granted KMS permissions. You can click <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b> to have the following custom policies created automatically.<br/>"kms:cmk:create",<br/>"kms:dek:create",<br/>"kms:cmk:get",<br/>"kms:dek:decrypt",<br/>"kms:cmk:list"</li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</li><li>To use an IAM agency, the following minimum permissions are required:<br/>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</li></ul> <p><b>WARNING</b><br/>The agency name can contain only letters (case-sensitive), digits, underscores (_), and hyphens (-). Otherwise, the backup will fail.</p> |

**Table 3-31** Setting automatic snapshot creation

| Parameter            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Snapshot Name Prefix | 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, <b>snapshot-1566921603720</b> .                                                                                                                                                                                                                                            |
| Time Zone            | Time zone for the backup time. Specify <b>Backup Started Time</b> based on the time zone.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Backup Start Time    | Specify the start time of auto backup.<br>Select a time from the drop-down list. The interval can be <b>Daily</b> , <b>Hourly</b> , or weekly (by selecting a specific day of the week), and the backup time can be set to any hour from 00:00 to 23:00 (full hours only).                                                                                                                                                                                                                                                             |
| Retained Snapshots   | Number of automatically created snapshots to be retained. The value ranges from 1 to 90. The system automatically deletes excess snapshots every half hour. (The expiration deletion policy applies only to the snapshots that were automatically created at the same frequency as the current automatic snapshot creation policy.)<br><b>NOTE</b><br>If the snapshot creation interval is short or if the data size of indexes is large, the number of automatic snapshots retained may not reach the value set using this parameter. |

- b. Configure advanced settings for the cluster. Select **Default** or **Custom**.
- **Default:** **VPC Endpoint**, **OpenSearch Dashboards Public Network Access**, and **Tags** are disabled by default. You can manually enable these settings after the cluster is created.
  - **Custom:** You can enable **VPC Endpoint**, **OpenSearch Dashboards Public Network Access**, and **Tags** as required.

#### VPC Endpoint

VPC Endpoint enables you to access resources across Virtual Private Clouds (VPCs) using a dedicated gateway, without exposing the network information of servers. When VPC Endpoint is enabled, a VPC endpoint will be created by default. You can select **Create Private Domain Name** if necessary. Users will be able to access this cluster across VPCs through node IP addresses or a private domain name.

#### NOTE

- If a shared VPC and a subnet within this shared VPC were selected earlier for the cluster on the **Network** page, **VPC Endpoint** cannot be enabled for the cluster.
- After the VPC endpoint service is enabled for a cluster, you will be billed per use for the service. For more information, see [Billing Modes](#).

**Table 3-32** Configuring VPC Endpoint

| Parameter                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create Private Domain Name | If <b>Create Private Domain Name</b> is selected, the system generates a node IP address and also automatically creates a private domain name, which enables users to access this cluster from within the same VPC. If it is not selected, only a node IP address is generated.                                                                                                                                                                                                                                                                                                                                                                                      |
| VPC Endpoint Whitelist     | <p>In <b>VPC Endpoint Whitelist</b>, you can add accounts that are allowed to access the cluster using a node IP address or private domain name.</p> <ul style="list-style-type: none"><li>Click <b>Add</b> to add accounts in <b>Account ID</b>. If no account is added or the account ID is set to *, all users are allowed to access the cluster through the VPC endpoint service.</li><li>Click <b>Delete</b> in the <b>Operation</b> column to delete accounts.</li></ul> <p><b>NOTE</b><br/>To obtain your authorized account ID, point to your username in the upper right corner, and choose <b>My Credentials</b>. Copy the value of <b>Account ID</b>.</p> |

### OpenSearch Dashboards Public Network Access

This parameter is available only when the security mode is enabled for the cluster. By enabling this option, you can obtain a public IP address for accessing OpenSearch Dashboards.

**Table 3-33** Configuring public network access for OpenSearch Dashboards

| Parameter      | Description                                                                                                                                                                                                                                                              |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bandwidth      | Bandwidth for accessing Kibana through a public IP address.<br>Value range: 1 to 100<br>Unit: Mbit/s                                                                                                                                                                     |
| Access Control | If you disable this function, all IP addresses can access OpenSearch Dashboards from the public network. If you enable this function, only IP addresses or IP address ranges that are on the whitelist can access OpenSearch Dashboards through from the public network. |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Whitelist | <p>IP addresses or IP address ranges that are allowed to access the cluster from the public network. Use commas (,) to separate different values. This parameter can be configured only when <b>Access Control</b> is enabled.</p> <p>You are advised to enable the whitelist.</p> <p><b>NOTE</b></p> <p>The whitelist that controls OpenSearch Dashboards public network access depends on whitelist support by the ELB service. After you update 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 in approximately 1 minute after these connections are disconnected.</p> |

### Tags

Adding tags to clusters helps you identify and manage your cluster resources. You can customize tags or use tags predefined by Tag Management Service (TMS). If you want to use the same tag to identify multiple cloud resources for better resource grouping, we recommend that you predefine tags in Tag Management Service (TMS). For details, see [Predefined Tags Overview](#).

#### NOTE

If your organization has configured tag policies for CSS, add tags to clusters based on these policies. Tags that do not comply with predefined tag policies will cause a cluster creation failure. Contact the administrator to learn about the tag policies.

9. Click **Next: Confirm Configuration**. Check the configuration and click **Next** to create a cluster.
10. Return to the cluster list and check the newly created cluster. If the cluster is created successfully, **Cluster Status** changes to **Available**. Cluster creation time depends on the number of nodes. Typically, this process takes less than 60 minutes, though clusters with a large number of nodes may require additional time.

If cluster creation fails, try creating the cluster again by rectifying the errors returned.

## Follow-up Operations

After an OpenSearch cluster is created, you are advised to optimize the query performance of the cluster to improve efficiency by referring to [Cluster Performance Tuning](#).

## 3.5 Accessing an OpenSearch Cluster



### 3.5.1 OpenSearch Cluster Access Methods

There are many ways to connect to an OpenSearch cluster. Select the most appropriate access method based on the programming language you prefer. The various access methods support different types of network configurations. The cluster access addresses vary depending on the network configuration you choose. Select an appropriate network configuration method based on your network environment.

#### Access Methods

**Table 3-34** Cluster access methods

| Access Method                       | Scenario                                                                                                                                                                    | Details                                                                                                                                                             |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OpenSearch Dashboards (recommended) | <ul style="list-style-type: none"><li>• A graphical user interface</li><li>• Instance monitoring</li><li>• Data management</li></ul> These methods are language-irrelevant. | <a href="#">Logging In to an OpenSearch Cluster Through OpenSearch Dashboards</a>                                                                                   |
| Cerebro                             |                                                                                                                                                                             | <a href="#">Logging In to an OpenSearch Cluster Through Cerebro</a>                                                                                                 |
| Open-source OpenSearch API          | Run Curl commands to access and manage OpenSearch clusters.                                                                                                                 | <a href="#">Accessing an OpenSearch Cluster Using cURL Commands</a>                                                                                                 |
| Java                                | Use Java to access and manage OpenSearch clusters.<br><br>To avoid compatibility issues, use the Java client that matches the software version of your OpenSearch cluster.  | <a href="#">Accessing an OpenSearch Cluster Using the High Level REST Client</a><br><a href="#">Accessing an OpenSearch Cluster Using the Low Level REST Client</a> |
| Python                              | Use Python to access and manage OpenSearch clusters.                                                                                                                        | <a href="#">Accessing an OpenSearch Cluster with the Python Client</a>                                                                                              |
| LDAP Authentication                 | Enable Lightweight Directory Access Protocol (LDAP) authentication for an OpenSearch cluster.                                                                               | <a href="#">Accessing an OpenSearch Cluster Using LDAP</a>                                                                                                          |

#### Network Configuration

Configure the network of the OpenSearch cluster, and ensure that you have a server that can connect to the cluster's network. Obtain the cluster access address based on the selected network. Use this address to access the cluster.

**Table 3-35** Network configuration for the cluster

| Network Environment | Access Method                                         | Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Details                                                                    |
|---------------------|-------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|
| Private network     | Private IP address                                    | <p>Obtain the private IP address of each node in the cluster and access the cluster by connecting to one of these IP addresses.</p> <p>How to obtain:</p> <ol style="list-style-type: none"><li>1. Log in to the <a href="#">CSS management console</a>.</li><li>2. In the navigation pane on the left, choose <b>Clusters &gt; OpenSearch</b>.</li><li>3. In the cluster list, obtain the target cluster's private IP address from the <b>Private IP Address</b> column. Generally, the IP address format is <i>&lt;host&gt;:&lt;port&gt;</i> or <i>&lt;host&gt;:&lt;port&gt;,&lt;host&gt;:&lt;port&gt;</i>. If the cluster has only one node, the IP address and port number of this one node are displayed, for example, <b>10.62.179.32:9200</b>. If the cluster has multiple nodes, the IP addresses and port numbers of all nodes are displayed, for example, <b>10.62.179.32:9200,10.62.179.33:9200</b>.</li></ol> | This is the default method.                                                |
| Private network     | IP address or private domain name of the VPC endpoint | <p>Configure a VPC endpoint for the cluster, enabling cross-VPC access over the internal network. The cluster can be accessed via the VPC endpoint IP address or a private domain name.</p> <p>How to obtain:</p> <ol style="list-style-type: none"><li>1. Log in to the <a href="#">CSS management console</a>.</li><li>2. In the navigation pane on the left, choose <b>Clusters &gt; OpenSearch</b>.</li><li>3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.</li><li>4. Choose <b>Cluster Access &gt; VPC Endpoint</b>. In the VPC endpoint list below, record <b>Service Address</b> or <b>Private Domain Name</b> of the target VPC endpoint.</li></ol>                                                                                                                                                                                                     | <a href="#">Configuring VPC Endpoint Service for an OpenSearch Cluster</a> |

| Network Environment | Access Method                                   | Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Details                                                                         |
|---------------------|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Private network     | Private IP address of a dedicated load balancer | <p>Associate a dedicated load balancer with the cluster to enable access through the private IP address of this dedicated load balancer.</p> <p>How to obtain:</p> <ol style="list-style-type: none"> <li>1. Log in to the <a href="#">CSS management console</a>.</li> <li>2. In the navigation pane on the left, choose <b>Clusters &gt; OpenSearch</b>.</li> <li>3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.</li> <li>4. Choose <b>Cluster Access &gt; Load Balancing</b>. On the <b>OpenSearch</b> tab, find and record the private or public IP address of the target load balancer.</li> </ol>                                    | <a href="#">Configuring a Dedicated Load Balancer for an OpenSearch Cluster</a> |
| Public network      | Public IP address                               | <p>Enable public network access for a security-mode cluster that uses HTTPS, and configure a whitelist for access control. The cluster can be accessed from the public network via its public IP address.</p> <p>How to obtain:</p> <ol style="list-style-type: none"> <li>1. Log in to the <a href="#">CSS management console</a>.</li> <li>2. In the navigation pane on the left, choose <b>Clusters &gt; OpenSearch</b>.</li> <li>3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.</li> <li>4. Click the <b>Overview</b> tab. In the <b>Configuration</b> area, find and record the IP address used for public network access.</li> </ol> | <a href="#">Configuring Public Network Access for an OpenSearch Cluster</a>     |

| Network Environment | Access Method                                  | Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Details                                                                         |
|---------------------|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Public network      | Public IP address of a dedicated load balancer | <p>Associate with the cluster a dedicated load balancer configured with an EIP to enable access through the public IP address of this dedicated load balancer.</p> <p>How to obtain:</p> <ol style="list-style-type: none"><li>1. Log in to the <a href="#">CSS management console</a>.</li><li>2. In the navigation pane on the left, choose <b>Clusters &gt; OpenSearch</b>.</li><li>3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.</li><li>4. Choose <b>Cluster Access &gt; Load Balancing</b>. On the <b>OpenSearch</b> tab, find and record the private or public IP address of the target load balancer.</li></ol> | <a href="#">Configuring a Dedicated Load Balancer for an OpenSearch Cluster</a> |

### 3.5.2 Logging In to an OpenSearch Cluster Through OpenSearch Dashboards

OpenSearch Dashboards is the data visualization and exploration platform for OpenSearch. It enables in-depth data analysis and interactive visualization. In CSS, OpenSearch Dashboards is pre-built for each OpenSearch cluster. You can start OpenSearch Dashboards with one click, without installing anything. OpenSearch Dashboards offers a comprehensive suite of dashboard features and visualization tools. It seamlessly integrates OpenSearch's analytical capabilities, supporting the full analytics process from data exploration to actionable business insights.

OpenSearch Dashboards supports multiple access methods. Steps needed to log in to an OpenSearch cluster vary depending on the access method you choose. See [Table 3-36](#).

**Table 3-36** Methods for logging In to an OpenSearch cluster through OpenSearch Dashboards

| OpenSearch Dashboards Access Method                                | Constraints                                                                                                                                      | Details                                                                                                                |
|--------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| One-click access to OpenSearch Dashboards from the service console | N/A                                                                                                                                              | <a href="#">Logging In to an OpenSearch Cluster by Accessing OpenSearch Dashboards Through the Console</a>             |
| Accessing OpenSearch Dashboards using a public IP address          | <ul style="list-style-type: none"><li>Only clusters in security mode support OpenSearch Dashboards access through a public IP address.</li></ul> | <a href="#">Logging In to an OpenSearch Cluster by Accessing OpenSearch Dashboards Using a Public IP Address</a>       |
| Accessing OpenSearch Dashboards using a private IP address         | The client that accesses OpenSearch Dashboards must be able to reach the private IP address of OpenSearch Dashboards.                            | <a href="#">Logging In to an OpenSearch Cluster by Accessing OpenSearch Dashboards Using a Private Network Address</a> |

## Constraints Concerning the Use of OpenSearch Dashboards

You can customize the username, role name, and tenant name in OpenSearch Dashboards. Chinese characters are not allowed.

## Logging In to an OpenSearch Cluster by Accessing OpenSearch Dashboards Through the Console

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
  - Non-security mode cluster: The OpenSearch Dashboards console is displayed.
  - Security-mode cluster: Enter the username and password on the login page and click **Log In** to go to the OpenSearch Dashboards console. The default username is **admin** and the password is the one specified during cluster creation.

- After logging in, you can use OpenSearch Dashboards to manage the cluster.

## Logging In to an OpenSearch Cluster by Accessing OpenSearch Dashboards Using a Public IP Address

### WARNING

- Only clusters in security mode support OpenSearch Dashboards access through a public IP address.
- The whitelist that controls OpenSearch Dashboards public network access depends on whitelist support by the ELB service. After you update 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 in approximately 1 minute after these connections are disconnected.
- If you disable OpenSearch Dashboards public network access and then re-enable it, the public IP address for accessing OpenSearch Dashboards may change. Exercise caution.

- Log in to the [CSS management console](#).
- In the navigation pane on the left, choose **Clusters > OpenSearch**.
- In the cluster list, click the name of the target cluster. The cluster information page is displayed.
- Choose **Cluster Access > Dashboards Public Network Access** to check whether OpenSearch Dashboards public network access is enabled for the cluster.
  - Yes: Go to [6](#).
  - No: Go to the next step.
- Enable OpenSearch Dashboards public network access for the OpenSearch cluster.
  - On the **Dashboards Public Network Access** page, toggle on the **Dashboards Public Network Access** button.
  - In the displayed dialog box, set the parameters.

**Table 3-37** Configuring public network access for OpenSearch Dashboards

| Parameter | Description                                                                                                   |
|-----------|---------------------------------------------------------------------------------------------------------------|
| Bandwidth | Bandwidth for accessing OpenSearch Dashboards from the public network.<br>Value range: 1 Mbit/s to 200 Mbit/s |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure Whitelist | <p>Control OpenSearch Dashboards public network access using a whitelist.</p> <ul style="list-style-type: none"><li>If a whitelist is configured, only IP addresses that are on this whitelist can access the cluster's OpenSearch Dashboards over the public network.<br/>Click <b>+ Add</b>. In the displayed text box, enter IP addresses or CIDR blocks that are allowed to access the cluster's OpenSearch Dashboards console from the public network. Separate them using commas (,). Each value must be unique. An example of valid values: <b>192.168.1.1,10.0.0.0/24</b>. Examples of invalid values: <b>0.0.0.0</b>, <b>xx.xx.xx.x0</b>, <b>172.16.0.0-172.16.255.255</b>, non-standard formats (e.g., <b>192.168.1</b>), and duplicate values.</li><li>If no whitelist is configured, all public IP addresses can access the cluster's OpenSearch Dashboards console. However, this can be a security risk and should be avoided.</li></ul> |

c. Click **OK** to confirm the settings.

If OpenSearch Dashboards public network access is already enabled, you can modify relevant settings.

- After OpenSearch Dashboards public network access is enabled, obtain **Dashboards Public IP Address** on the **Dashboards Public Network Access** page.

**Figure 3-3** Obtaining the OpenSearch Dashboards public IP address

Dashboards Public Network Access  For more information, see [this document](#).

Dashboards Public IP Address <https://100.85.126.56:5601>

- Enter the OpenSearch Dashboards public IP address in the browser address box to go to the OpenSearch Dashboards login page.  
Enter the username and password on the login page and click **Log In** to log in to the OpenSearch Dashboards console. The default username is **admin** and the password is the one specified during cluster creation.
- After the login is successful, you can access the OpenSearch cluster through OpenSearch Dashboards.

## Logging In to an OpenSearch Cluster by Accessing OpenSearch Dashboards Using a Private Network Address

### ⚠ CAUTION

The client that accesses OpenSearch Dashboards must be able to reach the private IP address of OpenSearch Dashboards.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Overview** tab. In the **Configuration** area, record **Private IPv4 Address**.

**Figure 3-4** Obtaining the private IP address

|                      |               |
|----------------------|---------------|
| Enterprise Project   | default       |
| HTTPS Access         | Disabled      |
| Private IPv4 Address | 10 . . . . 00 |

5. Obtain the private IP address of OpenSearch Dashboards.  
By changing the port number in the cluster's private IP address from **9200** to **5601**, you get the private IP address of OpenSearch Dashboards. For example, if the cluster's private IPv4 address is **192.168.0.\*\*\*:9200**, the private IP address of OpenSearch Dashboards is **192.168.0.\*\*\*:5601**.
6. (Optional) Configure a return route. To enable a client to access the cluster's OpenSearch Dashboards console across different networks, configure a route for the OpenSearch cluster.
  - a. Connect the client and OpenSearch Dashboards through a Direct Connect or VPC peering connection.
  - b. Configure the route connecting the OpenSearch cluster and the client.  
On the **Cluster Information** page, find **Cluster Route**, and click **Add Route** on the right. In the displayed dialog box, set **IP Address** and **Subnet Mask**.  
For details, see [Configuring Routes for an OpenSearch Cluster](#).
7. On the client, enter the private network address of OpenSearch Dashboards (for example, **https://192.168.0.xx:5601**) to go to the OpenSearch Dashboards login page.
  - Non-security mode cluster: The OpenSearch Dashboards console is displayed.
  - Security-mode cluster: Enter the username and password on the login page and click **Log In** to go to the OpenSearch Dashboards console. The default username is **admin** and the password is the one specified during cluster creation.



8. After the login is successful, you can access the OpenSearch cluster through OpenSearch Dashboards.

## Related Documents

- For how to use in-house developed OpenSearch Dashboards to access OpenSearch clusters in CSS, see [How Do I Connect In-house Developed OpenSearch Dashboards to an OpenSearch Cluster in CSS?](#).
- For routine O&M tasks (such as shard adjustment, index management, and performance monitoring), you are advised to use Cerebro for cluster login. For details, see [Logging In to an OpenSearch Cluster Through Cerebro](#).

### 3.5.3 Logging In to an OpenSearch Cluster Through Cerebro

Cerebro is an open-source Elasticsearch web admin tool, designed for real-time monitoring and efficient cluster operations. Cerebro works for OpenSearch just as well. In CSS, Cerebro is pre-built for each OpenSearch cluster. You can start Cerebro with one click, without installing anything. Cerebro allows you to quickly check a cluster's health status, node distribution, and index details. For routine O&M tasks, including shard adjustment, index management, and performance monitoring, we recommend using Cerebro. Its intuitive cluster topology views and comprehensive management features streamline O&M workflows, reducing complexity while improving efficiency. Cerebro is suitable for different roles, including developers, O&M teams, and data analysts.

#### Accessing a Cluster Through Cerebro on the CSS Console

The OpenSearch clusters created in CSS provide a built-in Cerebro component, using which you can quickly access these clusters.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. Find the target cluster, and choose **More > Cerebro** in the **Operation** column. The Cerebro login page is displayed.
  - 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.
4. After the login is successful, you can access clusters through Cerebro.

#### Accessing a Cluster Using an In-House Built Cerebro



#### CAUTION

- When using in-house built Cerebro to access a cluster, make sure the network between them is connected.
- Make sure the version of the in-house built Cerebro is compatible with the OpenSearch cluster version.

1. Obtain the cluster access address. For details about how to obtain the cluster access address for different network configurations, see [Network Configuration](#).
2. Start Cerebro, and enter the cluster access address.
  - For a cluster in security mode, enter `https://access address:9200`.  
To log in to a cluster in security mode, you need to provide a username and password.
  - For a cluster in non-security mode, enter `http://access address:9200`.

### 3.5.4 Accessing an OpenSearch Cluster Using cURL Commands

Open-source OpenSearch provides a series of RESTful APIs. You can run cURL commands to access these APIs using tools such as OpenSearch Dashboards and Postman. This topic describes how to run cURL commands to access an OpenSearch cluster.

#### Prerequisites

- The target CSS OpenSearch cluster is available.
- An ECS has been created, and it is in the same VPC and security group as the CSS cluster.
  - If they are not in the same security group, modify the ECS security group, or configure its inbound and outbound rules to allow all access from the cluster. For details, see [Configuring Security Group Rules](#).
  - For details about how to use an ECS, see [Purchasing a Custom ECS](#).

#### Accessing a Cluster

1. Obtain the private network address of the cluster. It is used to access the cluster.

This topic uses a private IP address as an example to describe how to access a cluster. The cluster access address varies depending on the network configurations used. For details, see [Network Configuration](#).

  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > OpenSearch**.
  - c. In the cluster list, obtain the target cluster's private IP address from the **Private IP Address** column. Generally, the IP address format is `<host>.<port>` or `<host>.<port>,<host>.<port>`.  
If the cluster has only one node, the IP address and port number of this 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**.
2. Run one of the following commands on the ECS to access the cluster. The access command varies depending on the security mode of the cluster.
  - For a cluster with the security mode disabled:  

```
curl "http://<host>.<port>"
```
  - For a security-mode cluster that uses HTTP:  

```
curl -u <user>:<password> "http://<host>.<port>"
```

- For a security-mode cluster that uses HTTPS:

```
curl -u <user>:<password> -k "https://<host>:<port>"
```

**Table 3-38** Variables

| Variable   | Description                                                                                                                                                                                                          |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <host>     | IP address of each node in the cluster. If the cluster contains multiple nodes, there will be multiple IP addresses. You can use any of them.                                                                        |
| <port>     | Port number for accessing a cluster node. Generally, the port number is 9200.                                                                                                                                        |
| <user>     | Username for accessing the cluster.                                                                                                                                                                                  |
| <password> | Password of the user.<br>If the password contains special characters, enclose the username and password in single quotation marks, for example, <b>curl -u "user:password!" "http://&lt;host&gt;:&lt;port&gt;"</b> . |

Example command:

```
curl "http://10.62.176.32:9200"
```

Example response:

```
HTTP/1.1 200 OK
content-type: application/json; charset=UTF-8
content-length: 513

{
 "name" : "css-xxx-ess-esn-3-1",
 "cluster_name" : "css-xxx",
 "cluster_uuid" : "xxx_uuid",
 "version" : {
 "number" : "7.10.2",
 "build_type" : "tar",
 "build_hash" : "unknown",
 "build_date" : "unknown",
 "build_snapshot" : true,
 "lucene_version" : "9.12.1",
 "minimum_wire_compatibility_version" : "7.10.0",
 "minimum_index_compatibility_version" : "7.0.0"
 },
 "tagline" : "The OpenSearch Project: https://opensearch.org/"
}
```

#### NOTE

For more commands, see the [OpenSearch documentation](#).

## 3.5.5 Accessing an OpenSearch Cluster Using Java

### 3.5.5.1 Accessing an OpenSearch Cluster Using the High Level REST Client

You can query and manage data in a CSS OpenSearch cluster using the High Level REST Client. The High Level REST Client encapsulates OpenSearch APIs. You only

need to construct the required structures to access an OpenSearch cluster. This simplifies the process of working with an OpenSearch cluster. For details, see [Java high-level REST client](#).

## Prerequisites

- The target OpenSearch cluster is available.
- The server that runs the Java code can communicate with the OpenSearch cluster.
- Depending on the network configuration method used, obtain the cluster access address. For details, see [Network Configuration](#).
- Java has been installed on the server and the JDK version is 11 or later. Download JDK 11 from [Java Archive Downloads](#).

## Introducing Dependencies

Introduce the required Java dependencies on the server where you run Java code.

CSS allows you to use the Elasticsearch 7.10 Java client to connect to OpenSearch clusters. To ensure better compatibility, however, you are advised to use an OpenSearch Java client that has the same version as the target OpenSearch cluster.

If your High Level REST Client version is later than the OpenSearch cluster version and there are incompatibility issues with some requests, you can use **`RestHighLevelClient.getLowLevelClient()`** to obtain the Low Level REST Client and customize OpenSearch requests. For details, see [Accessing an OpenSearch Cluster Using the Low Level REST Client](#).

Select a reference example based on the Java client you use.

Scenario 1 (recommended): Use the OpenSearch Java client to access an OpenSearch cluster and introduce dependencies.

- Maven:  
Replace 1.3.6 with the actual version number.

```
<dependency>
 <groupId>org.opensearch.client</groupId>
 <artifactId>opensearch-rest-high-level-client</artifactId>
 <version>1.3.6</version>
</dependency>
```

- Gradle:  
Replace 1.3.6 with the actual version number.

```
compile group: 'org.opensearch.client', name: 'opensearch-rest-high-level-client', version: '1.3.6'
```

Scenario 2: Use the Elasticsearch 7.10.2 Java client to access an OpenSearch cluster and introduce dependencies.

- Maven:

```
<dependency>
 <groupId>org.elasticsearch.client</groupId>
 <artifactId>elasticsearch-rest-high-level-client</artifactId>
 <version>7.10.2</version>
</dependency>
<dependency>
 <groupId>org.elasticsearch</groupId>
 <artifactId>elasticsearch</artifactId>
```

```
<version>7.10.2</version>
</dependency>
```

- Gradle:  
compile group: 'org.elasticsearch.client', name: 'elasticsearch-rest-high-level-client', version: '7.10.2'

### Accessing a Cluster

The sample code varies depending on the Java client and the security mode settings of the target OpenSearch cluster. Select the right reference document based on your service scenario.

**Table 3-39** Cluster access scenarios

| Java Client          | OpenSearch Cluster Security-Mode Settings     | Security Certificate | Details                                                                                                                      |
|----------------------|-----------------------------------------------|----------------------|------------------------------------------------------------------------------------------------------------------------------|
| OpenSearch           | Non-security mode                             | -                    | <a href="#">Connecting to a Non-Security Mode Cluster Using the OpenSearch High Level REST Client</a>                        |
| OpenSearch           | Security mode + HTTP<br>Security mode + HTTPS | No                   | <a href="#">Connecting to a Security-Mode Cluster Using the OpenSearch High Level REST Client (Without a Certificate)</a>    |
| OpenSearch           | Security mode + HTTPS                         | Yes                  | <a href="#">Connecting to a Security-Mode Cluster Using the OpenSearch High Level REST Client (With a Certificate)</a>       |
| Elasticsearch 7.10.2 | Non-security mode                             | -                    | <a href="#">Connecting to a Non-Security Mode Cluster Using the Elasticsearch High Level REST Client</a>                     |
| Elasticsearch 7.10.2 | Security mode + HTTP<br>Security mode + HTTPS | No                   | <a href="#">Connecting to a Security-Mode Cluster Using the Elasticsearch High Level REST Client (Without a Certificate)</a> |
| Elasticsearch 7.10.2 | Security mode + HTTPS                         | Yes                  | <a href="#">Connecting to a Security-Mode Cluster Using the Elasticsearch High Level REST Client (With a Certificate)</a>    |

## Connecting to a Non-Security Mode Cluster Using the OpenSearch High Level REST Client

Use the OpenSearch High Level REST Client to connect to an OpenSearch cluster for which the security mode is disabled, and query whether the **test** index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
import org.opensearch.client.RequestOptions;
import org.opensearch.client.RestClient;
import org.opensearch.client.RestClientBuilder;
import org.opensearch.client.RestHighLevelClient;
import org.opensearch.client.indices.GetIndexRequest;
import java.io.IOException;
import java.util.Arrays;
import java.util.List;
public class RestHighLevelClientExample {
 public static void main(String[] args) throws IOException {
 List<String> host = Arrays.asList("{Cluster access address}");
 RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, 9200, "http"));
 final RestHighLevelClient client = new RestHighLevelClient(builder);
 GetIndexRequest indexRequest = new GetIndexRequest("test");
 boolean exists = client.indices().exists(indexRequest, RequestOptions.DEFAULT);
 System.out.println(exists);
 client.close();
 }
 /**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
 public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
 }
}
```

This piece of code checks whether the **test** index exists in the cluster. If **true** or **false** is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode Cluster Using the OpenSearch High Level REST Client (Without a Certificate)

Use the OpenSearch High Level REST Client to connect to a security-mode OpenSearch cluster (HTTP or HTTPS) without loading a security certificate, and query whether the **test** index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
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.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.opensearch.action.admin.cluster.health.ClusterHealthRequest;
import org.opensearch.action.admin.cluster.health.ClusterHealthResponse;
import org.opensearch.client.RequestOptions;
import org.opensearch.client.RestClient;
import org.opensearch.client.RestClientBuilder;
import org.opensearch.client.RestHighLevelClient;
import org.opensearch.client.indices.GetIndexRequest;
import org.opensearch.common.Nullable;
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 java.util.Arrays;
import java.util.List;
import java.util.Objects;
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;
/**
 * Connect to a security-model cluster using the High Level REST Client (without a certificate).
 */
public class RestHighLevelClientExample {
 /**
 * Create a class for the client. Define the create function.
 */
 public static RestHighLevelClient create(List<String> host, int port, String protocol, int connectTimeout,
int connectionRequestTimeout, int socketTimeout, String username, String password) throws IOException{
 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(constructHttpHosts(host, port, protocol))
 .setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout)
 .setConnectionRequestTimeout(connectionRequestTimeout)
 .setSocketTimeout(socketTimeout))
 .setHttpClientConfigCallback(httpClientConfigCallback);
 final RestHighLevelClient client = new RestHighLevelClient(builder);
 logger.info("opensearch rest client build success {} ", client);

 ClusterHealthRequest request = new ClusterHealthRequest();
 ClusterHealthResponse response = client.cluster().health(request, RequestOptions.DEFAULT);
 logger.info("opensearch rest client health response {} ", response);
 return client;
 }

 /**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
 public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
 }

 /**
 * Configure trustAllCerts to ignore the certificate configuration.
 */
 public 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;
 }
 }
};

private static final Logger logger =
LogManager.getLogger(RestHighLevelClientExampleHttpWithSecurityNoCert.class);

static 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;
 }
}

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



```
 }
}

/**
 * The following is an example of the main function. Call the create function to create a client and check
whether the test index exists.
 */
public static void main(String[] args) throws IOException {
 RestHighLevelClient client = create(Arrays.asList("{host}"), 9200, "https", 30000, 30000, 30000,
"username", "password");
 GetIndexRequest indexRequest = new GetIndexRequest("test");
 boolean exists = client.indices().exists(indexRequest, RequestOptions.DEFAULT);
 System.out.println(exists);
 client.close();
}
}
```

**Table 3-40** Variables

| Parameter                | Description                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------|
| host                     | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port                     | Access port of the cluster. The default value is <b>9200</b> .                                           |
| protocol                 | Connection protocol, which can be <b>http</b> or <b>https</b> .                                          |
| connectTimeout           | Socket connection timeout (in ms).                                                                       |
| connectionRequestTimeout | Socket connection request timeout (in ms).                                                               |
| socketTimeout            | Socket request timeout (in ms).                                                                          |
| username                 | Username for accessing the cluster.                                                                      |
| password                 | Password of the user.                                                                                    |

This piece of code checks whether the **test** index exists in the cluster. If **true** or **false** is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode Cluster Using the OpenSearch High Level REST Client (With a Certificate)

Use the OpenSearch High Level REST Client to connect to a security-mode OpenSearch cluster that uses HTTPS with a security certificate loaded, and query whether the **test** index exists. The sample code is as follows:

---

### CAUTION

For how to obtain and upload a security certificate, see [Obtaining and Uploading a Security Certificate](#).

---

```
import org.apache.http.HttpHost;
import org.apache.http.auth.AuthScope;
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.client.CredentialsProvider;
import org.apache.http.conn.ssl.NoopHostnameVerifier;
import org.apache.http.impl.client.BasicCredentialsProvider;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOSessionStrategy;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.opensearch.action.admin.cluster.health.ClusterHealthRequest;
import org.opensearch.action.admin.cluster.health.ClusterHealthResponse;
import org.opensearch.client.RequestOptions;
import org.opensearch.client.RestClient;
import org.opensearch.client.RestClientBuilder;
import org.opensearch.client.RestHighLevelClient;
import org.opensearch.client.indices.GetIndexRequest;
import org.opensearch.common.Nullable;
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 java.util.Arrays;
import java.util.List;
import java.util.Objects;
import javax.net.ssl.SSLContext;
import javax.net.ssl.TrustManager;
import javax.net.ssl.TrustManagerFactory;
import javax.net.ssl.X509TrustManager;
/**
 * Use Hive Level REST Client to connect to a security-mode cluster (with an HTTPS certificate).
 */
public class RestHighLevelClientExample {
 public static RestHighLevelClient create(List<String> host, int port, String protocol, int connectTimeout,
int connectionRequestTimeout, int socketTimeout, String username, String password, String certFilePath,
String certPassword) throws IOException {

 final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
 credentialsProvider.setCredentials(AuthScope.ANY, new UsernamePasswordCredentials(username,
password));
 SSLContext sc = null;
 try {
 TrustManager[] tm = {new MyX509TrustManager(certFilePath, certPassword)};
 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 NoopHostnameVerifier());
 SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,
credentialsProvider);

 RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, port, protocol))
 .setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout)
 .setConnectionRequestTimeout(connectionRequestTimeout)
 .setSocketTimeout(socketTimeout))
 .setHttpClientConfigCallback(httpClientConfigCallback);
 final RestHighLevelClient client = new RestHighLevelClient(builder);
 logger.info("opensearch rest client build success {} ", client);

 ClusterHealthRequest request = new ClusterHealthRequest();
 ClusterHealthResponse response = client.cluster().health(request, RequestOptions.DEFAULT);
 }
}
```

```
 logger.info("opensearch rest client health response {} ", response);
 return client;
 }

 /**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
 public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
 }

 /**
 * SecuredHttpClientConfigCallback class definition
 */
 static class SecuredHttpClientConfigCallback implements RestClientBuilder.HttpClientConfigCallback {
 @Nullable
 private final CredentialsProvider credentialsProvider;

 private final SSLIOSessionStrategy sslStrategy;

 SecuredHttpClientConfigCallback(final SSLIOSessionStrategy sslStrategy,
 @Nullable final CredentialsProvider credentialsProvider) {
 this.sslStrategy = Objects.requireNonNull(sslStrategy);
 this.credentialsProvider = credentialsProvider;
 }

 @Nullable
 CredentialsProvider getCredentialsProvider() {
 return credentialsProvider;
 }

 SSLIOSessionStrategy getSSLStrategy() {
 return sslStrategy;
 }

 @Override
 public HttpAsyncClientBuilder customizeHttpClient(final HttpAsyncClientBuilder httpClientBuilder) {
 httpClientBuilder.setSSLStrategy(sslStrategy);
 if (credentialsProvider != null) {
 httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
 }
 return httpClientBuilder;
 }
 }

 private static final Logger logger = LogManager.getLogger(RestHighLevelClient.class);

 public static class MyX509TrustManager implements X509TrustManager {
 X509TrustManager sunJSSEX509TrustManager;

 MyX509TrustManager(String certFilePath, String certPassword) throws Exception {
 File file = new File(certFilePath);
 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, certPassword.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];
}
}

/**
 * The following is an example of the main function. Call the create function to create a client and check
 * whether the test index exists.
 */
public static void main(String[] args) throws IOException {
 RestHighLevelClient client = create(Arrays.asList("{host}"), 9200, "https", 30000, 30000, 30000,
 "username", "password", "certFilePath", "certPassword");
 GetIndexRequest indexRequest = new GetIndexRequest("test");
 boolean exists = client.indices().exists(indexRequest, RequestOptions.DEFAULT);
 System.out.println(exists);
 client.close();
}
}
```

**Table 3-41** Function parameters

| Parameter                | Description                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------|
| host                     | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port                     | Access port of the cluster. The default value is <b>9200</b> .                                           |
| protocol                 | Connection protocol. Set this parameter to <b>https</b> .                                                |
| connectTimeout           | Socket connection timeout (in ms).                                                                       |
| connectionRequestTimeout | Socket connection request timeout (in ms).                                                               |
| socketTimeout            | Socket request timeout (in ms).                                                                          |
| username                 | Username for accessing the cluster.                                                                      |
| password                 | Password of the user.                                                                                    |
| certFilePath             | Path for storing the security certificate.                                                               |
| certPassword             | Password of the security certificate.                                                                    |

This piece of code checks whether the **test** index exists in the cluster. If **true** or **false** is returned, it indicates that the cluster is connected.

## Connecting to a Non-Security Mode Cluster Using the Elasticsearch High Level REST Client

Use the Elasticsearch 7.10.2 High Level REST Client to connect to an OpenSearch cluster for which the security mode is disabled, and query whether the **test** index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
import org.elasticsearch.client.RequestOptions;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;
import org.elasticsearch.client.indices.GetIndexRequest;
import java.io.IOException;
import java.util.Arrays;
import java.util.List;
public class RestHighLevelClientExample {
 public static void main(String[] args) throws IOException {
 List<String> host = Arrays.asList("{Cluster access address}");
 RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, 9200, "http"));
 final RestHighLevelClient client = new RestHighLevelClient(builder);
 GetIndexRequest indexRequest = new GetIndexRequest("test");
 boolean exists = client.indices().exists(indexRequest, RequestOptions.DEFAULT);
 System.out.println(exists);
 client.close();
 }
 /**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
 public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
 }
}
```

This piece of code checks whether the **test** index exists in the cluster. If **true** or **false** is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode Cluster Using the Elasticsearch High Level REST Client (Without a Certificate)

Use the Elasticsearch 7.10.2 High Level REST Client to connect to a security-mode OpenSearch cluster (HTTP or HTTPS) without loading a security certificate, and query whether the **test** index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
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.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.elasticsearch.action.admin.cluster.health.ClusterHealthRequest;
import org.elasticsearch.action.admin.cluster.health.ClusterHealthResponse;
import org.elasticsearch.client.RequestOptions;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;
import org.elasticsearch.client.indices.GetIndexRequest;
import org.elasticsearch.common.Nullable;
```

```
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 java.util.Arrays;
import java.util.List;
import java.util.Objects;

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;

/**
 * Connect to a security-model cluster using the High Level REST Client (without a certificate).
 */
public class RestHighLevelClientExample {
 /**
 * Create a class for the client. Define the create function.
 */
 public static RestHighLevelClient create(List<String> host, int port, String protocol, int connectTimeout,
int connectionRequestTimeout, int socketTimeout, String username, String password) throws IOException{
 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);
 RestClient.Builder builder = RestClient.builder(constructHttpHosts(host, port, protocol))
 .setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout)
 .setConnectionRequestTimeout(connectionRequestTimeout)
 .setSocketTimeout(socketTimeout))
 .setHttpClientConfigCallback(httpClientConfigCallback);
 final RestHighLevelClient client = new RestHighLevelClient(builder);
 logger.info("es rest client build success {} ", client);

 ClusterHealthRequest request = new ClusterHealthRequest();
 ClusterHealthResponse response = client.cluster().health(request, RequestOptions.DEFAULT);
 logger.info("es rest client health response {} ", response);
 return client;
 }

 /**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
 public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
 }

 /**
 * Configure trustAllCerts to ignore the certificate configuration.
 */
 public 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;
 }
 }
};

private static final Logger logger =
LogManager.getLogger(RestHighLevelClientExampleHttpWithSecurityNoCert.class);

static 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;
 }
}
```

```
}

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

/**
 * The following is an example of the main function. Call the create function to create a client and check
 * whether the test index exists.
 */
public static void main(String[] args) throws IOException {
 RestHighLevelClient client = create(Arrays.asList("{host}"), 9200, "https", 30000, 30000, 30000,
 "username", "password");
 GetIndexRequest indexRequest = new GetIndexRequest("test");
 boolean exists = client.indices().exists(indexRequest, RequestOptions.DEFAULT);
 System.out.println(exists);
 client.close();
}
}
```

**Table 3-42** Variables

| Parameter                | Description                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------|
| host                     | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port                     | Access port of the cluster. The default value is <b>9200</b> .                                           |
| protocol                 | Connection protocol, which can be <b>http</b> or <b>https</b> .                                          |
| connectTimeout           | Socket connection timeout (in ms).                                                                       |
| connectionRequestTimeout | Socket connection request timeout (in ms).                                                               |
| socketTimeout            | Socket request timeout (in ms).                                                                          |
| username                 | Username for accessing the cluster.                                                                      |
| password                 | Password of the user.                                                                                    |

This piece of code checks whether the **test** index exists in the cluster. If **true** or **false** is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode Cluster Using the Elasticsearch High Level REST Client (With a Certificate)

Use the Elasticsearch 7.10.2 High Level REST Client to connect to a security-mode OpenSearch cluster that uses HTTPS with a security certificate loaded, and query whether the **test** index exists. The sample code is as follows:



 **CAUTION**

For how to obtain and upload a security certificate, see [Obtaining and Uploading a Security Certificate](#).

```
import org.apache.http.HttpHost;
import org.apache.http.auth.AuthScope;
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.client.CredentialsProvider;
import org.apache.http.conn.ssl.NoopHostnameVerifier;
import org.apache.http.impl.client.BasicCredentialsProvider;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOSessionStrategy;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.elasticsearch.action.admin.cluster.health.ClusterHealthRequest;
import org.elasticsearch.action.admin.cluster.health.ClusterHealthResponse;
import org.elasticsearch.client.RequestOptions;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;
import org.elasticsearch.client.indices.GetIndexRequest;
import org.elasticsearch.common.Nullable;

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 java.util.Arrays;
import java.util.List;
import java.util.Objects;

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

/**
 * Use Hive Level REST Client to connect to a security-mode cluster (with an HTTPS certificate).
 */
public class RestHighLevelClientExample {
 public static RestHighLevelClient create(List<String> host, int port, String protocol, int connectTimeout,
int connectionRequestTimeout, int socketTimeout, String username, String password, String certFilePath,
String certPassword) throws IOException {

 final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
 credentialsProvider.setCredentials(AuthScope.ANY, new UsernamePasswordCredentials(username,
password));
 SSLContext sc = null;
 try {
 TrustManager[] tm = {new MyX509TrustManager(certFilePath, certPassword)};
 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 NoopHostnameVerifier());
 SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,
credentialsProvider);

 RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, port, protocol))
```

```
.setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout)
 .setConnectionRequestTimeout(connectionRequestTimeout)
 .setSocketTimeout(socketTimeout))
.setHttpClientConfigCallback(httpClientConfigCallback);
final RestHighLevelClient client = new RestHighLevelClient(builder);
logger.info("es rest client build success {}", client);

ClusterHealthRequest request = new ClusterHealthRequest();
ClusterHealthResponse response = client.cluster().health(request, RequestOptions.DEFAULT);
logger.info("es rest client health response {}", response);
return client;
}

/**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
}

/**
 * SecuredHttpClientConfigCallback class definition
 */
static class SecuredHttpClientConfigCallback implements RestClientBuilder.HttpClientConfigCallback {
 @Nullable
 private final CredentialsProvider credentialsProvider;

 private final SSLIOSessionStrategy sslStrategy;

 SecuredHttpClientConfigCallback(final SSLIOSessionStrategy sslStrategy,
 @Nullable final CredentialsProvider credentialsProvider) {
 this.sslStrategy = Objects.requireNonNull(sslStrategy);
 this.credentialsProvider = credentialsProvider;
 }

 @Nullable
 CredentialsProvider getCredentialsProvider() {
 return credentialsProvider;
 }

 SSLIOSessionStrategy getSSLStrategy() {
 return sslStrategy;
 }

 @Override
 public HttpAsyncClientBuilder customizeHttpClient(final HttpAsyncClientBuilder httpClientBuilder) {
 httpClientBuilder.setSSLStrategy(sslStrategy);
 if (credentialsProvider != null) {
 httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
 }
 return httpClientBuilder;
 }
}

private static final Logger logger = LogManager.getLogger(RestHighLevelClient.class);

public static class MyX509TrustManager implements X509TrustManager {
 X509TrustManager sunJSSEX509TrustManager;

 MyX509TrustManager(String certFilePath, String certPassword) throws Exception {
 File file = new File(certFilePath);
 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, certPassword.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];
}
}

/**
 * The following is an example of the main function. Call the create function to create a client and check
 * whether the test index exists.
 */
public static void main(String[] args) throws IOException {
 RestHighLevelClient client = create(Arrays.asList("{host}"), 9200, "https", 30000, 30000, 30000,
 "username", "password", "certFilePath", "certPassword");
 GetIndexRequest indexRequest = new GetIndexRequest("test");
 boolean exists = client.indices().exists(indexRequest, RequestOptions.DEFAULT);
 System.out.println(exists);
 client.close();
}
}
```

**Table 3-43** Function parameters

| Parameter                | Description                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------|
| host                     | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port                     | Access port of the cluster. The default value is <b>9200</b> .                                           |
| protocol                 | Connection protocol. Set this parameter to <b>https</b> .                                                |
| connectTimeout           | Socket connection timeout (in ms).                                                                       |
| connectionRequestTimeout | Socket connection request timeout (in ms).                                                               |
| socketTimeout            | Socket request timeout (in ms).                                                                          |
| username                 | Username for accessing the cluster.                                                                      |
| password                 | Password of the user.                                                                                    |

| Parameter    | Description                                |
|--------------|--------------------------------------------|
| certFilePath | Path for storing the security certificate. |
| certPassword | Password of the security certificate.      |

This piece of code checks whether the **test** index exists in the cluster. If **true** or **false** is returned, it indicates that the cluster is connected.

## Obtaining and Uploading a Security Certificate

To access a security-mode OpenSearch cluster that uses HTTPS, a security certificate must be loaded. Perform the following steps to obtain the security certificate and upload it to the client:

1. Obtain the security certificate **CloudSearchService.cer**.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > OpenSearch**.
  - c. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  - d. Click the **Overview** tab. In the **Configuration** area, click **Download Certificate** next to **HTTPS Access**.

**Figure 3-5** Downloading a security certificate



2. Convert the security certificate **CloudSearchService.cer**. Upload the downloaded security certificate to the client and use keytool to convert the **.cer** certificate into a **.jks** certificate that can be read by Java.
  - In Linux, run the following command to convert the certificate:  
`keytool -import -alias newname -keystore ./truststore.jks -file ./CloudSearchService.cer`
  - In Windows, run the following command to convert the certificate:  
`keytool -import -alias newname -keystore .\truststore.jks -file .\CloudSearchService.cer`

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

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

### 3.5.5.2 Accessing an OpenSearch Cluster Using the Low Level REST Client

You can query and manage data in a CSS OpenSearch cluster using the Low Level REST Client. The Low Level REST Client encapsulates OpenSearch APIs. You only need to construct the required structures to access an OpenSearch cluster. This simplifies the process of working with an OpenSearch cluster. The Low Level REST Client allows you to customize the request structure, which is more flexible and

supports all the request formats of OpenSearch, such as GET, POST, DELETE, and HEAD.

You can access an OpenSearch cluster with the Low Level REST Client using either of the following ways: directly creating the Low Level REST Client; or creating the High Level REST Client first and then calling **RestHighLevelClient.getLowLevelClient()** to obtain the Low Level REST Client.

## Prerequisites

- The target OpenSearch cluster is available.
- The server that runs the Java code can communicate with the OpenSearch cluster.
- Depending on the network configuration method used, obtain the cluster access address. For details, see [Network Configuration](#).
- Java has been installed on the server and the JDK version is 11 or later. Download JDK 11 from [Java Archive Downloads](#).

## Introducing Dependencies

Introduce the required Java dependencies on the server where you run Java code.

CSS allows you to use the Elasticsearch 7.10 Java client to connect to OpenSearch clusters. To ensure better compatibility, however, you are advised to use an OpenSearch Java client that has the same version as the target OpenSearch cluster.

Select a reference example based on the Java client you use.

- Scenario 1 (recommended): Use the OpenSearch Java client to access the OpenSearch cluster and introduce the Apache version using Maven.

Replace 1.3.6 with the actual version number.

```
<dependency>
 <groupId>org.opensearch.client</groupId>
 <artifactId>opensearch-rest-high-level-client</artifactId>
 <version>1.3.6</version>
</dependency>
```

- Scenario 2: Use the Elasticsearch 7.10.2 Java client to access the OpenSearch cluster and introduce the Apache version using Maven.

```
<dependency>
 <groupId>org.elasticsearch.client</groupId>
 <artifactId>elasticsearch-rest-client</artifactId>
 <version>7.10.2</version>
</dependency>
<dependency>
 <groupId>org.elasticsearch</groupId>
 <artifactId>elasticsearch</artifactId>
 <version>7.10.2</version>
</dependency>
```

## Accessing a Cluster

The sample code varies depending on the Java client and the security mode settings of the target OpenSearch cluster. Select the right reference document based on your service scenario.

**Table 3-44** Cluster access scenarios

| Java Client          | How Low Level REST Client Is Created                                                                                       | OpenSearch Cluster Security-Mode Settings     | Security Certificate | Details                                                                                                                   |
|----------------------|----------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|----------------------|---------------------------------------------------------------------------------------------------------------------------|
| OpenSearch           | Directly create the Low Level REST Client                                                                                  | Non-security mode                             | -                    | <a href="#">Connecting to a Non-Security Mode Cluster Using the OpenSearch Low Level REST Client</a>                      |
|                      |                                                                                                                            | Security mode + HTTP<br>Security mode + HTTPS | No                   | <a href="#">Connecting to a Security-Mode Cluster Using the OpenSearch Low Level REST Client (Without a Certificate)</a>  |
|                      |                                                                                                                            | Security mode + HTTPS                         | Yes                  | <a href="#">Connecting to a Security-Mode Cluster Using the OpenSearch Low Level REST Client (With a Certificate)</a>     |
|                      | Create the High Level REST Client first and then call <code>getLowLevelClient()</code> to obtain the Low Level REST Client | Non-security mode                             | -                    | <a href="#">Connecting to a Non-Security Mode Cluster Using the OpenSearch High Level REST Client</a>                     |
|                      |                                                                                                                            | Security mode + HTTP<br>Security mode + HTTPS | No                   | <a href="#">Connecting to a Security-Mode Cluster Using the OpenSearch High Level REST Client (Without a Certificate)</a> |
|                      |                                                                                                                            | Security mode + HTTPS                         | Yes                  | <a href="#">Connecting to a Security-Mode Cluster Using the OpenSearch High Level REST Client (With a Certificate)</a>    |
| Elasticsearch 7.10.2 | Directly create the Low Level REST Client                                                                                  | Non-security mode                             | -                    | <a href="#">Connecting to a Non-Security Mode Cluster Using the Elasticsearch Low Level REST Client</a>                   |

| Java Client | How Low Level REST Client Is Created                                                                                       | OpenSearch Cluster Security-Mode Settings | Security Certificate | Details                                                                                                                      |
|-------------|----------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|----------------------|------------------------------------------------------------------------------------------------------------------------------|
|             |                                                                                                                            | Security mode + HTTP                      | No                   | <a href="#">Connecting to a Security-Mode Cluster Using the Elasticsearch Low Level REST Client (Without a Certificate)</a>  |
|             |                                                                                                                            | Security mode + HTTPS                     | Yes                  | <a href="#">Connecting to a Security-Mode Cluster Using the Elasticsearch Low Level REST Client (With a Certificate)</a>     |
|             | Create the High Level REST Client first and then call <code>getLowLevelClient()</code> to obtain the Low Level REST Client | Non-security mode                         | -                    | <a href="#">Connecting to a Non-Security Mode Cluster Using the Elasticsearch High Level REST Client</a>                     |
|             |                                                                                                                            | Security mode + HTTP                      | No                   | <a href="#">Connecting to a Security-Mode Cluster Using the Elasticsearch High Level REST Client (Without a Certificate)</a> |
|             |                                                                                                                            | Security mode + HTTPS                     | Yes                  | <a href="#">Connecting to a Security-Mode Cluster Using the Elasticsearch High Level REST Client (With a Certificate)</a>    |

## Connecting to a Non-Security Mode Cluster Using the OpenSearch Low Level REST Client

Use the OpenSearch Low Level REST Client to connect to an OpenSearch cluster for which the security mode is disabled, and query whether the **test** index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
import org.opensearch.client.Request;
import org.opensearch.client.Response;
import org.opensearch.client.RestClient;
import org.opensearch.client.RestClientBuilder;

import java.io.IOException;
import java.util.Arrays;
import java.util.List;
public class RestLowLevelClientExample {
 public static void main(String[] args) throws IOException {
```

```
List<String> host = Arrays.asList("{Cluster access address}");
RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, 9200, "http"));
/**
 * Create the Low Level Rest Client.
 */
RestClient lowLevelClient = builder.build();
/**
 * Check whether the test index exists. If the index exists, 200 is returned. If the index does not exist,
404 is returned.
 */
Request request = new Request("HEAD", "/test");
Response response = lowLevelClient.performRequest(request);
System.out.println(response.getStatusLine().getStatusCode());
lowLevelClient.close();
}

/**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
}
}
```

This piece of code checks whether the **test** index exists in the cluster. If 200 or 404 is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode Cluster Using the OpenSearch Low Level REST Client (Without a Certificate)

Use the OpenSearch Low Level REST Client to connect to a security-mode OpenSearch cluster (HTTP or HTTPS) without loading a security certificate, and query whether the **test** index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
import org.apache.http.HttpResponse;
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.impl.client.DefaultConnectionKeepAliveStrategy;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.apache.http.protocol.HttpContext;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.opensearch.client.Request;
import org.opensearch.client.Response;
import org.opensearch.client.RestClient;
import org.opensearch.client.RestClientBuilder;
import org.opensearch.common.Nullable;
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 java.util.Arrays;
import java.util.List;
import java.util.Objects;
import java.util.concurrent.TimeUnit;
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 RestLowLevelClientExample {
```



```
/**
 * Create a class for the client. Define the create function.
 */
public static RestClient create(List<String> host, int port, String protocol, int connectTimeout, int
connectionRequestTimeout, int socketTimeout, String username, String password) throws IOException {
 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(constructHttpHosts(host, port, protocol))
 .setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout)
 .setConnectionRequestTimeout(connectionRequestTimeout)
 .setSocketTimeout(socketTimeout))
 .setHttpClientConfigCallback(httpClientConfigCallback);
 final RestClient client = builder.build();
 logger.info("opensearch rest client build success {} ", client);
 return client;
}

/**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
}

/**
 * Configure trustAllCerts to ignore the certificate configuration.
 */
public 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;
 }
 }
};

/**
 * The CustomConnectionKeepAliveStrategy function is used to set the connection keepalive during when
there are a large number of short connections or when there are not many data requests.
 */
public static class CustomConnectionKeepAliveStrategy extends DefaultConnectionKeepAliveStrategy {
 public static final CustomConnectionKeepAliveStrategy INSTANCE = new
CustomConnectionKeepAliveStrategy();

 private CustomConnectionKeepAliveStrategy() {
 super();
 }
}
```

```
}

/**
 * Maximum keepalive time (in minutes)
 * The default value is 10 minutes. You can set it based on the number of TCP connections in
TIME_WAIT state. If there are too many TCP connections, you can increase this value.
 */
private final long MAX_KEEP_ALIVE_MINUTES = 10;

@Override
public long getKeepAliveDuration(HttpResponse response, HttpContext context) {
 long keepAliveDuration = super.getKeepAliveDuration(response, context);
 // <0 indicates that the keepalive period is unlimited.
 // Change the period from unlimited to a default period.
 if (keepAliveDuration < 0) {
 return TimeUnit.MINUTES.toMillis(MAX_KEEP_ALIVE_MINUTES);
 }
 return keepAliveDuration;
}
}

private static final Logger logger = LogManager.getLogger(RestLowLevelClientExample.class);

static 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;
 }
}

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

/**
 * The following is an example of the main function. Call the create function to create a Low Level REST
 * Client and check whether the test index exists.
 */
public static void main(String[] args) throws IOException {
 RestClient lowLevelClient = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200, "http",
30000, 30000, 30000, "username", "password");
 Request request = new Request("HEAD", "/test");
 Response response = lowLevelClient.performRequest(request);
 System.out.println(response.getStatusLine().getStatusCode());
 lowLevelClient.close();
}
}
```

**Table 3-45** Variables

| Parameter                | Description                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------|
| host                     | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port                     | Access port of the cluster. The default value is <b>9200</b> .                                           |
| protocol                 | Connection protocol, which can be <b>http</b> or <b>https</b> .                                          |
| connectTimeout           | Socket connection timeout (in ms).                                                                       |
| connectionRequestTimeout | Socket connection request timeout (in ms).                                                               |
| socketTimeout            | Socket request timeout (in ms).                                                                          |
| username                 | Username for accessing the cluster.                                                                      |
| password                 | Password of the user.                                                                                    |

This piece of code checks whether the **test** index exists in the cluster. If 200 or 404 is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode Cluster Using the OpenSearch Low Level REST Client (With a Certificate)

Use the OpenSearch Low Level REST Client to connect to a security-mode OpenSearch cluster that uses HTTPS with a security certificate loaded, and query whether the **test** index exists. The sample code is as follows:

---

### CAUTION

For how to obtain and upload a security certificate, see [Obtaining and Uploading a Security Certificate](#).

---

```
import org.apache.http.HttpHost;
import org.apache.http.auth.AuthScope;
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.client.CredentialsProvider;
import org.apache.http.conn.ssl.NoopHostnameVerifier;
import org.apache.http.impl.client.BasicCredentialsProvider;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.opensearch.client.Request;
import org.opensearch.client.Response;
import org.opensearch.client.RestClient;
import org.opensearch.client.RestClientBuilder;
import org.opensearch.common.Nullable;

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 java.util.Arrays;
import java.util.List;
import java.util.Objects;

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

public class RestLowLevelClientExample {

 private static final Logger logger = LogManager.getLogger(RestLowLevelClientExample.class);

 /**
 * Create a class for the client. Define the create function.
 */
 public static RestClient create(List<String> host, int port, String protocol, int connectTimeout, int
connectionRequestTimeout, int socketTimeout, String username, String password, String certFilePath, String
certPassword) throws IOException {
 final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
 credentialsProvider.setCredentials(AuthScope.ANY, new UsernamePasswordCredentials(username,
password));
 SSLContext sc = null;
 try {
 TrustManager[] tm = {new MyX509TrustManager(certFilePath, certPassword)};
 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 NoopHostnameVerifier());
SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,
credentialsProvider);

RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, port, protocol))
.setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout)
.setConnectionRequestTimeout(connectionRequestTimeout)
.setSocketTimeout(socketTimeout))
.setHttpClientConfigCallback(httpClientConfigCallback);
final RestClient client = builder.build();
logger.info("opensearch rest client build success {} ", client);
return client;
}

/**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);}

static class SecuredHttpClientConfigCallback implements RestClientBuilder.HttpClientConfigCallback {
@Nullable
private final CredentialsProvider credentialsProvider;

private final SSLIOSessionStrategy sslStrategy;

SecuredHttpClientConfigCallback(final SSLIOSessionStrategy sslStrategy,
@Nullable final CredentialsProvider credentialsProvider) {
this.sslStrategy = Objects.requireNonNull(sslStrategy);
this.credentialsProvider = credentialsProvider;
}

@Nullable
CredentialsProvider getCredentialsProvider() {
return credentialsProvider;
}

SSLIOSessionStrategy getSSLStrategy() {
return sslStrategy;
}

@Override
public HttpAsyncClientBuilder customizeHttpClient(final HttpAsyncClientBuilder httpClientBuilder) {
httpClientBuilder.setSSLStrategy(sslStrategy);
if (credentialsProvider != null) {
httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
}
return httpClientBuilder;
}}

public static class MyX509TrustManager implements X509TrustManager {
X509TrustManager sunJSSEX509TrustManager;

MyX509TrustManager(String certFilePath, String certPassword) throws Exception {
File file = new File(certFilePath);
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, certPassword.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];
 }
}

/**
 * The following is an example of the main function. Call the create function to create a Low Level REST
 * Client and check whether the test index exists.
 */
public static void main(String[] args) throws IOException {
 RestClient lowLevelClient = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200, "https",
30000, 30000, 30000, "username", "password", "certFilePath", "certPassword");
 Request request = new Request("HEAD", "test");
 Response response = lowLevelClient.performRequest(request);
 System.out.println(response.getStatusLine().getStatusCode());
 lowLevelClient.close();
}
```

**Table 3-46** Function parameters

| Parameter                | Description                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------|
| host                     | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port                     | Access port of the cluster. The default value is <b>9200</b> .                                           |
| protocol                 | Connection protocol. Set this parameter to <b>https</b> .                                                |
| connectTimeout           | Socket connection timeout (in ms).                                                                       |
| connectionRequestTimeout | Socket connection request timeout (in ms).                                                               |
| socketTimeout            | Socket request timeout (in ms).                                                                          |
| username                 | Username for accessing the cluster.                                                                      |
| password                 | Password of the user.                                                                                    |

| Parameter    | Description                                |
|--------------|--------------------------------------------|
| certFilePath | Path for storing the security certificate. |
| certPassword | Password of the security certificate.      |

This piece of code checks whether the **test** index exists in the cluster. If 200 or 404 is returned, it indicates that the cluster is connected.

## Connecting to a Non-Security Mode Cluster Using the OpenSearch High Level REST Client

Use the OpenSearch High Level REST Client to obtain the Low Level REST Client by calling **getLowLevelClient()**, use the low-level client to connect to an OpenSearch cluster for which the security mode is disabled, and query whether the **test** index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
import org.opensearch.client.Request;
import org.opensearch.client.Response;
import org.opensearch.client.RestClient;
import org.opensearch.client.RestClientBuilder;
import org.opensearch.client.RestHighLevelClient;

import java.io.IOException;
import java.util.Arrays;
import java.util.List;
public class RestLowLevelClientExample {

 public static void main(String[] args) throws IOException {
 List<String> host = Arrays.asList("{Cluster access address}");
 RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, 9200, "http"));
 final RestHighLevelClient restHighLevelClient = new RestHighLevelClient(builder);
 /**
 * Create a high-level client and then call getLowLevelClient() to obtain a low-level client. The code
 * differs from the client creation code only in the following line:
 */
 final RestClient lowLevelClient = restHighLevelClient.getLowLevelClient();
 /**
 * Check whether the test index exists. If the index exists, 200 is returned. If the index does not exist,
 * 404 is returned.
 */
 Request request = new Request("HEAD", "/test");
 Response response = lowLevelClient.performRequest(request);
 System.out.println(response.getStatusLine().getStatusCode());
 lowLevelClient.close();
 }

 /**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
 public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
 }
}
```

This piece of code checks whether the **test** index exists in the cluster. If 200 or 404 is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode Cluster Using the OpenSearch High Level REST Client (Without a Certificate)

Use the OpenSearch High Level REST Client to obtain the Low Level REST Client by calling `getLowLevelClient()`, use the low-level client to connect to a security-mode OpenSearch cluster that uses HTTP or HTTPS without loading a security certificate, and query whether the `test` index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
import org.apache.http.HttpResponse;
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.impl.client.DefaultConnectionKeepAliveStrategy;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOSessionStrategy;
import org.apache.http.protocol.HttpContext;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.opensearch.client.Request;
import org.opensearch.client.Response;
import org.opensearch.client.RestClient;
import org.opensearch.client.RestClientBuilder;
import org.opensearch.client.RestHighLevelClient;
import org.opensearch.common.Nullable;

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 java.util.Arrays;
import java.util.List;
import java.util.Objects;
import java.util.concurrent.TimeUnit;

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 RestLowLevelClientExample {

 /**
 * Create a class for the client. Define the create function.
 */
 public static RestHighLevelClient create(List<String> host, int port, String protocol, int connectTimeout,
int connectionRequestTimeout, int socketTimeout, String username, String password) throws IOException {

 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(constructHttpHosts(host, port, protocol))
.setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout))
```



```
.setConnectionRequestTimeout(connectionRequestTimeout)
.setSocketTimeout(socketTimeout))
.setHttpClientConfigCallback(httpClientConfigCallback);
final RestHighLevelClient client = new RestHighLevelClient(builder);
logger.info("opensearch rest client build success {}", client);
return client;
}

/**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
}

/**
 * Configure trustAllCerts to ignore the certificate configuration.
 */
public 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;
 }
 }
};

/**
 * The CustomConnectionKeepAliveStrategy function is used to set the connection keepalive during when
there are a large number of short connections or when there are not many data requests.
 */
public static class CustomConnectionKeepAliveStrategy extends DefaultConnectionKeepAliveStrategy {
 public static final CustomConnectionKeepAliveStrategy INSTANCE = new
CustomConnectionKeepAliveStrategy();

 private CustomConnectionKeepAliveStrategy() {
 super();
 }

 /**
 * Maximum keepalive time (in minutes)
 * The default value is 10 minutes. You can set it based on the number of TCP connections in
TIME_WAIT state. If there are too many TCP connections, you can increase this value.
 */
 private final long MAX_KEEP_ALIVE_MINUTES = 10;

 @Override
 public long getKeepAliveDuration(HttpResponse response, HttpContext context) {
 long keepAliveDuration = super.getKeepAliveDuration(response, context);
 // <0 indicates an unlimited keepalive period.
 // Change the period from unlimited to a default period.
 if (keepAliveDuration < 0) {
 return TimeUnit.MINUTES.toMillis(MAX_KEEP_ALIVE_MINUTES);
 }
 return keepAliveDuration;
 }
}

private static final Logger logger = LogManager.getLogger(RestLowLevelClientExample.class);
```

```
static 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;
 }
}

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

/**
 * The following is an example of the main function. Call the create function to create the high-level client,
 * then call the getLowLevelClient() function to obtain the low-level client, and query whether the test index
 * exists.
 */
public static void main(String[] args) throws IOException {
```

```
RestHighLevelClient client = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200, "http",
30000, 30000, 30000, "username", "password");
RestClient lowLevelClient = client.getLowLevelClient();
Request request = new Request("HEAD", "test");
Response response = lowLevelClient.performRequest(request);
System.out.println(response.getStatusLine().getStatusCode());
lowLevelClient.close();
}
```

**Table 3-47** Variables

| Parameter                | Description                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------|
| host                     | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port                     | Access port of the cluster. The default value is <b>9200</b> .                                           |
| protocol                 | Connection protocol, which can be <b>http</b> or <b>https</b> .                                          |
| connectTimeout           | Socket connection timeout (in ms).                                                                       |
| connectionRequestTimeout | Socket connection request timeout (in ms).                                                               |
| socketTimeout            | Socket request timeout (in ms).                                                                          |
| username                 | Username for accessing the cluster.                                                                      |
| password                 | Password of the user.                                                                                    |

This piece of code checks whether the **test** index exists in the cluster. If 200 or 404 is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode Cluster Using the OpenSearch High Level REST Client (With a Certificate)

Use the OpenSearch High Level REST Client to connect to a security-mode OpenSearch cluster that uses HTTPS with a security certificate loaded, and query whether the **test** index exists. The sample code is as follows:

### CAUTION

For how to obtain and upload a security certificate, see [Obtaining and Uploading a Security Certificate](#).

```
import org.apache.http.HttpHost;
import org.apache.http.auth.AuthScope;
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.client.CredentialsProvider;
import org.apache.http.conn.ssl.NoopHostnameVerifier;
```

```
import org.apache.http.impl.client.BasicCredentialsProvider;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.opensearch.action.admin.cluster.health.ClusterHealthRequest;
import org.opensearch.action.admin.cluster.health.ClusterHealthResponse;
import org.opensearch.client.Request;
import org.opensearch.client.RequestOptions;
import org.opensearch.client.Response;
import org.opensearch.client.RestClient;
import org.opensearch.client.RestClientBuilder;
import org.opensearch.client.RestHighLevelClient;
import org.opensearch.common.Nullable;
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 java.util.Arrays;
import java.util.List;
import java.util.Objects;
import javax.net.ssl.SSLContext;
import javax.net.ssl.TrustManager;
import javax.net.ssl.TrustManagerFactory;
import javax.net.ssl.X509TrustManager;

public class RestLowLevelClientExample {
 private static final Logger logger = LogManager.getLogger(RestLowLevelClientExample.class);

 /**
 * Create a class for the client. Define the create function.
 */
 public static RestHighLevelClient create(List<String> host, int port, String protocol, int connectTimeout,
int connectionRequestTimeout, int socketTimeout, String username, String password, String certFilePath,
String certPassword) throws IOException {
 final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
 credentialsProvider.setCredentials(AuthScope.ANY, new UsernamePasswordCredentials(username,
password));
 SSLContext sc = null;
 try {
 TrustManager[] tm = {new MyX509TrustManager(certFilePath, certPassword)};
 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();
 }

 SSLIOStrategy sessionStrategy = new SSLIOStrategy(sc, new NoopHostnameVerifier());
 SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,
credentialsProvider);

 RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, port, protocol))
 .setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout))
 .setConnectionRequestTimeout(connectionRequestTimeout)
 .setSocketTimeout(socketTimeout)
 .setHttpClientConfigCallback(httpClientConfigCallback);
 final RestHighLevelClient client = new RestHighLevelClient(builder);
 logger.info("opensearch rest client build success {} ", client);

 ClusterHealthRequest request = new ClusterHealthRequest();
 ClusterHealthResponse response = client.cluster().health(request, RequestOptions.DEFAULT);
 logger.info("opensearch rest client health response {} ", response);
 return client;
 }
}
```

```
}

/**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);}

static class SecuredHttpClientConfigCallback implements RestClientBuilder.HttpClientConfigCallback {
 @Nullable
 private final CredentialsProvider credentialsProvider;

 private final SSLIOSessionStrategy sslStrategy;

 SecuredHttpClientConfigCallback(final SSLIOSessionStrategy sslStrategy,
 @Nullable final CredentialsProvider credentialsProvider) {
 this.sslStrategy = Objects.requireNonNull(sslStrategy);
 this.credentialsProvider = credentialsProvider;
 }

 @Nullable
 CredentialsProvider getCredentialsProvider() {
 return credentialsProvider;
 }

 SSLIOSessionStrategy getSSLStrategy() {
 return sslStrategy;
 }

 @Override
 public HttpAsyncClientBuilder customizeHttpClient(final HttpAsyncClientBuilder httpClientBuilder) {
 httpClientBuilder.setSSLStrategy(sslStrategy);
 if (credentialsProvider != null) {
 httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
 }
 return httpClientBuilder;
 }
}

public static class MyX509TrustManager implements X509TrustManager {
 X509TrustManager sunJSSEX509TrustManager;

 MyX509TrustManager(String certFilePath, String certPassword) throws Exception {
 File file = new File(certFilePath);
 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, certPassword.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];
 }
}

/**
 * The following is an example of the main function. Call the create function to create the high-level client,
 * then call the getLowLevelClient() function to obtain the low-level client, and query whether the test index
 * exists.
 */
public static void main(String[] args) throws IOException {
 RestHighLevelClient client = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200, "https",
30000, 30000, 30000, "username", "password", "certFilePath", "certPassword");
 RestClient lowLevelClient = client.getLowLevelClient();
 Request request = new Request("HEAD", "test");
 Response response = lowLevelClient.performRequest(request);
 System.out.println(response.getStatusLine().getStatusCode());
 lowLevelClient.close();
}
}
```

**Table 3-48** Function parameters

| Parameter                | Description                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------|
| host                     | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port                     | Access port of the cluster. The default value is <b>9200</b> .                                           |
| protocol                 | Connection protocol. Set this parameter to <b>https</b> .                                                |
| connectTimeout           | Socket connection timeout (in ms).                                                                       |
| connectionRequestTimeout | Socket connection request timeout (in ms).                                                               |
| socketTimeout            | Socket request timeout (in ms).                                                                          |
| username                 | Username for accessing the cluster.                                                                      |
| password                 | Password of the user.                                                                                    |
| certFilePath             | Path for storing the security certificate.                                                               |
| certPassword             | Password of the security certificate.                                                                    |

This piece of code checks whether the **test** index exists in the cluster. If 200 or 404 is returned, it indicates that the cluster is connected.

## Connecting to a Non-Security Mode Cluster Using the Elasticsearch Low Level REST Client

Use the Elasticsearch 7.10.2 Low Level REST Client to connect to an OpenSearch cluster for which the security mode is disabled, and query whether the **test** index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
import org.elasticsearch.client.Request;
import org.elasticsearch.client.Response;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;

import java.io.IOException;
import java.util.Arrays;
import java.util.List;
public class RestLowLevelClientExample {
 public static void main(String[] args) throws IOException {
 List<String> host = Arrays.asList("{Cluster access address}");
 RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, 9200, "http"));
 /**
 * Create the Low Level Rest Client.
 */
 RestClient lowLevelClient = builder.build();
 /**
 * Check whether the test index exists. If the index exists, 200 is returned. If the index does not exist,
 404 is returned.
 */
 Request request = new Request("HEAD", "/test");
 Response response = lowLevelClient.performRequest(request);
 System.out.println(response.getStatusLine().getStatusCode());
 lowLevelClient.close();
 }

 /**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
 public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
 }
}
```

This piece of code checks whether the **test** index exists in the cluster. If 200 or 404 is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode Cluster Using the Elasticsearch Low Level REST Client (Without a Certificate)

Use the Elasticsearch 7.10.2 Low Level REST Client to connect to a security-mode OpenSearch cluster (HTTP or HTTPS) without loading a security certificate, and query whether the **test** index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
import org.apache.http.HttpResponse;
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.impl.client.DefaultConnectionKeepAliveStrategy;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.apache.http.protocol.HttpContext;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.elasticsearch.client.Request;
```

```
import org.elasticsearch.client.Response;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.common.Nullable;

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 java.util.Arrays;
import java.util.List;
import java.util.Objects;
import java.util.concurrent.TimeUnit;
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 RestLowLevelClientExample {
 /**
 * Create a class for the client. Define the create function.
 */
 public static RestClient create(List<String> host, int port, String protocol, int connectTimeout, int
connectionRequestTimeout, int socketTimeout, String username, String password) throws IOException {
 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(constructHttpHosts(host, port, protocol))
 .setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout)
 .setConnectionRequestTimeout(connectionRequestTimeout)
 .setSocketTimeout(socketTimeout))
 .setHttpClientConfigCallback(httpClientConfigCallback);
 final RestClient client = builder.build();
 logger.info("es rest client build success {} ", client);
 return client;
 }

 /**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
 public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
 }

 /**
 * Configure trustAllCerts to ignore the certificate configuration.
 */
 public 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;
 }
 }
};

/**
 * The CustomConnectionKeepAliveStrategy function is used to set the connection keepalive during when
 * there are a large number of short connections or when there are not many data requests.
 */
public static class CustomConnectionKeepAliveStrategy extends DefaultConnectionKeepAliveStrategy {
 public static final CustomConnectionKeepAliveStrategy INSTANCE = new
CustomConnectionKeepAliveStrategy();

 private CustomConnectionKeepAliveStrategy() {
 super();
 }

 /**
 * Maximum keepalive time (in minutes)
 * The default value is 10 minutes. You can set it based on the number of TCP connections in
TIME_WAIT state. If there are too many TCP connections, you can increase this value.
 */
 private final long MAX_KEEP_ALIVE_MINUTES = 10;

 @Override
 public long getKeepAliveDuration(HttpResponse response, HttpContext context) {
 long keepAliveDuration = super.getKeepAliveDuration(response, context);
 // <0 indicates an unlimited keepalive period.
 // Change the period from unlimited to a default period.
 if (keepAliveDuration < 0) {
 return TimeUnit.MINUTES.toMillis(MAX_KEEP_ALIVE_MINUTES);
 }
 return keepAliveDuration;
 }
}

private static final Logger logger = LogManager.getLogger(RestLowLevelClientExample.class);

static 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 SSLIOStrategy} that will be added to the HTTP client.
 *
 * @return Never {@code null}.
 */
SSLIOStrategy getSSLStrategy() {
 return sslStrategy;
}
/**
 * Sets the {@linkplain HttpAsyncClientBuilder#setDefaultCredentialsProvider(CredentialsProvider)
credentials 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;
}

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

/**
 * The following is an example of the main function. Call the create function to create a Low Level REST
Client and check whether the test index exists.
 */
public static void main(String[] args) throws IOException {
 RestClient lowLevelClient = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200, "http",
30000, 30000, 30000, "username", "password");
 Request request = new Request("HEAD", "/test");
 Response response = lowLevelClient.performRequest(request);
 System.out.println(response.getStatusLine().getStatusCode());
 lowLevelClient.close();
}
}
```

**Table 3-49** Variables

| Parameter      | Description                                                                                              |
|----------------|----------------------------------------------------------------------------------------------------------|
| host           | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port           | Access port of the cluster. The default value is <b>9200</b> .                                           |
| protocol       | Connection protocol, which can be <b>http</b> or <b>https</b> .                                          |
| connectTimeout | Socket connection timeout (in ms).                                                                       |

| Parameter                | Description                                |
|--------------------------|--------------------------------------------|
| connectionRequestTimeout | Socket connection request timeout (in ms). |
| socketTimeout            | Socket request timeout (in ms).            |
| username                 | Username for accessing the cluster.        |
| password                 | Password of the user.                      |

This piece of code checks whether the **test** index exists in the cluster. If 200 or 404 is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode Cluster Using the Elasticsearch Low Level REST Client (With a Certificate)

Use the Elasticsearch 7.10.2 Low Level REST Client to connect to a security-mode OpenSearch cluster that uses HTTPS with a security certificate loaded, and query whether the **test** index exists. The sample code is as follows:

---

### CAUTION

For how to obtain and upload a security certificate, see [Obtaining and Uploading a Security Certificate](#).

---

```
import org.apache.http.HttpHost;
import org.apache.http.auth.AuthScope;
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.client.CredentialsProvider;
import org.apache.http.conn.ssl.NoopHostnameVerifier;
import org.apache.http.impl.client.BasicCredentialsProvider;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.elasticsearch.client.Request;
import org.elasticsearch.client.Response;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.common.Nullable;

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 java.util.Arrays;
import java.util.List;
import java.util.Objects;

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

```
public class RestLowLevelClientExample {

 private static final Logger logger = LogManager.getLogger(RestLowLevelClientExample.class);

 /**
 * Create a class for the client. Define the create function.
 */
 public static RestClient create(List<String> host, int port, String protocol, int connectTimeout, int
connectionRequestTimeout, int socketTimeout, String username, String password, String certFilePath, String
certPassword) throws IOException {
 final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
 credentialsProvider.setCredentials(AuthScope.ANY, new UsernamePasswordCredentials(username,
password));
 SSLContext sc = null;
 try {
 TrustManager[] tm = {new MyX509TrustManager(certFilePath, certPassword)};
 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 NoopHostnameVerifier());
 SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,
credentialsProvider);

 RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, port, protocol))
 .setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout)
 .setConnectionRequestTimeout(connectionRequestTimeout)
 .setSocketTimeout(socketTimeout))
 .setHttpClientConfigCallback(httpClientConfigCallback);
 final RestClient client = builder.build();
 logger.info("es rest client build success {} ", client);
 return client;
 }

 /**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
 public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
 }

 static class SecuredHttpClientConfigCallback implements RestClientBuilder.HttpClientConfigCallback {
 @Nullable
 private final CredentialsProvider credentialsProvider;

 private final SSLIOSessionStrategy sslStrategy;

 SecuredHttpClientConfigCallback(final SSLIOSessionStrategy sslStrategy,
 @Nullable final CredentialsProvider credentialsProvider) {
 this.sslStrategy = Objects.requireNonNull(sslStrategy);
 this.credentialsProvider = credentialsProvider;
 }

 @Nullable
 CredentialsProvider getCredentialsProvider() {
 return credentialsProvider;
 }

 SSLIOSessionStrategy getSSLStrategy() {
 return sslStrategy;
 }

 @Override
 public HttpAsyncClientBuilder customizeHttpClient(final HttpAsyncClientBuilder httpClientBuilder) {
 httpClientBuilder.setSSLStrategy(sslStrategy);
 if (credentialsProvider != null) {

```

```
 httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
 }
 return httpClientBuilder;
}}

public static class MyX509TrustManager implements X509TrustManager {
 X509TrustManager sunJSSEX509TrustManager;

 MyX509TrustManager(String certFilePath, String certPassword) throws Exception {
 File file = new File(certFilePath);
 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, certPassword.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];
 }
}

/**
 * The following is an example of the main function. Call the create function to create a Low Level REST
 * Client and check whether the test index exists.
 */
public static void main(String[] args) throws IOException {
 RestClient lowLevelClient = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200, "https",
30000, 30000, 30000, "username", "password", "certFilePath", "certPassword");
 Request request = new Request("HEAD", "test");
 Response response = lowLevelClient.performRequest(request);
 System.out.println(response.getStatusLine().getStatusCode());
 lowLevelClient.close();
}
}
```

**Table 3-50** Function parameters

| Parameter | Description                                                                                              |
|-----------|----------------------------------------------------------------------------------------------------------|
| host      | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |

| Parameter                | Description                                                    |
|--------------------------|----------------------------------------------------------------|
| port                     | Access port of the cluster. The default value is <b>9200</b> . |
| protocol                 | Connection protocol. Set this parameter to <b>https</b> .      |
| connectTimeout           | Socket connection timeout (in ms).                             |
| connectionRequestTimeout | Socket connection request timeout (in ms).                     |
| socketTimeout            | Socket request timeout (in ms).                                |
| username                 | Username for accessing the cluster.                            |
| password                 | Password of the user.                                          |
| certFilePath             | Path for storing the security certificate.                     |
| certPassword             | Password of the security certificate.                          |

This piece of code checks whether the **test** index exists in the cluster. If 200 or 404 is returned, it indicates that the cluster is connected.

## Connecting to a Non-Security Mode Cluster Using the Elasticsearch High Level REST Client

Use the Elasticsearch 7.10.2 High Level REST Client obtain the Low Level REST Client by calling **getLowLevelClient()**, use the low-level client to connect to an OpenSearch cluster for which the security mode is disabled, and query whether the **test** index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
import org.elasticsearch.client.Request;
import org.elasticsearch.client.Response;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;

import java.io.IOException;
import java.util.Arrays;
import java.util.List;
public class RestLowLevelClientExample {

 public static void main(String[] args) throws IOException {
 List<String> host = Arrays.asList("{Cluster access address}");
 RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, 9200, "http"));
 final RestHighLevelClient restHighLevelClient = new RestHighLevelClient(builder);
 /**
 * Create a high-level client and then call getLowLevelClient() to obtain a low-level client. The code
 * differs from the client creation code only in the following line:
 */
 final RestClient lowLevelClient = restHighLevelClient.getLowLevelClient();
 /**
 * Check whether the test index exists. If the index exists, 200 is returned. If the index does not exist,
 * 404 is returned.
 */
 Request request = new Request("HEAD", "/test");
```

```
Response response = lowLevelClient.performRequest(request);
System.out.println(response.getStatusLine().getStatusCode());
lowLevelClient.close();
}

/**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
}
}
```

This piece of code checks whether the **test** index exists in the cluster. If 200 or 404 is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode Cluster Using the Elasticsearch High Level REST Client (Without a Certificate)

Use the Elasticsearch 7.10.2 High Level REST Client to obtain the Low Level REST Client by calling **getLowLevelClient()**, use the low-level client to connect to a security-mode OpenSearch cluster that uses HTTP or HTTPS without loading a security certificate, and query whether the **test** index exists. The sample code is as follows:

```
import org.apache.http.HttpHost;
import org.apache.http.HttpResponse;
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.impl.client.DefaultConnectionKeepAliveStrategy;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.apache.http.protocol.HttpContext;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.elasticsearch.client.Request;
import org.elasticsearch.client.Response;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;
import org.elasticsearch.common.Nullable;

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 java.util.Arrays;
import java.util.List;
import java.util.Objects;
import java.util.concurrent.TimeUnit;

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 RestLowLevelClientExample {

 /**
 * Create a class for the client. Define the create function.
 */
 public static RestHighLevelClient create(List<String> host, int port, String protocol, int connectTimeout,
```

```
int connectionRequestTimeout, int socketTimeout, String username, String password) throws IOException {

 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(constructHttpHosts(host, port, protocol))
 .setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout)
 .setConnectionRequestTimeout(connectionRequestTimeout)
 .setSocketTimeout(socketTimeout))
 .setHttpClientConfigCallback(httpClientConfigCallback);
 final RestHighLevelClient client = new RestHighLevelClient(builder);
 logger.info("es rest client build success {}", client);
 return client;
}

/**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
}

/**
 * Configure trustAllCerts to ignore the certificate configuration.
 */
public 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;
 }
 }
};

/**
 * The CustomConnectionKeepAliveStrategy function is used to set the connection keepalive during when
there are a large number of short connections or when there are not many data requests.
 */
public static class CustomConnectionKeepAliveStrategy extends DefaultConnectionKeepAliveStrategy {
 public static final CustomConnectionKeepAliveStrategy INSTANCE = new
CustomConnectionKeepAliveStrategy();

 private CustomConnectionKeepAliveStrategy() {
 super();
 }
}

/**
```



```
* Maximum keepalive time (in minutes)
* The default value is 10 minutes. You can set it based on the number of TCP connections in
TIME_WAIT state. If there are too many TCP connections, you can increase this value.
*/
private final long MAX_KEEP_ALIVE_MINUTES = 10;

@Override
public long getKeepAliveDuration(HttpResponse response, HttpContext context) {
 long keepAliveDuration = super.getKeepAliveDuration(response, context);
 // <0 indicates an unlimited keepalive period.
 // Change the period from unlimited to a default period.
 if (keepAliveDuration < 0) {
 return TimeUnit.MINUTES.toMillis(MAX_KEEP_ALIVE_MINUTES);
 }
 return keepAliveDuration;
}
}

private static final Logger logger = LogManager.getLogger(RestLowLevelClientExample.class);

static 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;
}

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

/**
 * The following is an example of the main function. Call the create function to create the high-level client,
 * then call the getLowLevelClient() function to obtain the low-level client, and query whether the test index
 * exists.
 */
public static void main(String[] args) throws IOException {
 RestHighLevelClient client = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200, "http",
30000, 30000, 30000, "username", "password");
 RestClient lowLevelClient = client.getLowLevelClient();
 Request request = new Request("HEAD", "test");
 Response response = lowLevelClient.performRequest(request);
 System.out.println(response.getStatusLine().getStatusCode());
 lowLevelClient.close();
}
```

**Table 3-51** Variables

| Parameter                | Description                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------|
| host                     | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port                     | Access port of the cluster. The default value is <b>9200</b> .                                           |
| protocol                 | Connection protocol, which can be <b>http</b> or <b>https</b> .                                          |
| connectTimeout           | Socket connection timeout (in ms).                                                                       |
| connectionRequestTimeout | Socket connection request timeout (in ms).                                                               |
| socketTimeout            | Socket request timeout (in ms).                                                                          |
| username                 | Username for accessing the cluster.                                                                      |
| password                 | Password of the user.                                                                                    |

This piece of code checks whether the **test** index exists in the cluster. If 200 or 404 is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode Cluster Using the Elasticsearch High Level REST Client (With a Certificate)

Use the Elasticsearch 7.10.2 High Level REST Client to obtain the Low Level REST Client by calling **getLowLevelClient()**, use the low-level client to connect to a security-mode OpenSearch cluster that uses HTTPS with a security certificate loaded, and query whether the **test** index exists. The sample code is as follows:

---

### CAUTION

For how to obtain and upload a security certificate, see [Obtaining and Uploading a Security Certificate](#).

---

```
import org.apache.http.HttpHost;
import org.apache.http.auth.AuthScope;
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.client.CredentialsProvider;
import org.apache.http.conn.ssl.NoopHostnameVerifier;
import org.apache.http.impl.client.BasicCredentialsProvider;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.apache.logging.log4j.LogManager;
import org.apache.logging.log4j.Logger;
import org.elasticsearch.action.admin.cluster.health.ClusterHealthRequest;
import org.elasticsearch.action.admin.cluster.health.ClusterHealthResponse;
import org.elasticsearch.client.Request;
import org.elasticsearch.client.RequestOptions;
import org.elasticsearch.client.Response;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;
import org.elasticsearch.common.Nullable;

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 java.util.Arrays;
import java.util.List;
import java.util.Objects;

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

public class RestLowLevelClientExample {
 private static final Logger logger = LogManager.getLogger(RestLowLevelClientExample.class);

 /**
 * Create a class for the client. Define the create function.
 */
 public static RestHighLevelClient create(List<String> host, int port, String protocol, int connectTimeout,
 int connectionRequestTimeout, int socketTimeout, String username, String password, String certFilePath,
 String certPassword) throws IOException {
 final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
 credentialsProvider.setCredentials(AuthScope.ANY, new UsernamePasswordCredentials(username,
 password));
 SSLContext sc = null;
 try {
 TrustManager[] tm = {new MyX509TrustManager(certFilePath, certPassword)};
```

```
 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 NoopHostnameVerifier());
 SecuredHttpClientConfigCallback httpClientConfigCallback = new
 SecuredHttpClientConfigCallback(sessionStrategy,
 credentialsProvider);

 RestClientBuilder builder = RestClient.builder(constructHttpHosts(host, port, protocol))
 .setRequestConfigCallback(requestConfig -> requestConfig.setConnectTimeout(connectTimeout)
 .setConnectionRequestTimeout(connectionRequestTimeout)
 .setSocketTimeout(socketTimeout))
 .setHttpClientConfigCallback(httpClientConfigCallback);
 final RestHighLevelClient client = new RestHighLevelClient(builder);
 logger.info("es rest client build success {}", client);

 ClusterHealthRequest request = new ClusterHealthRequest();
 ClusterHealthResponse response = client.cluster().health(request, RequestOptions.DEFAULT);
 logger.info("es rest client health response {}", response);
 return client;
}

/**
 * Use the constructHttpHosts function to convert the node IP address list of the host cluster.
 */
public static HttpHost[] constructHttpHosts(List<String> host, int port, String protocol) {
 return host.stream().map(p -> new HttpHost(p, port, protocol)).toArray(HttpHost[]::new);
}

static class SecuredHttpClientConfigCallback implements RestClientBuilder.HttpClientConfigCallback {
 @Nullable
 private final CredentialsProvider credentialsProvider;

 private final SSLIOSessionStrategy sslStrategy;

 SecuredHttpClientConfigCallback(final SSLIOSessionStrategy sslStrategy,
 @Nullable final CredentialsProvider credentialsProvider) {
 this.sslStrategy = Objects.requireNonNull(sslStrategy);
 this.credentialsProvider = credentialsProvider;
 }

 @Nullable
 CredentialsProvider getCredentialsProvider() {
 return credentialsProvider;
 }

 SSLIOSessionStrategy getSSLStrategy() {
 return sslStrategy;
 }

 @Override
 public HttpAsyncClientBuilder customizeHttpClient(final HttpAsyncClientBuilder httpClientBuilder) {
 httpClientBuilder.setSSLStrategy(sslStrategy);
 if (credentialsProvider != null) {
 httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
 }
 return httpClientBuilder;
 }
}

public static class MyX509TrustManager implements X509TrustManager {
 X509TrustManager sunJSSEX509TrustManager;

 MyX509TrustManager(String certFilePath, String certPassword) throws Exception {
 File file = new File(certFilePath);
 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, certPassword.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];
}
}

/**
 * The following is an example of the main function. Call the create function to create the high-level client,
 * then call the getLowLevelClient() function to obtain the low-level client, and query whether the test index
 * exists.
 */
public static void main(String[] args) throws IOException {
 RestHighLevelClient client = create(Arrays.asList("xxx.xxx.xxx.xxx", "xxx.xxx.xxx.xxx"), 9200, "https",
30000, 30000, 30000, "username", "password", "certFilePath", "certPassword");
 RestClient lowLevelClient = client.getLowLevelClient();
 Request request = new Request("HEAD", "test");
 Response response = lowLevelClient.performRequest(request);
 System.out.println(response.getStatusLine().getStatusCode());
 lowLevelClient.close();
}
}
```

**Table 3-52** Function parameters

| Parameter      | Description                                                                                              |
|----------------|----------------------------------------------------------------------------------------------------------|
| host           | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port           | Access port of the cluster. The default value is <b>9200</b> .                                           |
| protocol       | Connection protocol. Set this parameter to <b>https</b> .                                                |
| connectTimeout | Socket connection timeout (in ms).                                                                       |

| Parameter                | Description                                |
|--------------------------|--------------------------------------------|
| connectionRequestTimeout | Socket connection request timeout (in ms). |
| socketTimeout            | Socket request timeout (in ms).            |
| username                 | Username for accessing the cluster.        |
| password                 | Password of the user.                      |
| certFilePath             | Path for storing the security certificate. |
| certPassword             | Password of the security certificate.      |

This piece of code checks whether the **test** index exists in the cluster. If 200 or 404 is returned, it indicates that the cluster is connected.

## Obtaining and Uploading a Security Certificate

To access a security-mode OpenSearch cluster that uses HTTPS, a security certificate must be loaded. Perform the following steps to obtain the security certificate and upload it to the client:

1. Obtain the security certificate **CloudSearchService.cer**.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > OpenSearch**.
  - c. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  - d. Click the **Overview** tab. In the **Configuration** area, click **Download Certificate** next to **HTTPS Access**.

**Figure 3-6** Downloading a security certificate



2. Convert the security certificate **CloudSearchService.cer**. Upload the downloaded security certificate to the client and use keytool to convert the **.cer** certificate into a **.jks** certificate that can be read by Java.
  - In Linux, run the following command to convert the certificate:  
`keytool -import -alias newname -keystore ./truststore.jks -file ./CloudSearchService.cer`
  - In Windows, run the following command to convert the certificate:  
`keytool -import -alias newname -keystore .\truststore.jks -file .\CloudSearchService.cer`

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

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

## 3.5.6 Accessing an OpenSearch Cluster with the Python Client

Connect to a CSS OpenSearch cluster with Python to query and manage data.

OpenSearch Python Client (or Elasticsearch Python Client) is the official Python client library provided for interacting with OpenSearch clusters using Python.

### Prerequisites

- The target OpenSearch cluster is available.
- The server that runs the Python code can communicate with the OpenSearch cluster.
- Depending on the network configuration method used, obtain the cluster access address. For details, see [Network Configuration](#).
- Python 3 has been installed on the server.

### Installing the Python Client Library

Install the Python client library on the server that runs the Python code.

CSS allows you to use the Elasticsearch 7.10 Python client to connect to OpenSearch clusters. To ensure better compatibility, however, you are advised to use a Python client library that has the same version as the target OpenSearch cluster.

Select a reference example based on the Python client you use.

- Scenario 1 (recommended): Use the OpenSearch Python client to access an OpenSearch cluster and install the OpenSearch Python client library.  
`pip install opensearch-py`
- Scenario 2: Use the Elasticsearch 7.10 Python client to access an OpenSearch cluster and install the Elasticsearch 7.10 Python client library.  
`pip install Elasticsearch==7.10`

### Cluster Access Scenarios

The sample code varies depending on the Python client and the security mode settings of the target OpenSearch cluster. Select the right reference document based on your service scenario.

**Table 3-53** Cluster access scenarios

| Python Client Type | OpenSearch Cluster Security-Mode Settings | Reference Document                                                                            |
|--------------------|-------------------------------------------|-----------------------------------------------------------------------------------------------|
| OpenSearch         | Non-security mode                         | <a href="#">Connecting to a Non-Security Mode Cluster with the OpenSearch Python Client</a>   |
|                    | Security mode + HTTP                      | <a href="#">Connecting to a Security-Mode +HTTP Cluster with the OpenSearch Python Client</a> |

| Python Client Type   | OpenSearch Cluster Security-Mode Settings | Reference Document                                                                                |
|----------------------|-------------------------------------------|---------------------------------------------------------------------------------------------------|
|                      | Security mode + HTTPS                     | <a href="#">Connecting to a Security-Mode +HTTPS Cluster with the OpenSearch Python Client</a>    |
| Elasticsearch 7.10.2 | Non-security mode                         | <a href="#">Connecting to a Non-Security Mode Cluster with the Elasticsearch Python Client</a>    |
|                      | Security mode + HTTP                      | <a href="#">Connecting to a Security-Mode +HTTP Cluster with the Elasticsearch Python Client</a>  |
|                      | Security mode + HTTPS                     | <a href="#">Connecting to a Security-Mode +HTTPS Cluster with the Elasticsearch Python Client</a> |

## Connecting to a Non-Security Mode Cluster with the OpenSearch Python Client

Use the OpenSearch Python client to connect to an OpenSearch cluster for which the security mode is disabled, and query whether the **test** index exists.

The following is an example of the code for creating a client instance using the cluster connection information:

```
from opensearchpy import OpenSearch

class OpenSearchFactory(object):

 def __init__(self, host: list, port: str, username: str, password: str):
 self.port = port
 self.host = host
 self.username = username
 self.password = password

 def create(self) -> OpenSearch:
 addrs = []
 for host in self.host:
 addr = {'host': host, 'port': self.port}
 addrs.append(addr)

 if self.username and self.password:
 opensearch = OpenSearch(addrs, http_auth=(self.username, self.password))
 else:
 opensearch = OpenSearch(addrs)
 return opensearch

os = OpenSearchFactory(["{cluster address}"], "9200", None, None).create()
print(os.indices.exists(index='test'))
```

This piece of code checks whether the **test** index exists in the cluster. If **true** or **false** is returned, it indicates that the cluster is connected.



## Connecting to a Security-Mode+HTTP Cluster with the OpenSearch Python Client

Use the OpenSearch Python client to connect to a security-mode OpenSearch cluster that uses HTTP, and query whether the **test** index exists.

The following is an example of the code for creating a client instance using the cluster connection information:

```
from opensearchpy import OpenSearch

class OpenSearchFactory(object):

 def __init__(self, host: list, port: str, username: str, password: str):
 self.port = port
 self.host = host
 self.username = username
 self.password = password

 def create(self) -> OpenSearch:
 addrs = []
 for host in self.host:
 addr = {'host': host, 'port': self.port}
 addrs.append(addr)

 if self.username and self.password:
 opensearch = OpenSearch(addrs, http_auth=(self.username, self.password))
 else:
 opensearch = OpenSearch(addrs)
 return opensearch

os = OpenSearchFactory(["xxx.xxx.xxx.xxx"], "9200", "username", "password").create()
print(os.indices.exists(index='test'))
```

**Table 3-54** Variables

| Parameter | Description                                                                                              |
|-----------|----------------------------------------------------------------------------------------------------------|
| host      | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port      | Access port of the cluster. Enter <b>9200</b> .                                                          |
| username  | Username for accessing the cluster.                                                                      |
| password  | Password of the user.                                                                                    |

This piece of code checks whether the **test** index exists in the cluster. If **true** or **false** is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode+HTTPS Cluster with the OpenSearch Python Client

Use the OpenSearch Python client to connect to a security-mode OpenSearch cluster that uses HTTPS, and query whether the **test** index exists.

The following is an example of the code for creating a client instance using the cluster connection information:

```
from opensearchpy import OpenSearch
import ssl

class OpenSearchFactory(object):

 def __init__(self, host: list, port: str, username: str, password: str):
 self.port = port
 self.host = host
 self.username = username
 self.password = password

 def create(self) -> OpenSearch:
 context = ssl._create_unverified_context()

 addrs = []
 for host in self.host:
 addr = {'host': host, 'port': self.port}
 addrs.append(addr)

 if self.username and self.password:
 opensearch = OpenSearch(addrs, http_auth=(self.username, self.password), scheme="https",
ssl_context=context)
 else:
 opensearch = OpenSearch(addrs)
 return opensearch
os = OpenSearchFactory(["xxx.xxx.xxx.xxx"], "9200", "username", "password").create()
print(os.indices.exists(index='test'))
```

**Table 3-55** Variables

| Parameter | Description                                                                                              |
|-----------|----------------------------------------------------------------------------------------------------------|
| host      | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port      | Access port of the cluster. Enter <b>9200</b> .                                                          |
| username  | Username for accessing the cluster.                                                                      |
| password  | Password of the user.                                                                                    |

This piece of code checks whether the **test** index exists in the cluster. If **true** or **false** is returned, it indicates that the cluster is connected.

## Connecting to a Non-Security Mode Cluster with the Elasticsearch Python Client

Use the Elasticsearch 7.10.2 Python client to connect to an OpenSearch cluster for which the security mode is disabled, and query whether the **test** index exists.

The following is an example of the code for creating a client instance using the cluster connection information:

```
from elasticsearch import Elasticsearch

class ElasticFactory(object):

 def __init__(self, host: list, port: str, username: str, password: str):
 self.port = port
 self.host = host
```

```
self.username = username
self.password = password

def create(self) -> Elasticsearch:
 addrs = []
 for host in self.host:
 addr = {'host': host, 'port': self.port}
 addrs.append(addr)

 if self.username and self.password:
 elasticsearch = Elasticsearch(addrs, http_auth=(self.username, self.password))
 else:
 elasticsearch = Elasticsearch(addrs)
 return elasticsearch

es = ElasticFactory(["{cluster address}"], "9200", None, None).create()
print(es.indices.exists(index='test'))
```

This piece of code checks whether the **test** index exists in the cluster. If **true** or **false** is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode+HTTP Cluster with the Elasticsearch Python Client

Use the Elasticsearch 7.10.2 Python client to connect to a security-mode OpenSearch cluster that uses HTTP, and query whether the **test** index exists.

The following is an example of the code for creating a client instance using the cluster connection information:

```
from elasticsearch import Elasticsearch

class ElasticFactory(object):

 def __init__(self, host: list, port: str, username: str, password: str):
 self.port = port
 self.host = host
 self.username = username
 self.password = password

 def create(self) -> Elasticsearch:
 addrs = []
 for host in self.host:
 addr = {'host': host, 'port': self.port}
 addrs.append(addr)

 if self.username and self.password:
 elasticsearch = Elasticsearch(addrs, http_auth=(self.username, self.password))
 else:
 elasticsearch = Elasticsearch(addrs)
 return elasticsearch

es = ElasticFactory(["xxx.xxx.xxx.xxx"], "9200", "username", "password").create()
print(es.indices.exists(index='test'))
```

**Table 3-56** Variables

| Parameter | Description                                                                                              |
|-----------|----------------------------------------------------------------------------------------------------------|
| host      | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port      | Access port of the cluster. Enter <b>9200</b> .                                                          |

| Parameter | Description                         |
|-----------|-------------------------------------|
| username  | Username for accessing the cluster. |
| password  | Password of the user.               |

This piece of code checks whether the **test** index exists in the cluster. If **true** or **false** is returned, it indicates that the cluster is connected.

## Connecting to a Security-Mode+HTTPS Cluster with the Elasticsearch Python Client

Use the Elasticsearch 7.10.2 Python client to connect to a security-mode OpenSearch cluster that uses HTTPS, and query whether the **test** index exists.

The following is an example of the code for creating a client instance using the cluster connection information:

```
from elasticsearch import Elasticsearch
import ssl

class ElasticFactory(object):

 def __init__(self, host: list, port: str, username: str, password: str):
 self.port = port
 self.host = host
 self.username = username
 self.password = password

 def create(self) -> Elasticsearch:
 context = ssl._create_unverified_context()

 addrs = []
 for host in self.host:
 addr = {'host': host, 'port': self.port}
 addrs.append(addr)

 if self.username and self.password:
 elasticsearch = Elasticsearch(addrs, http_auth=(self.username, self.password), scheme="https",
ssl_context=context)
 else:
 elasticsearch = Elasticsearch(addrs)
 return elasticsearch

es = ElasticFactory(["xxx.xxx.xxx.xxx"], "9200", "username", "password").create()
print(es.indices.exists(index='test'))
```

**Table 3-57** Variables

| Parameter | Description                                                                                              |
|-----------|----------------------------------------------------------------------------------------------------------|
| host      | IP address for accessing the cluster. If there are multiple IP addresses, separate them with commas (,). |
| port      | Access port of the cluster. Enter <b>9200</b> .                                                          |
| username  | Username for accessing the cluster.                                                                      |

| Parameter | Description           |
|-----------|-----------------------|
| password  | Password of the user. |

This piece of code checks whether the **test** index exists in the cluster. If **true** or **false** is returned, it indicates that the cluster is connected.

### 3.5.7 Accessing an OpenSearch Cluster Using LDAP

The Light Directory Access Protocol (LDAP) is a lightweight version of the directory access protocol based on the X.500 standard. An LDAP service provides user authentication and authorization. Using the Security plugin for Open Distro for Elasticsearch, CSS adds Active Directory as an authentication backend for clusters, connecting them seamlessly to the LDAP service. This topic describes the steps needed to connect a CSS cluster to an LDAP service.

It also describes how to enable LDAP authentication for an OpenSearch cluster to allow access by LDAP users of specific roles.

#### Preparations

- A security-mode OpenSearch cluster has been created in CSS and its status is available.
- The LDAP service that is in the same VPC as the OpenSearch cluster and the necessary user data have been prepared. For details, see [the OpenLDAP document: A Quick-Start Guide](#).

#### Accessing a Cluster

**Step 1** Install an LDAP service on an ECS. If the LDAP service and user data have already been prepared, skip this step.

1. Create an ECS. The ECS must run a Windows OS and must be in the same VPC and security group as the security-mode Elasticsearch cluster of CSS. The Windows Server running on the ECS provides the built-in Active Directory service that supports the LDAP protocol.


For how to create an ECS, see [Purchasing an ECS in Custom Config Mode](#).

2. Log in to the ECS, and enable the Active Directory service. Create a domain, administrator, users, and user groups.

**Step 2** Modify the parameter settings of the security-mode OpenSearch cluster on CSS. Configure a static parameter in **opensearch.yml** to connect the cluster to the LDAP service.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Cluster Settings > Parameter Settings**, and click **Edit**. Expand **Custom**, and add the following parameter and value.

- **Parameter:**  
plugins.security.unsupported.restapi.allow\_securityconfig\_modification
  - **Value:** true
5. Click **Submit** above. In the displayed **Submit Configuration** dialog box, select the box that says "I understand that the modification will take effect after the cluster is restarted." and click **Yes**.  
If **Status** is **Succeeded** in the parameter change list, the change has been saved. Up to 20 change records can be displayed.
  6. After saving the change, click **Restart** in the upper-right corner of the page to restart the cluster, thus making the change take effect.

**Figure 3-7** Restarting a cluster

- You need to restart the cluster after the change, or **Configuration not updated** will be displayed in the **Task Status** column in the cluster list.
- If you restart the cluster after the change, and **Task Status** displays **Configuration error** in the cluster list, the parameter configuration file has failed to be modified.

**Step 3** Configure a route for an OpenSearch cluster on the CSS console to connect the cluster to the LDAP service.

1. On the cluster information page, click the **Overview** tab.
2. In the **Configuration** area, click **Add Route** next to **Cluster Route**.
3. In the displayed dialog box, configure the route information.

**Table 3-58** Adding a route

| Parameter   | Description                                                                                                         |
|-------------|---------------------------------------------------------------------------------------------------------------------|
| IP Address  | Enter the IP address of the LDAP server. If the LDAP service on the ECS is used, enter the IP address of the ECS.   |
| Subnet Mask | Enter the subnet mask of the LDAP server. If the LDAP service on the ECS is used, enter the subnet mask of the ECS. |

4. Click OK.

**Step 4** Configure LDAP authentication for a security-mode OpenSearch cluster.

1. On the cluster information page, click **Dashboards** in the upper-right corner to log in to OpenSearch Dashboards.
2. On OpenSearch Dashboards, expand the menu in the upper-left corner, and choose **Dev Tools**.
3. Run the following commands to configure LDAP authentication.

 NOTE

- Concepts used in an X.500 directory access protocol (including LDAP):

- CN = Common Name
- OU = Organizational Unit
- DC = Domain Component
- DN = Distinguished Name

The CN, OU, and DC must be provided in the correct order. Otherwise, authentication will fail.

- The configuration file consists of two parts: **authc** and **authz**.
  - **authc** (authentication): verifies whether a user is truly who they claim they are (password verification).
  - **authz** (authorization): verifies what the current user has access to.

```
PUT _plugins/_security/api/securityconfig/config
{
 "dynamic": {
 "authc": {
 "basic_internal_auth_domain": {
 "description": "Authenticate via HTTP Basic against internal users database",
 "http_enabled": true,
 "transport_enabled": true,
 "order": 1,
 "http_authenticator": {
 "type": "basic",
 "challenge": true
 },
 "authentication_backend": {
 "type": "intern"
 }
 },
 "ldap": {
 "description": "Authenticate via LDAP or Active Directory",
 "http_enabled": true,
 "transport_enabled": true,
 "order": 2,
 "http_authenticator": {
 "type": "basic",
 "challenge": false
 },
 "authentication_backend": {
 "type": "ldap",
 "config": {
 "enable_ssl": false,
 "enable_start_tls": false,
 "enable_ssl_client_auth": false,
 "verify_hostnames": true,
 "hosts": ["10.0.XXX.XXX:389"],
 "bind_dn": "CN=adminAD,DC=test,DC=ldap,DC=com",
 "password": "<password>",
 "userbase": "OU=ITDepartment,DC=test,DC=ldap,DC=com",
 "usersearch": "(sAMAccountName={0})",
 "username_attribute": "uid"
 }
 }
 }
 },
 "authz": {
 "roles_from_myldap": {
 "description": "Authorize via LDAP or Active Directory",
```

```
"http_enabled": true,
"transport_enabled": true,
"authorization_backend": {
 "type": "ldap",
 "config": {
 "enable_ssl": false,
 "enable_start_tls": false,
 "enable_ssl_client_auth": false,
 "verify_hostnames": true,
 "hosts": ["10.0.XXX.XXX:389"],
 "bind_dn": "CN=adminAD,DC=test,DC=ldap,DC=com",
 "password": "<password>",
 "rolebase": "OU=groups,DC=test,DC=ldap,DC=com",
 "rolesearch": "(member={0})",
 "userroleattribute": null,
 "userrolename": "disabled",
 "rolename": "CN",
 "resolve_nested_roles": true,
 "userbase": "OU=ITDepartment,DC=test,DC=ldap,DC=com",
 "usersearch": "(uid={0})"
 }
}
```

The parameters in [Table 3-59](#) need to be modified based on the actual environment.


**Table 3-59** Parameter description

| Parameter | Description                                                                                                                                                                   |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| hosts     | Address of the LDAP service. The port number is 389. If the LDAP service on the ECS is used, enter the IP address of the ECS.                                                 |
| bind_dn   | It is similar to the LDAP user name (CN - OU - DC) and is used to access the LDAP server. Select a user name from the user data of the LDAP service.                          |
| password  | Password of the LDAP user configured using <b>bind_dn</b> .                                                                                                                   |
| userbase  | After the LDAP service is connected, the DN that the user belongs to is obtained. In this example, all user information in the <b>ITDepartment</b> directory is synchronized. |
| rolebase  | The collection of permissions that can be configured for the <b>userbase</b> user group of the LDAP service.                                                                  |

**Step 5** Configure the mapping between LDAP user permissions and OpenSearch permissions in the OpenSearch security-mode cluster to enable fine-grained access control.

The rolebase permissions group of the LDAP server must be mapped to the roles in the OpenSearch cluster. [Figure 3-8](#) illustrates the mapping. For details about the configuration, see [Creating Users for an OpenSearch Cluster and Granting Cluster Access](#).



**Figure 3-8** Permissions mapping

1. On OpenSearch Dashboards, expand the menu in the upper-left corner, and choose **Security**. The **Security** page is displayed.
2. Click **Roles** to go to the Open Distro Security Roles page. Click **Create Role**, set **Name**, **Cluster Permissions**, **Index Permissions**, and **Tenant Permissions**. Then click **Save Role Definition** to save the role settings. The parameters are as follows:
  - Name (name of the role)
  - Cluster Permissions
  - Index permissions
  - Tenant permissions
3. Click the newly created role, select **Mapped users**, enter a permissions group of the LDAP service in **Backend roles**, and click **Map**.
4. Check the configuration result.

**Figure 3-9** Permissions mapping

|                                      |                          |                   |                |                 |                               |               |
|--------------------------------------|--------------------------|-------------------|----------------|-----------------|-------------------------------|---------------|
| Q reserved:false                     | Cluster permissions      | Index permissions | Internal users | Backend roles   | Tenants                       | Customization |
| <input type="checkbox"/> Role        | Cluster permissions      | Index permissions | Internal users | Backend roles   | Tenants                       | Customization |
| <input type="checkbox"/> role_google | cluster_composite_ops_ro | index*            | —              | groupGoogleInAD | global_tenant<br>admin_tenant | Custom        |

5. Repeat [Step 5.2](#) to [Step 5.4](#) to map other permissions groups.

**Step 6** Log in to OpenSearch Dashboards using the LDAP user to verify the configuration.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column. Use the LDAP user to log in to OpenSearch Dashboards.  
If the login is successful, the configuration is successful, and users can access the OpenSearch cluster through LDAP. The specific permissions authorized are controlled by role permissions configured in OpenSearch.

-----End

## 3.6 Importing Data to an OpenSearch Cluster

## 3.6.1 Different Ways to Ingest Data into an OpenSearch Cluster

### Introduction

OpenSearch clusters support multiple data ingestion methods, as listed in [Table 3-60](#). Select one that fits your needs the best. Before starting to ingest data, determine whether to enhance the data ingestion performance of OpenSearch clusters first. For details, see [Enhancing the Data Ingestion Performance of OpenSearch Clusters](#).

**Table 3-60** Different ways to ingest data into an OpenSearch cluster

| Data Ingestion Method      | Scenario                                                                                                                                                                                                                        | Supported Data Formats | Details                                                                                   |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|-------------------------------------------------------------------------------------------|
| Open-source Logstash       | Open-source Logstash offers a server-side, real-time data processing pipeline, which supports data ingestion from multiple sources. It can be used to ingest various types of data, such as logs, monitoring data, and metrics. | JSON, CSV, and text    | <a href="#">Using In-house Built Logstash to Ingest Data into an OpenSearch Cluster</a>   |
| Open-source OpenSearch API | Open-source Elasticsearch APIs can be used to ingest data. This method is flexible, as you can write your own application code.                                                                                                 | JSON                   | <a href="#">Using Open Source OpenSearch APIs to Import Data to an OpenSearch Cluster</a> |
| Cloud Data Migration (CDM) | You can use CDM for batch data migration. For example, if data is stored in OBS or an Oracle database, CDM is recommended.                                                                                                      | JSON                   | <a href="#">Using CDM to Import Data to an OpenSearch Cluster</a>                         |

| Data Ingestion Method          | Scenario                                                                          | Supported Data Formats            | Details                                                                |
|--------------------------------|-----------------------------------------------------------------------------------|-----------------------------------|------------------------------------------------------------------------|
| Data Replication Service (DRS) | DRS can be used for online database migration and real-time data synchronization. | Relational Database Service (RDS) | <a href="#">Using DRS to Ingest Data from a Database to OpenSearch</a> |

## Using DRS to Ingest Data from a Database to OpenSearch

DRS is an easy-to-use, stable, and efficient cloud service for online database migration and real-time database synchronization. Real-time data synchronization refers to the real-time replication of data from one source to another while ensuring data consistency.

DRS can be used to ingest data from multiple types of relational databases to OpenSearch clusters. For details about the supported software versions for source databases and destination clusters, see [Table 3-61](#).

**Table 3-61** Using DRS to ingest data from a database to an OpenSearch cluster

| Scenario                                                                           | Source DB                                                             | Destination OpenSearch Cluster | Details                                                                        |
|------------------------------------------------------------------------------------|-----------------------------------------------------------------------|--------------------------------|--------------------------------------------------------------------------------|
| Ingesting data from an RDS for MySQL database into a CSS OpenSearch cluster        | RDS for MySQL 5.5, 5.6, 5.7, or 8.0                                   | OpenSearch 1.3.6               | <a href="#">(Real-Time Synchronization) From MySQL to CSS/ES</a>               |
| Ingesting data from a TaurusDB database into a CSS OpenSearch cluster              | Primary/standby TaurusDB instances                                    | OpenSearch 1.3.6               | <a href="#">(Real-Time Synchronization) From TaurusDB to CSS/Elasticsearch</a> |
| Ingesting data from an in-house built MySQL database into a CSS OpenSearch cluster | MySQL database 5.5, 5.6, 5.7, or 8.0 created on a local server or ECS | OpenSearch 1.3.6               | <a href="#">(Real-Time Synchronization) From MySQL to CSS/ES</a>               |

## 3.6.2 Using In-house Built Logstash to Ingest Data into an OpenSearch Cluster

With CSS, you can use in-house developed Logstash to ingest data into OpenSearch for efficient search and exploration. Supported data formats include JSON and CSV.

Logstash is an open-source, server-side data processing pipeline that ingests data from multiple sources simultaneously, processes and transforms the data, and then sends it to OpenSearch. For more information about Logstash, visit <https://www.elastic.co/guide/en/logstash/current/getting-started-with-logstash.html>

Depending on where Logstash is deployed, there are two data ingestion scenarios:

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


### Prerequisites

- To facilitate operations, you are advised to deploy Logstash on a host that runs a Linux operating system (OS).
- Logstash must use an OSS version that is consistent with that of the CSS cluster. To download Logstash, go to <https://www.elastic.co/downloads/logstash-oss>.
- The JDK must be installed before Logstash is installed. In a Linux OS, you can run the `yum -y install java-1.8.0` command to install JDK 1.8.0. In a Windows OS, you can download the required JDK version from the [official website of JDK](#), and install it by following the installation guide.
- After installing Logstash, perform the steps below to ingest data. For details about how to install Logstash, visit the following website: <https://www.elastic.co/guide/en/logstash/current/installing-logstash.html>
- In the [Ingesting Data When Logstash Is Deployed Using an ECS](#) scenario, ensure that the ECS and the destination Elasticsearch cluster are in the same VPC.

### Ingesting Data When Logstash Is Deployed on an External Network

**Figure 3-10** illustrates the data ingestion process when Logstash is deployed on an external network.

**Figure 3-10** Data ingestion process 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 a Linux OS and an EIP has been associated with it.
  - The jump host resides in the same VPC as the destination 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 jump host via its EIP.
3. Run the following command to configure port mapping to forward requests sent to the opened port on the jump host to the destination 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>
```

  - In the preceding command, *<Local port of the jump host>* refers to the jump host port configured 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, 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 accessible from the public network, 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
```

4. Log in to the server where Logstash is deployed and store the files to be ingested on this server.

For example, data file **access\_20181029\_log** needs to be ingested, the file storage path is **/tmp/access\_log/** (create the **access\_log** folder if it does not already exist), and the data file contains 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     |

5. On 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
```

6. Enter the following content in **logstash-simple.conf**:

```
input {
 Location of data
}
filter {
 Related data processing
}
output {
 elasticsearch {
 hosts => "<EIP of the jump host>:<Port number that is accessible from the public network 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 its usage, visit the following website: <https://www.elastic.co/guide/en/logstash/current/input-plugins.html>
- The **filter** parameter specifies the data processing method. For example, extract and process logs to convert unstructured information into structured information. For details about the **filter** parameter and its 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 its usage, visit <https://www.elastic.co/guide/en/logstash/current/output-plugins.html>. Replace <EIP address of the jump host> with the IP address (with **EIP**) of the created jump host in the **IP Address** column in the ECS list on the ECS management console. <The port accessible from the public network on the jump host> is the port number 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 ingestion starts from the first row of the 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 destination index is **myindex**. Edit the configuration file as follows, and enter **:wq** to save the change 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** to try and rectify the error. If the cluster has the security mode enabled, you need to download a certificate first.

a. **Obtaining the Security Certificate.**

- b. Store the downloaded 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 ingestion starts from the first row of the 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. The name of the destination index 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 change 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" # Username for accessing the security-mode cluster
 password => "*****" # Password for accessing the security-mode cluster
 cacert => "/logstash/logstash6.8/config/CloudSearchService.cer"
 manager_template => false
 ilm_enabled => false
 ssl => true
 ssl_certificate_verification => false
 }
}
```

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

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

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


8. Log in to the **CSS management console**.
9. In the navigation pane on the left, choose **Clusters > OpenSearch**.
10. In the cluster list, find the destination cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
11. In the left navigation pane, choose **Dev Tools**.
12. On the **Console** page of OpenSearch Dashboards, search for the ingested data.

Run the following command to search for data. Check the search results. If they are consistent with the ingested data, data ingestion has been successful.

```
GET myindex/_search
```

## Ingesting Data When Logstash Is Deployed Using an ECS

**Figure 3-11** illustrates the data ingestion process when Logstash is deployed on an ECS that resides in the same VPC as the destination cluster.

**Figure 3-11** Data ingestion process when Logstash is deployed on an ECS

1. Make sure the ECS where Logstash is deployed and the destination cluster reside in the same VPC, port **9200** is opened in the ECS's security group to allow external network access, and an EIP has been associated with the ECS.
  - If there are multiple servers in a VPC, you only need to associate an EIP with one of these servers. Switch to the node where Logstash is deployed from the node with which the EIP is associated.
  - If a private line or VPN is available, there is no need for an EIP.

2. Use PuTTY to log in to the ECS.

For example, the file **access\_20181029\_log** is stored in the **/tmp/access\_log/** path of the ECS, and the file contains 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     |

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

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

Enter 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 its



usage, visit the following website: <https://www.elastic.co/guide/en/logstash/current/input-plugins.html>

- The **filter** parameter specifies the data processing method. For example, extract and process logs to convert unstructured information into structured information. For details about the **filter** parameter and its 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 its 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 ingestion starts from the first row of the 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 destination cluster are **192.168.0.81** and **9200**, respectively, and the name of the destination index is **myindex**. Edit the configuration file as follows, and enter **:wq** to save the change and exit.

```
input {
 file{
 path => "/tmp/access_log/*"
 start_position => "beginning"
 }
}
filter {
}
output {
 elasticsearch {
 hosts => "192.168.0.81:9200"
 index => "myindex"
 }
}
```

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

- Obtaining the Security Certificate.**
- Store the downloaded certificate to the server where Logstash is deployed.
- Modify the **logstash-simple.conf** configuration file.

Consider the data files in the **/tmp/access\_log/** path mentioned in [step 2](#) as an example. Assume that data ingestion starts from the first row of the 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, The

name of the destination index 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 change 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" # Username for accessing the security-mode cluster
 password => "*****" # Password for accessing the security-mode cluster
 cacert => "/logstash/logstash6.8/config/CloudSearchService.cer"
 manager_template => false
 ilm_enabled => false
 ssl => true
 ssl_certificate_verification => false
 }
}
```

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 > OpenSearch**.
7. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
8. In the left navigation pane, choose **Dev Tools**.
9. On the **Console** page of OpenSearch Dashboards, search for the ingested data.

Run the following command to search for data. Check the search results. If they are consistent with the ingested data, data ingestion has been successful.

```
GET myindex/_search
```

## Obtaining the Security Certificate

To access a security-mode OpenSearch cluster that uses HTTPS, a security certificate must be loaded. To obtain this security certificate (CloudSearchService.cer), follow these steps:

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Overview** tab. In the **Configuration** area, click **Download Certificate** next to **HTTPS Access**.

**Figure 3-12** Downloading a security certificate

### 3.6.3 Using Open Source OpenSearch APIs to Import Data to an OpenSearch Cluster

With CSS, you can use open-source OpenSearch APIs on OpenSearch Dashboards or an ECS server to import data to an OpenSearch cluster. JSON files are supported. For details about OpenSearch APIs, see [REST API reference](#).

- **Using an Open-Source OpenSearch API to Import Data on OpenSearch Dashboards:** Run POST commands to import data.
- **Using an Open-Source OpenSearch API to Import Data on an ECS Server:** Run cURL commands to import data.

#### Using an Open-Source OpenSearch API to Import Data on OpenSearch Dashboards

On OpenSearch Dashboards, you can run POST commands to import single pieces of data using an open-source OpenSearch API.

1. Log in to the OpenSearch Dashboards.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > OpenSearch**.
  - c. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
  - d. In the left navigation pane, choose **Dev Tools**.
2. Run the following command on OpenSearch Dashboards to check whether the cluster has indexes:

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

  - If indexes are available in the cluster to which you want to import data, you do not need to create an index. Go to [7](#).
  - If no indexes are available in the cluster, go to the next step to create an index.
3. Run the following command on OpenSearch Dashboards to create an index for storing imported data and create a custom mapping to define the data type.

For example, run the following command to create index **my\_store**:

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

```
 "size": {
 "type": "keyword"
 }
 }
}
```

4. Run commands on OpenSearch Dashboards to import data. In the example below, only a single piece of data is imported:

```
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 3-13](#). If the value of the **errors** field in the result is **false**, the data is successfully imported.

**Figure 3-13** 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 }
```

## Using an Open-Source OpenSearch API to Import Data on an ECS Server

On an ECS server, you can run cURL commands to use an open-source OpenSearch API to import JSON files.

In the example below, a cluster in non-security mode is used to describe how to import data using cURL commands. For the commands to use with a security cluster, see [Accessing an OpenSearch Cluster Using cURL Commands](#).

1. Obtain the private network address of the cluster.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > OpenSearch**.
  - c. In the cluster list, obtain the target cluster's private IP address from the **Private IP Address** column. Generally, the IP address format is `<host>.<port>` or `<host>.<port>,<host>.<port>`.

If the cluster has only one node, the IP address and port number of this 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**.

2. Log in to the ECS that you are going to use to import data.

For how to buy and use an ECS, see [Purchasing and Logging In to a Linux ECS](#).

3. Upload a JSON file to the ECS.

For example, save the following data as a JSON file and upload the file to the ECS:

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

4. Run the following command in the path where the JSON file is stored in the ECS to import the JSON file to an OpenSearch cluster.

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

Replace **{Private network address and port number of the node}** with the private network address and port number of a node in the cluster. 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. **test.json** indicates the JSON file to be imported.

#### NOTE

In the case of the failure of a cluster node, if the cluster contains multiple nodes, you can replace **{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 this node and execute the command again.

**Example:** In this example, assume that you need to import data in the **test.json** file to an OpenSearch 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.

- a. Run the following command to create an index named **my\_store**:

```
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 **test.json** file:

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

In this case, if the following information is displayed, the data is successfully imported:

```
{
 "took": 204,
 "errors": false,
 "items": [
 {
 "index": {
 "_index": "my_store",
 "_type": "_doc",
 "_id": "DJQkBlwBbJvUd2769Wi-",
 "_version": 1,
 "result": "created",
 "_shards": {
 "total": 2,
 "successful": 1,
 "failed": 0,
 "_seq_no": 0,
 "_primary_term": 1,
 "status": 201
 }
 }
 },
 {
 "index": {
 "_index": "my_store",
 "_type": "_doc",
 "_id": "DZQkBlwBbJvUd2769Wi_",
 "_version": 1,
 "result": "create",
 "_shards": {
 "total": 2,
 "successful": 1,
 "failed": 0,
 "_seq_no": 1,
 "_primary_term": 1,
 "status": 201
 }
 }
 }
]
}
```

### 3.6.4 Using CDM to Import Data to an OpenSearch Cluster

With CSS, you can use the web user interface of the Cloud Data Migration (CDM) service to import data from an Oracle database or OBS to an OpenSearch cluster. Only JSON files are supported.

**Table 3-62** Using CDM to import data to CSS

| Scenario                                      | Source Data                            | Destination Cluster |
|-----------------------------------------------|----------------------------------------|---------------------|
| Importing data from an Oracle database to CSS | A local or third-party Oracle database | OpenSearch 1.3.6    |
| Importing data from OBS to CSS                | JSON files in OBS buckets              | OpenSearch 1.3.6    |

## Preparations

1. Confirm the source data.

For example, the source data can be the following JSON files:

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

2. Obtain data source information.
  - If the data source is an Oracle database, you need to obtain the IP address, database name, username, and password of the database.
  - If the data source is JSON files in an OBS bucket, you need to obtain its domain name, port, AK, and SK.
3. If the data source is an Oracle database, make sure the Oracle database can be accessed using a public IP address, or a VPN or Direct Connect connection has been established between the customer's data center and Huawei Cloud.

## Importing Data

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, obtain the target cluster's private IP address from the **Private IP Address** column. Generally, the IP address format is `<host>:<port>` or `<host>:<port>,<host>:<port>`.

If the cluster has only one node, the IP address and port number of this 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**.

4. In the cluster list, locate the destination cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
5. In the left navigation pane, choose **Dev Tools**.
6. Run the following command on OpenSearch Dashboards to check whether the cluster has indexes:

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

- If indexes are available in the cluster to which you want to import data, you do not need to create an index. Go to [8](#).
- If no indexes are available in the cluster, go to the next step to create an index.

7. Run the following command to create an index for storing imported data and create a custom mapping to define the data type.

For example, run the following command to create index **demo**:

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

The command is successfully executed if the following information is displayed.

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

8. Use CDM to import data from Oracle or OBS to the OpenSearch cluster.
  - If the data source is an Oracle database, see [Migrating Data from Oracle to CSS](#).
  - If the source data is JSON files in an OBS bucket, see [Migrating Data from OBS to CSS](#).
9. After the data migration is complete, go to the OpenSearch Dashboards console of the OpenSearch cluster again, and search for the imported data.

Run the following command to search for data. If the imported data is consistent with the source data, data importing is successful.

```
GET demo/_search
```

**demo** is the name of the created index. Replace it if another index is used.

The command is successfully executed if the following information is displayed.

```
{
 "took": 18,
```

```
"timed_out": false,
"_shards": {
 "total": 1,
 "successful": 1,
 "skipped": 0,
 "failed": 0
},
"hits": {
 "total": 2,
 "max_score": 1,
 "hits": [
 {
 "_index": "demo",
 "_type": "products",
 "_id": "g6UepnEBuvdFwWkRmn4V",
 "_score": 1,
 "_source": {
 "size": """"size": "L""""",
 "productName": """"{"productName": "Latest art shirts for women in autumn 2019""""
 }
 },
 {
 "_index": "demo",
 "_type": "products",
 "_id": "hKUepnEBuvdFwWkRmn4V",
 "_score": 1,
 "_source": {
 "size": """"size": "M""""",
 "productName": """"{"productName": "Latest art shirts for women in autumn 2019""""
 }
 }
]
}
```

### 3.6.5 Enhancing the Data Ingestion Performance of OpenSearch Clusters

#### Overview

This unique CSS feature significantly enhances data ingestion performance and reduces write rejections through bulk route optimization, text index acceleration, and word segmentation acceleration. You are advised to use this feature for clusters that contain a large number of index shards and text indexes, or have a high inbound data throughput.



**Table 3-63** Ways to improve data ingestion performance

| Method                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Details                                       |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|
| Bulk route optimization       | <p>According to the default routing rule of OpenSearch, each piece of data in a bulk request is routed to a different shard. When a large amount of data is written in and a large number of index shards exist, excessive internal request forwarding may trigger bulk rejection. Additionally, in a large-scale cluster, the long tail effect causes a high latency for bulk requests.</p> <p>You can specify the <b>index.bulk_routing</b> configuration item to enable bulk route optimization. You can use it to reduce the number of requests that need to be internally forwarded. For a cluster that has a large number of shards, this setting can improve write performance and reduce bulk rejections.</p>                                                                                                                                                                                             | <a href="#">Bulk Route Optimization</a>       |
| Bulk aggregation optimization | <p>You can specify the <b>index.aggr_perf_batch_size</b> configuration item to enable or disable bulk aggregation optimization. When bulk aggregation optimization is enabled, documents in a bulk request are written in batches, rather than individually. This helps to reduce the overheads of memory requests, lock requests, and other calls, enhancing data ingestion performance.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <a href="#">Bulk Aggregation Optimization</a> |
| Text index acceleration       | <p>You can configure <b>index.native_speed_up</b> to enable or disable text index acceleration. This setting optimizes the indexing process and memory access to accelerate index building for text fields (text and keyword). When text index acceleration is enabled, you can configure <b>index.native_analyzer</b> to enable word segmentation acceleration as well. For texts that need common word segmentation, you can use the analyzer to accelerate word segmentation.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>Word segmentation acceleration can be enabled (set <b>index.native_analyzer</b> to <b>true</b>) only when text index acceleration is enabled (set <b>index.native_speed_up</b> to <b>true</b>). Otherwise, word segmentation acceleration will not take effect.</li><li>Text index acceleration cannot be enabled for indexes that contain nested fields.</li></ul> | <a href="#">Text Index Acceleration</a>       |

| Method                        | Description                                                                                                                                                                                                                                                                                                                                                                           | Details                                       |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|
| Index merge task optimization | Enabling the three optimization options above will increase the number of index merge tasks in the cluster. You can configure <b>index.merge.scheduler.max_thread_count</b> to reduce the impact of index merge tasks on data ingestion performance. The optimization is about increasing the number of shard merging threads and thereby alleviate the throttling of data ingestion. | <a href="#">Index Merge Task Optimization</a> |

## Constraints

Only OpenSearch 1.3.6 and 2.19.0 clusters support data ingestion performance enhancement.

## Prerequisites

The target cluster is available.

## Logging In to OpenSearch Dashboards

Log in to OpenSearch Dashboards and go to the command execution page. OpenSearch clusters support multiple access methods. This topic uses OpenSearch Dashboards as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
4. In the left navigation pane, choose **Dev Tools**.

## Bulk Route Optimization

You can specify the **index.bulk\_routing** configuration item to enable bulk route optimization. You can use it to reduce the number of requests that need to be internally forwarded. For a cluster that has a large number of shards, this setting can improve write performance and reduce bulk rejections.

---

### CAUTION

After bulk route optimization is enabled (that is, **index.bulk\_routing** is set to **pack** or **local\_pack**), data writes are no longer routed based on **\_id**, and routing-related functionality may be affected. For example, ID-based GET requests may fail.

---

Run the following command to enable bulk route optimization:

```
PUT my_index
{
```

```
"settings": {
 "index.bulk_routing": "local_pack"
}
```

The options for the **index.bulk\_routing** configuration item are as follows:

- **default**: Use the default routing mechanism. Records in a bulk request are split and then routed separately.
- **pack**: All data of a single bulk request is randomly routed to the same shard.
- **local\_pack**: The data of a single bulk request is routed to a local shard of the data node that has received the bulk request. If the node does not contain the corresponding index shard, the data is randomly routed to another node that contains the index shard. This solution depends on the random distribution of client bulk requests and the balanced distribution of primary shards.

## Bulk Aggregation Optimization

You can specify the **index.aggr\_perf\_batch\_size** configuration item to enable or disable bulk aggregation optimization. When bulk aggregation optimization is enabled, documents in a bulk request are written in batches, rather than individually. This helps to reduce the overheads of memory requests, lock requests, and other calls, enhancing data ingestion performance.

Run the following command to enable bulk aggregation optimization:

```
PUT my_index
{
 "settings": {
 "index.aggr_perf_batch_size": "128"
 }
}
```

The value range of the **index.aggr\_perf\_batch\_size** configuration item is [1, Integer.MAX\_VALUE]. The default value is 1, indicating that bulk aggregation optimization is disabled. If the value is greater than 1, bulk aggregation optimization is enabled, the value of **MIN(bulk\_doc\_size, aggr\_perf\_batch\_size)** indicates the bulk size.

## Text Index Acceleration

You can configure **index.native\_speed\_up** to enable or disable text index acceleration. This setting optimizes the indexing process and memory access to accelerate index building for text fields (text and keyword). When text index acceleration is enabled, you can configure **index.native\_analyzer** to enable word segmentation acceleration as well. For texts that need common word segmentation, you can use the analyzer to accelerate word segmentation.

---

### CAUTION

- Word segmentation acceleration can be enabled (set **index.native\_analyzer** to **true**) only when text index acceleration is enabled (set **index.native\_speed\_up** to **true**). Otherwise, word segmentation acceleration will not take effect.
  - Text index acceleration cannot be enabled for indexes that contain nested fields.
-

Run the following command to enable text index acceleration:

```
PUT my_index
{
 "settings": {
 "index.native_speed_up": true,
 "index.native_analyzer": true
 }
}
```

The values of **index.native\_speed\_up** and **index.native\_analyzer** are **true** or **false**. The default value is **false**.

## Index Merge Task Optimization

Enabling the three optimization options above will increase the number of index merge tasks in the cluster. You can configure **index.merge.scheduler.max\_thread\_count** to reduce the impact of index merge tasks on data ingestion performance. The optimization is about increasing the number of shard merging threads and thereby alleviate the throttling of data ingestion.

Run the following command to start index merge task optimization:

```
PUT my_index
{
 "settings": {
 "index.merge.scheduler.max_thread_count": 8
 }
}
```

The value range of **index.merge.scheduler.max\_thread\_count** is [1, node.processors/2]. The default value is 4, and the recommended value is 8.

## 3.7 Searching Data in an OpenSearch Cluster

### 3.7.1 Using DSL to Search for Data in OpenSearch

DSL is the specified query language for OpenSearch. It is the best language for interaction between OpenSearch clusters and clients. OpenSearch DSL is a JSON-based language. Other languages, such as SQL, are translated into OpenSearch DSL before they can be used for interacting with OpenSearch clusters.

This topic lists some of the most commonly used OpenSearch DSL query statements. For more, see [Query DSL](#).

#### DSL Usage Example

Compile the request content in JSON format on the Dev Tools page of OpenSearch Dashboards and execute the search request.

For example, run the following command to retrieve all documents in the **test** index:

```
GET /test/_search
{
 "query": {
 "match_all": {}
 }
}
```

The search result is also in JSON format.

## Common DSL Query Statements

- Sets the query filters, which is equivalent to **where** in the SQL language.  
In the command below, there is no index filter in front of **\_search**, so all indexes are queried. A bool query allows you to combine multiple search queries with boolean conditions. **filter** forcibly filters documents whose **status** field is **published** and **publish\_date** is later than **2015-01-01**. **must** specifies that both **title** and **content** must include **Search**.

### NOTE

The difference between **must** and **filter** is that **filter** is equivalent to **where** in SQL but its results are not used for scoring. The **must** field is also a mandatory filter criteria, but the matching documents are scored based on relevance. The most relevant documents are displayed at the top.

```
GET /_search
{
 "query": {
 "bool": {
 "must": [
 {
 "match": {
 "title": "Search"
 }
 },
 {
 "match": {
 "content": "search"
 }
 }
],
 "filter": [
 {
 "term": {
 "status": "published"
 }
 },
 {
 "range": {
 "publish_date": {
 "gte": "2015-01-01"
 }
 }
 }
]
 }
 }
}
```

- Aggregations** are similar to **Group by** in SQL.  
An aggregation summarizes your data as metrics, statistics, or other analytics. In the example below, the results are aggregated based on the title field in the **test** index. If **title** is of the text (including keyword) type, use **title.keyword** for aggregation. By default, a cluster cannot directly aggregate data of the text type. **titles** is only an example name of the aggregation. You can name the aggregation **titleaggs** instead.

```
GET /test/_search
{
 "aggs": {
 "titles": {
 "terms": {
```

```
 "field": "title.keyword"
 }
}
```

The example above for query aggregation includes all documents in the **test** index. That is, **match\_all** is used. You can set search criteria to narrow the scope of the aggregation to specific documents.

### 3.7.2 Using SQL to Search for Data in OpenSearch

For OpenSearch clusters, Open Distro for OpenSearch SQL lets you write queries in SQL rather than the OpenSearch query domain-specific language (DSL).

Users who are familiar with SQL can use SQL to search for data in OpenSearch clusters.

#### SQL Usage Example

- Use SQL to search for data on OpenSearch Dashboards (recommended).

Log in to OpenSearch Dashboards and, on the **Dev Tools** page, send requests using request parameters or the request body to **\_plugins/\_sql**.

For example, run the following command to retrieve 50 records from the **my-index** index:

```
POST _plugins/_sql
{
 "query": "SELECT * FROM my-index LIMIT 50"
}
```

By default, the query result is returned in JSON format. To return data in CSV format, set the **format** parameter in the command as follows:

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

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

- Run cURL commands on an ECS to execute SQL statements to search for data.

For example, run the following command to retrieve 10 records from the **opensearch\_dashboards\_sample\_data\_flights** index.

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

**localhost** indicates the cluster access address. **username** and **password** indicate the username and password for accessing the security-mode cluster, respectively.

#### Supported SQL Operations

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

- Statements

**Table 3-64** 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 3-65** 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 ('Bob', 'David')    |
| Not            | SELECT * FROM my-index WHERE name NOT IN ('Bob')         |
| 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 3-66** Aggregations

| Aggregation  | Example                                      |
|--------------|----------------------------------------------|
| avg()        | SELECT avg(age) FROM my-index                |
| count()<br>) | 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 3-67** Include and exclude fields

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

- Functions

**Table 3-68** 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 3-69** 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 constraints, see [Joins](#).

- Show

Show commands display indexes and mappings that match an index pattern. You can use \* or % as a wildcard.

**Table 3-70** Show

| Show             | Example                 |
|------------------|-------------------------|
| Show tables like | SHOW TABLES LIKE logs-* |

## Joins

Open Distro for OpenSearch 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 each 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 OpenSearch with your business intelligence (BI) applications.

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

## 3.8 Enhancing Search Capabilities for OpenSearch Clusters

### 3.8.1 Search Enhancement Features for OpenSearch Clusters

Compared with open-source OpenSearch, OpenSearch clusters in CSS have many enhanced features. [Table 3-71](#) lists these enhanced features and the corresponding cluster versions.

**Table 3-71** Search enhancement features for OpenSearch clusters

| Enhanced Feature | Description                                                                                                                                                                                          | Cluster Version                     | Details                                 |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|-----------------------------------------|
| Vector search    | Unstructured data, such as images, videos, and language corpora, is converted into vectors, which are searched based on similarity using either an exact or approximate nearest neighbors algorithm. | OpenSearch 1.3.6, OpenSearch 2.19.0 | <a href="#">About CSS Vector Search</a> |

| Enhanced Feature           | Description                                                                                                                                                                                                                                                                                                                                                                           | Cluster Version                     | Details                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|-------------------------------------------------------------------------------------|
| Storage-compute decoupling | CSS stores new data as hot data on SSDs to ensure optimal query performance, and historical data as cold data in OBS to cut storage costs.<br><br>Compared with cold/hot storage switchover, storage-compute decoupling is a better option for use cases that are not particularly demanding in terms of search performance, as cold data is stored in OBS, which cuts storage costs. | OpenSearch 1.3.6, OpenSearch 2.19.0 | <a href="#">Configuring Decoupled Storage and Compute for an OpenSearch Cluster</a> |
| Flow control               | Protects clusters from overload through flow control policies, such as client request throttling, shard indexing backpressure, and traffic pattern analysis, ensuring proper resource allocation and risk prevention.                                                                                                                                                                 | OpenSearch 2.19.0                   | <a href="#">Configuring Flow Control for an OpenSearch Cluster</a>                  |
| Large query isolation      | Helps effectively manage memory-intensive, time-consuming query requests, ensuring cluster stability.                                                                                                                                                                                                                                                                                 | OpenSearch 2.19.0                   | <a href="#">Configuring Large Query Isolation for an OpenSearch Cluster</a>         |
| Enhanced aggregation       | CSS OpenSearch clusters enhance aggregation performance in the face of large data volumes by leveraging vectorization and optimized clustering, enabling faster analytics and decision-making in complex situations.                                                                                                                                                                  | OpenSearch 2.19.0                   | <a href="#">Configuring Enhanced Aggregation for an OpenSearch Cluster</a>          |
| Read/write splitting       | Read/write splitting directs writes to the leader cluster and queries to the follower cluster, separating the loads. If the leader becomes unavailable, the follower automatically upgrades to leader to ensure service continuity. You may use this feature where high availability and high query performance are prioritized.                                                      | OpenSearch 2.19.0                   | <a href="#">Configuring Read/Write Splitting Between Two OpenSearch Clusters</a>    |

| Enhanced Feature                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Cluster Version                                                       | Details                                                                          |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------|----------------------------------------------------------------------------------|
| Switchover between hot and cold storage | <p>You can keep hot data on high-performance servers to ensure fast query response times (in seconds). For historical data that requires a query response time of minutes, you can keep it on large-capacity, low-specs servers as cold data. This allows you to cut storage costs and improve search efficiency.</p> <p>Compared with storage-compute decoupling, cold/hot storage switchover is a better option for use cases that are demanding in terms of search performance. Cold data is stored on local cold data nodes in the cluster. The storage capacity available depends on the number of cold data nodes and their disk capacity. The storage cost is higher than that of OBS.</p> | This feature is supported as long as the cluster has cold data nodes. | <a href="#">Switching Between Hot and Cold Storage for an OpenSearch Cluster</a> |
| Index recycle bin                       | To prevent data loss caused by accidental deletion, CSS provides an index recycle bin. When enabled, deleted indexes are temporarily stored in the recycle bin, allowing for recovery before they are permanently removed. This feature improves data reliability and operational security.                                                                                                                                                                                                                                                                                                                                                                                                       | OpenSearch 2.19.0                                                     | <a href="#">Configuring an Index Recycle Bin for an OpenSearch Cluster</a>       |

## 3.8.2 Configuring Vector Search for OpenSearch Clusters

### 3.8.2.1 About CSS Vector Search

In the age of AI, unstructured data such as images, videos, audios, and text is growing rapidly. Traditional keyword-based search cannot effectively handle unstructured data because it cannot capture deep semantic or visual features. To enable effective search across unstructured data, CSS provides a vector search solution that enables high-performance, high-accuracy nearest neighbor or approximate nearest neighbor search. Typical use cases include image search, video search, similar product recommendations, semantic text search, and cross-modal search (for example, searching for images using text).

#### Advantages

- Efficient and reliable: The built-in CSS vector search engine provides high search throughput and low latency. Based on the OpenSearch native

distributed architecture, CSS delivers enterprise-grade reliability guarantees, including multi-replica, snapshots, and permission control.

- Flexible adaptation: CSS supports multiple indexing algorithms, including brute-force search, graph-based algorithms (such as HNSW), product quantization (PQ), and IVF-HNSW; as well as multiple similarity measurement methods, including Euclidean, inner product, cosine, and Hamming.
- Elasticsearch ecosystem compatibility: CSS vector search is compatible with open-source OpenSearch query language and APIs. It seamlessly integrates OpenSearch ecosystem tools, such as Cerebro, OpenSearch Dashboards, and Logstash; and it supports mainstream client languages such as Python, Java, Go, and C++, simplifying development and integration.

## How It Works

CSS vector search uses approximate nearest neighbor (ANN) search to mitigate the intensive computational load of k-Nearest Neighbors (k-NN) search, perfectly balancing search efficiency and accuracy. Key points include:

- Reducing the candidate set: Traditional text search filters out irrelevant documents through inverted indexes, while vector search quickly retrieves potentially relevant vectors using index structures like HNSW graphs or IVF-PQ, avoiding scanning the full dataset. For example, HNSW uses a multi-layer structure to quickly find nearest neighbors to the query vector.
- Reducing computational complexity: The funnel model first performs coarse quantization (such as IVF-PQ) on vectors to quickly obtain the candidate set. Then, it performs fine retrieval or reranking (such as cosine similarity) on the candidate set. Quantization uses the product quantization (PQ) algorithm to encode high-dimensional vectors into compact codes, reducing storage and computational overheads.
- Balancing performance and accuracy: Indexing parameters (such as the number of HNSW layers and the number of IVF clusters) can be dynamically tuned to trade off between recall and query latency.

## Procedure

1. Data preparation: Use an AI model (such as CNN and Transformer) to process your unstructured data (such as images, videos, and text) and extract feature vectors.
2. Index creation: Create vector indexes in your OpenSearch cluster and define vector field mappings, including specifying vector dimensions, indexing algorithms, and similarity measurement methods.
3. Data write: Store feature vectors (typically along with the original data or metadata) into these indexes.
4. Vector search: Use the standard OpenSearch query DSL (such as KNN query) to provide the query vector (generated by the same model), and specify the number ( $k$ ) of nearest neighbors you expect to return.
5. Result: The CSS vector search engine performs an efficient ANN search and returns the  $k$  most relevant results and their similarity scores. Your application can then process these results (for example, showing similar images or recommending relevant products).

## Constraints

The built-in CSS vector search engine is available for OpenSearch 1.3.6 and 2.19.0 clusters only.

## Related Documents

- CSS's OpenSearch clusters support the open-source vector search. For details, see [k-NN query](#).
- To learn how to quickly get started with CSS's vector search service, see [Using Elasticsearch for Vector Search](#).
- To learn more about the CSS vector database performance, see [Testing the Performance of CSS's Elasticsearch Vector Search](#).

### 3.8.2.2 Creating an OpenSearch Vector Cluster

Integrating efficient indexing techniques, the CSS vector database delivers a high-performance, low-cost, scalable solution for high-dimensional vector search. An OpenSearch vector cluster converts unstructured data into high-dimensional vectors and uses vector indexing algorithms (such as HNSW graph-based indexing and product quantization) to enable approximate nearest neighbor (ANN) search. This significantly reduces computational complexity while ensuring a high recall rate.

This topic focus on the memory capacity requirements and planning for an OpenSearch vector cluster, so as to provide guidance on how to choose cluster nodes of the appropriate specifications. All other procedures are the same as those for other regular search clusters. For details, see [OpenSearch Cluster Planning Suggestions](#).

## Memory Planning

Before creating an OpenSearch vector cluster, properly plan the cluster's memory capacity based on the data size, vector dimensions, and index types.

- **Node type**

The CSS vector search engine relies heavily on in-memory computing, and vector indexing relies on off-heap memory. For a vector cluster, we recommend **memory-optimized** nodes.

- **Memory calculation formula**

Estimate the required off-heap memory based on the index algorithms used, vector dimensions, and data volume.

- FLAT index:  $\text{mem\_size} = \text{dim} \times \text{dim\_size} \times \text{num} + \text{delta}$
- GRAPH index:  $\text{mem\_size} = (\text{dim} \times \text{dim\_size} + \text{neighbors} \times 4) \times \text{num} + \text{delta}$
- GRAPH\_PQ index:  $\text{mem\_size} = (\text{fragment\_num} + \text{neighbors} \times 4) \times \text{num} + \text{delta}$
- GRAPH\_SQ8 index:  $\text{mem\_size} = (\text{dim} \times 2 + \text{neighbors} \times 4) \times \text{num} + \text{delta}$
- GRAPH\_SQ4 index:  $\text{mem\_size} = (\text{dim} + \text{neighbors} \times 4) \times \text{num} + \text{delta}$
- IVF\_GRAPH and IVF\_GRAPH\_PQ indexes do not require resident memory, so there is need to estimate their memory requirements.

**Table 3-72** Parameter description

| Parameter    | Description                                                                                                                                                                                                                                                                                               |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| mem_size     | Size of off-heap memory required by vector indexes. (This estimate assumes there are no replicas. If index replicas are involved, the required memory size multiplies with the number of replicas.)                                                                                                       |
| dim          | Number of vector dimensions.                                                                                                                                                                                                                                                                              |
| dim_size     | Number of bytes required by each dimension. By default, each dimension is a 4-byte float value.                                                                                                                                                                                                           |
| neighbors    | Number of neighbors for each vector in a graph index. The default value is 64.                                                                                                                                                                                                                            |
| num          | Total number of vectors.                                                                                                                                                                                                                                                                                  |
| fragment_num | Number of segments a vector is split into when product quantization is used.<br>If fragment_num is not configured during index creation, the value is determined by <b>dim</b> .<br>if dim <= 256:<br>fragment_num = dim / 4<br>elif dim <= 512:<br>fragment_num = dim / 8<br>else :<br>fragment_num = 64 |
| delta        | Metadata size. This factor can be ignored.                                                                                                                                                                                                                                                                |

When selecting node 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's physical memory, and the maximum size is **31 GB**.

Example:

When creating a graph index based on the SIFT10M dataset (128-dimensional vector, 10 million records), if **neighbors** is set to 32, the off-heap memory required is as follows: **mem\_size = (128 x 4 + 32 x 4) x 10,000,000 + delta ≈ 6 GB**. If one index replica is needed, at least 12 GB (6 x 2) of off-heap memory is required. In this case, select one 8U32G or two 8U16G node for the cluster.

## Creating a Cluster

The procedure for creating a vector cluster is the same as that for creating any other regular search cluster. For details, see [Creating an OpenSearch Cluster \(New Version\)](#).

Pay attention to the following key parameters:

- **Cluster Version:** Select 1.3.6 or 2.19.0. Only OpenSearch clusters of this version support the built-in CSS vector search engine.

- **Node Specifications:** Select the node specifications based on the result of cluster memory planning.

## (Optional) Configuring the Circuit Breaker

To mitigate out-of-memory (OOM) errors and maintain optimal vector query performance, a circuit breaker mechanism is employed. When the cluster's off-heap memory usage exceeds a predefined threshold, this mechanism automatically blocks vector data writes to the cluster. The purposes of this mechanism are as follows:

- Preventing memory overload: Write throttling lowers off-heap memory usage.
- Maintaining query performance: Optimal vector query performance can be maintained by preventing memory overload.

The off-heap memory circuit breaker is enabled by default. You can enable or disable it and adjust its threshold based on service requirements. The command is as follows:

```
PUT _cluster/settings
{
 "persistent": {
 "native.cache.circuit_breaker.enabled": "true",
 "native.cache.circuit_breaker.cpu.limit": "80%"
 }
}
```

**Table 3-73** Parameter description

| Parameter                            | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| native.cache.circuit_breaker.enabled | Boolean | Whether to enable the off-heap memory circuit breaker.<br>Value range: <ul style="list-style-type: none"><li>• <b>true</b> (default value): Enable the off-heap memory circuit breaker. When the off-heap memory usage reaches the circuit breaker threshold, write requests are blocked.</li><li>• <b>false</b>: Disable the off-heap memory circuit breaker. OOM errors may occur in case of excessive off-heap memory usage.</li></ul> |



| Parameter                              | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| native.cache.circuit_breaker.cpu.limit | String | <p>Circuit breaker threshold in terms of maximum off-heap memory usage.</p> <p>This parameter is available only when <b>native.cache.circuit_breaker.enabled=true</b>.</p> <p>Value range: a value in percentage</p> <p>Default value: 80%</p> <p>Assume a cluster uses 128 GB memory. The required heap memory is 31 GB, and the default circuit breaker threshold is 80%, then: <math>(128 - 31) \times 80\% = 77.6</math> GB. This means when the off-heap memory usage exceeds 77.6 GB, the circuit breaker is triggered to block write operations.</p> |

### 3.8.2.3 Creating a Vector Index

Create a vector index in your OpenSearch cluster and define a mapping that contains vector fields, including vector dimensions, indexing algorithm, and similarity measurement methods. Then store vectors (typically along with the original data or metadata) into this index.

## Logging In to OpenSearch Dashboards

Log in to OpenSearch Dashboards and go to the command execution page. OpenSearch clusters support multiple access methods. This topic uses OpenSearch Dashboards as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
4. In the left navigation pane, choose **Dev Tools**.

## Creating a Vector Index

Run the following command on OpenSearch Dashboards to create a vector index.

For example, create an index named **my\_index**. This index contains a vector field named **my\_vector** and a text field named **my\_label**. A graph-based index is created for the vector field, and Euclidean distance is used for similarity measurement.

```
PUT my_index
{
 "settings": {
 "index": {
 "vector": true,
```

```
{
 "number_of_shards": 1,
 "number_of_replicas": 1
},
"mappings": {
 "properties": {
 "my_vector": {
 "type": "vector",
 "dimension": 2,
 "indexing": true,
 "algorithm": "GRAPH",
 "metric": "euclidean"
 },
 "my_label": {
 "type": "keyword"
 }
 }
}
}
```

**Table 3-74** settings parameters

| Parameter                           | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------|-----------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.vector                        | Yes       | Boolean | Whether to enable vector indexes.<br>Set this parameter to <b>true</b> .<br>Otherwise, vector indexes cannot be created.                                                                                                                                                                                                                                                                                                                  |
| index.number_of_shards              | No        | Integer | Number of index shards. This value should be divisible by the number of cluster nodes.<br>Value range: 1–1024<br>Default value: <b>1</b>                                                                                                                                                                                                                                                                                                  |
| index.number_of_replicas            | No        | Integer | Number of index replicas. Replicas improve data availability.<br>Value range: 0 to the number of nodes minus 1<br>Default value: <b>1</b>                                                                                                                                                                                                                                                                                                 |
| index.vector.exact_search_threshold | No        | Integer | Threshold for automatically switching from pre-filtering search to brute-force search. When the size of the intermediate result set in a segment is lower than this threshold, a brute-force search is performed.<br>Value range: null (disables auto switchover from pre-filtering search to brute-force search) or a positive integer<br>Default value: null (disables auto switchover from pre-filtering search to brute-force search) |

**Table 3-75** mappings parameters

| Parameter     | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------|-----------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| type          | Yes       | String  | Data type of a field.<br>Set this parameter to <b>vector</b> to indicate vector fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| dimension     | Yes       | Integer | Number of vector dimensions.<br>Value range: 1–4096                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| indexing      | No        | Boolean | Whether to enable vector index acceleration.<br>Value range: <ul style="list-style-type: none"><li>• <b>true</b>: Enables vector index acceleration. When this parameter is set to <b>true</b>, an extra vector index is created. The indexing algorithm is specified by the <b>algorithm</b> field and the index supports vector search.</li><li>• <b>false</b>: Disables vector index acceleration. When this parameter is set to <b>false</b>, vector data is written only to docvalues, and only <b>ScriptScore</b> and <b>Rescore</b> can be used for vector search.</li></ul>                                                                                                 |
| lazy_indexing | No        | Boolean | Whether to enable lazy indexing, where vector indexes are not built immediately when data is ingested. The aim to speed up write operations.<br>Lazy indexing takes effect only when <b>indexing</b> is set to <b>true</b> and the cluster version is OpenSearch 2.19.0.<br>Value range: <ul style="list-style-type: none"><li>• <b>true</b>: Enables lazy indexing. Vector indexes (such as graph-based indexes) are not built immediately when data is ingested. Instead, <b>offline index building</b> needs to be performed to build vector indexes. The vector indexes can then be used for vector search.</li><li>• <b>false</b> (default): Disables lazy indexing.</li></ul> |

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| algorithm | No        | String | <p>Vector indexing algorithm.</p> <p>This parameter is valid only when <b>indexing</b> is set to <b>true</b>. When this parameter is set to <b>IVF_GRAPH</b> or <b>IVF_GRAPH_PQ</b>, <b>(Optional) Pre-Building and Registering Centroid Vectors</b> is required.</p> <p>Value range:</p> <ul style="list-style-type: none"> <li>• <b>FLAT</b>: 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 can reach 100%. You can use this algorithm if you require high recall accuracy.</li> <li>• <b>GRAPH</b> (default): Hierarchical Navigable Small Worlds (HNSW) algorithm for graph-based indexes. This algorithm is mainly used when high performance and precision are required and the number of documents in a single shard reaches 10 million.</li> <li>• <b>GRAPH_PQ</b>: a combination of the HNSW algorithm and the PQ algorithm. The PQ algorithm reduces the storage overhead of the original vectors, so that HNSW can easily search through hundreds of millions of records.</li> <li>• <b>GRAPH_SQ8</b>: a combination of the HNSW algorithm and the scalar quantization (SQ) algorithm. By quantizing float32 values into int8, this algorithm reduces the storage overhead of the original vectors and improves build and query efficiency. The downside is a slightly decreased recall rate. Only OpenSearch 2.19.0 clusters support this algorithm.</li> <li>• <b>GRAPH_SQ4</b>: a combination of the HNSW algorithm and the SQ algorithm. By quantizing float32 values into int4, this algorithm</li> </ul> |

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------|-----------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           |           |        | <p>reduces the storage overhead of the original vectors and improves build and query efficiency. The downside is a slightly decreased recall rate. SQ4 has a higher quantization/compression ratio and higher computational efficiency than SQ8, but also a large decrease in recall rate. Only OpenSearch 2.19.0 clusters support this algorithm.</p> <ul style="list-style-type: none"><li>• <b>IVF_GRAPH</b>: a 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 through hundreds of millions of records.</li><li>• <b>IVF_GRAPH_PQ</b>: a 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 when there are more than 1 billion documents in shards and high retrieval performance is required.</li></ul> <p>When the indexing algorithm is set to GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4, CSS provides additional parameters, as shown in <a href="#">Table 3-76</a> and <a href="#">Table 3-77</a>, that you can choose to configure for enhanced query performance and accuracy.</p> |
| dim_type  | No        | String | <p>Vector data type.</p> <p>Value range:</p> <ul style="list-style-type: none"><li>• binary: binary value</li><li>• float (default): floating-point number</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| metric    | No        | String | <p>Vector distance metric, which measures the similarity or distance between vectors.</p> <p>Value range:</p> <ul style="list-style-type: none"><li>• <b>euclidean</b> (default): Euclidean distance</li><li>• <b>inner_product</b>: inner product distance</li><li>• <b>cosine</b>: cosine distance</li><li>• <b>hamming</b>: Hamming distance, which can be used only when <b>dim_type</b> is set to <b>binary</b>.</li></ul> |

**Table 3-76** Optional parameters for the GRAPH indexing algorithm

| Parameter | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                          |
|-----------|-----------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| neighbors | No        | Integer | <p>Number of neighbors of each vector in the graph index. A larger value results in higher query accuracy but slower index building and query.</p> <p>This parameter is available only when <b>indexing</b> is set to <b>true</b>, and <b>algorithm</b> is GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 20–255</p> <p>Default value: <b>64</b></p>                  |
| shrink    | No        | Float   | <p>How aggressively the HNSW graph removes redundant edges during construction. This setting affects the structure of the HNSW graph.</p> <p>Default value: 1.0f</p> <p>This parameter is available only when <b>indexing</b> is set to <b>true</b>, and <b>algorithm</b> is GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 0.1–10</p> <p>Default value: <b>1</b></p> |

| Parameter    | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                            |
|--------------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| scaling      | No        | Integer | <p>Scaling ratio for the number of upper-layer graph nodes in the HNSW graph. This setting affects the layers of the HNSW graph.</p> <p>This parameter is available only when <b>indexing</b> is set to <b>true</b>, and <b>algorithm</b> is GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 0-128</p> <p>Default value: 50</p>                                          |
| efc          | No        | Integer | <p>How many nearest neighbors to explore when inserting a new vector into the HNSW graph. A larger value results in higher accuracy but slower index building.</p> <p>This parameter is available only when <b>indexing</b> is set to <b>true</b>, and <b>algorithm</b> is GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 0-100000</p> <p>Default value: <b>200</b></p> |
| max_scan_num | No        | Integer | <p>Maximum number of nodes to be scanned during search or index building. A larger value results in higher query accuracy but slower indexing.</p> <p>This parameter is available only when <b>indexing</b> is set to <b>true</b>, and <b>algorithm</b> is GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 0 to 1000000</p> <p>Default value: <b>10000</b></p>           |

**Table 3-77** Optional parameters for the GRAPH\_PQ indexing algorithm

| Parameter    | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| centroid_num | No        | Integer | Number of centroids used during the coarse quantization stage of the algorithm. It affects quantization granularity and storage.<br><br>This parameter is available only when <b>indexing</b> is set to <b>true</b> , and <b>algorithm</b> is GRAPH_PQ.<br><br>Value range: 0–65535<br>Default value: 255                                                                                                                                                                              |
| fragment_num | No        | Integer | Number of fragments each vector is split into. It affects the PQ quantization granularity.<br><br>The default value is <b>0</b> . The plug-in automatically sets the number of fragments based on the vector length.<br><br>This parameter is available only when <b>indexing</b> is set to <b>true</b> , and <b>algorithm</b> is GRAPH_PQ.<br><br>Value range: 0–4096<br>Default value: <b>0</b> . The plug-in automatically sets the number of fragments based on the vector length. |

## (Optional) Pre-Building and Registering Centroid Vectors

When the IVF\_GRAPH or IVF\_GRAPH\_PQ algorithm is used for vector indexing, you need to pre-build and register centroid vectors.

IVF\_GRAPH and IVF\_GRAPH\_PQ help to accelerate indexing and queries in ultra-large-scale clusters. These two algorithms allow you to narrow down the query scope by dividing a vector space into subspaces through clustering or random sampling. Before pre-build, you need to obtain all centroid vectors through clustering or random sampling. Centroid vectors are pre-built into a GRAPH or GRAPH\_PQ index and then registered with the Elasticsearch cluster. All nodes in the cluster can share this index. Reuse of the centroid index among shards can effectively reduce the training overhead and the number of centroid index queries, improving write and query performance.

1. Create a centroid index.

For example, run the following command on OpenSearch Dashboards to create a centroid index named **my\_dict**:

```
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"
 }
 }
}
```

For detailed parameter configuration, see [Creating a Vector Index](#). Pay attention to the following mandatory parameters:

- **index.number\_of\_shards**: The number of index shards must be set to **1**. Otherwise, the centroid index cannot be registered.
- **indexing**: This parameter must be set to **true** to enable vector index acceleration.
- **algorithm**: Set the indexing algorithm. Set it to **GRAPH** for the IVF\_GRAPH algorithm, and GRAPH\_PQ if the IVF\_GRAPH\_PQ algorithm is used.

2. Write centroid vectors to the created index. Write the centroid vectors obtained through sampling or clustering into the newly created index **my\_dict**.
3. Call the registration API.

Run the following command on OpenSearch Dashboards to register the centroid index as a Dict object with a globally unique name (**dict\_name**):

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

4. Create an IVF\_GRAPH or IVF\_GRAPH\_PQ vector index.

When creating the vector index, you do not need to specify **dimension** or **metric**. Rather, you specify the registered Dict object. [Table 3-78](#) describes key parameters for specifying a Dict object. For details about other parameters, see [Creating a Vector Index](#).

For example, run the following command to create an IVF\_GRAPH vector index:

```
PUT my_index
{
 "settings": {
 "index": {
 "vector": true,
 "sort.field": "my_vector.centroid" # Set the centroid subfield of each vector field as a ranking field.
 }
 },
 "mappings": {
 "properties": {
 "my_vector": {
 "type": "vector",
 "indexing": true,

```

```
"algorithm": "IVF_GRAPH",
"dict_name": "my_dict",
"offload_ivf": true
}
}
}
```

**Table 3-78** Field mappings parameters

| Parameter   | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------|-----------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dict_name   | Yes       | String  | Name of the centroid index. For example, <b>my_dict</b> . The vector dimensions and metrics of the index must be the same as those of the Dict index.                                                                                                                                                                                                                                                                                               |
| offload_ivf | Yes       | Boolean | Whether to offload the IVF inverted index to Elasticsearch.<br>Value: <b>true</b> or <b>false</b> . <ul style="list-style-type: none"><li>• <b>true</b> (recommended value): Offloads the IVF inverted index implemented by the underlying index to Elasticsearch. This reduces the use of off-heap memory and the overhead of write and merge operations.</li><li>• <b>false</b> (default value): Not to offload the IVF inverted index.</li></ul> |

### 3.8.2.4 Importing Vector Data

Importing vector data is the process of ingesting data to the CSS vector database. When writing vector data to a vector index, you need to specify the vector field (for example, **my\_vector**) and the corresponding data format. The CSS vector database supports two common formats: floating-point arrays and Base64.

- Floating-point arrays: transmitting readable arrays directly.
- Base64: encoding vectors (little-endian byte order) into character strings to reduce network transmission overhead and improve efficiency in handling high-dimensional/binary vectors.

Choose a format based on the characteristics of your data. Also, choose an appropriate data importing method.

- Importing a single record: Use for small-scale applications or testing.
- Bulk import: Use for large-scale applications, where write requests are merged to reduce network overhead.

## Constraints

- Ensure that the vector field names and vector dimensions are consistent with those defined for the index.
- Base64 encoding must use the little-endian byte order. Otherwise, parsing errors may occur.
- In the case of bulk imports, you are advised to submit 100 to 1000 records per request. This balances throughput and memory overhead.

## Importing a Single Record

- Floating-point array

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

- Base64

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

## Bulk Import

- Floating-point array

```
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"}
```

- Base64

```
POST my_index/_bulk
{"index": {}}
{"my_vector": "AACAPwAAAEA=", "my_label": "red"}
{"index": {}}
{"my_vector": "AAAAQAAAAEA=", "my_label": "green"}
{"index": {}}
{"my_vector": "AAAAQAAQEA=", "my_label": "red"}
```

For details about how to use the Bulk API, see [Bulk API](#).

## (Optional) Post-processing after Data Ingestion: Offline Index Building

### WARNING

- Use offline index creation via an API only when real-time data is not required or crucial and the cluster version is OpenSearch 2.19.0.
- If lazy\_indexing is enabled, offline index building must be performed after data ingestion. Otherwise, the system will return error code 500 for standard vector query, with the error message "Load native index failed exception." To solve this problem, perform offline index building before vectors queries.

OpenSearch uses an LSM (Log-Structured Merge) tree-like model to accelerate write operations. As data is continuously written in and updated, numerous small

index segments are generated and later merged via a backend task to enhance query performance. As vector indexing is computationally intensive, frequent index merging while vector data is being written in consumes significant CPU resources. Therefore, where real-time data is not crucial, it is advisable to set **lazy\_indexing** to **true** for vector fields. This allows a final vector index to be created via a non-real time API after all data has been written in. This approach significantly reduces index merges, thereby improving overall write and index merging performance.

Offline index building consists of two steps:

1. Merge index segments.
2. Create the final vector index based on the final index segments.

The API used for offline index building is as follows:

```
POST _vector/indexing/{index_name}
{
 "field": "{field_name}"
}
```

where, {index\_name} indicates the name of the index to create. {field\_name} indicates the name of the vector field for which **lazy\_indexing** has been set to **true**.

### 3.8.2.5 Performing Vector Search

The CSS vector database supports a variety of query methods, including standard query, hybrid query, script\_score query, rescore query, and painless syntax extension. They enable efficient vector data search by accommodating diverse search needs.

- **Standard Query**: retrieves documents that are most similar to the query vector.
- **Hybrid Query**: combines vector search with traditional OpenSearch queries, such as pre-filtering and Boolean queries.
- **Script Score Query**: enables custom similarity calculations for vector searches by executing a custom script.
- **Rescore Query**: rescores and reranks the top results returned by an initial query to improve recall.
- **Painless Syntax Extension**: allows the use of vector distance or similarity calculation functions in custom scripts.

#### Standard Query

Standard query is used to retrieve documents that are most similar to the query vector.

The following command will return  $k$  (specified by **size/topk**) records that are the closest matches to the query vector.

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

```
"topk":2
}
}
}
}
```

**Table 3-79** Parameters for standard query

| Parameter              | Mandatory | Type         | Description                                                                                                                                                                                                                                                                                                              |
|------------------------|-----------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| size                   | Yes       | Integer      | Number of search results to return.<br>Default value: <b>10</b>                                                                                                                                                                                                                                                          |
| _source                | No        | Boolean      | Whether to return the source text in documents. To reduce data transmission and improve query performance, set this parameter to <b>false</b> .<br>Value range: <ul style="list-style-type: none"><li>• <b>true</b> (default): Returns the source text.</li><li>• <b>false</b>: Not to return the source text.</li></ul> |
| query                  | Yes       | Map          | Specifies the query vector.<br>Parameters:<br><b>vector</b> (mandatory): indicates a vector query (vector similarity-based search), including the vector field and query vector value.<br><b>my_vector</b> (mandatory): queried vector field (for example, <b>my_vector</b> ).                                           |
| vector (sub-parameter) | Yes       | Array/String | Query vector value. It is used to calculate the similarity between indexed vectors and the query vector.<br>The value can be an array (for example, [1, 1]) or Base64-encoded value (for example, <b>AAABAAACAAAD</b> ).                                                                                                 |
| topk                   | Yes       | Integer      | The number of the most similar or relevant results to be returned.<br>Default value: same as <b>size</b> .                                                                                                                                                                                                               |

| Parameter    | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                       |
|--------------|-----------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ef           | No        | Integer | <p>How many nearest neighbors to explore when inserting a new vector into the graph. A larger value indicates a higher query accuracy yet slower query speed.</p> <p>This parameter is available only when <b>algorithm</b> is set to GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 0-100000</p> <p>Default value: <b>200</b></p> |
| max_scan_num | No        | Integer | <p>Maximum number of graph nodes to scan during search. A larger value indicates a higher query accuracy yet slower query speed.</p> <p>This parameter is available only when <b>algorithm</b> is set to GRAPH, GRAPH_PQ, GRAPH_SQ8, or GRAPH_SQ4.</p> <p>Value range: 0-1000000</p> <p>Default value: <b>10000</b></p>                           |
| nprobe       | No        | Integer | <p>Number of centroids to explore during an IVF index query. A larger value indicates a higher query accuracy yet slower query speed.</p> <p>This parameter is available only when <b>algorithm</b> is IVF_GRAPH or IVF_GRAPH_PQ.</p> <p>Value range: 0-100000</p> <p>Default value: <b>100</b></p>                                               |

## Hybrid Query

Hybrid query combines vector search with traditional OpenSearch queries, such as pre-filtering and Boolean queries.

---

### WARNING

Only OpenSearch 2.19.0 clusters support pre-filtering.

---

In the following example, the top 10 records whose **my\_label** value is **red** are returned.

- **Pre-filtering query**

First, filters are applied to retrieve matching results. Then, vector search is performed on these results to retrieve the most relevant vectors based on similarity.

The following is an example:

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

**Table 3-80** Parameters for pre-filtering query

| Parameter | Mandatory | Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------|-----------|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| filter    | Yes       | Map  | Vector query filters. Standard OpenSearch query filters are supported, such as term and range. If <b>filter</b> is too restrictive, leading to a small intermediate result set, you can set the <b>index.vector.exact_search_threshold</b> parameter, so that when the intermediate result set is smaller than this threshold, pre-filtering query automatically switches over to brute-force query (FLAT algorithm), which ensures a high recall rate. For more information, see <a href="#">Creating a Vector Index</a> . |
| term      | No        | Map  | Term query is a type of exact query. Documents that contain the exact term will be returned. For example, <b>{"term": {"my_label": "red"}}</b> means only to return documents whose <b>my_label</b> value is <b>red</b> .                                                                                                                                                                                                                                                                                                   |

- **Boolean query**

A Boolean query is in fact a post-filtering query method. Filtering and vector similarity-based search are performed separately. Then, the results of the two are combined using Boolean logic defined by clauses like **must**, **should**, and **filter**.

The following is an example:

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

**Table 3-81** Boolean query parameters

| Parameter | Mandatory | Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------|-----------|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| bool      | Yes       | Map  | <p>A compound query clause that combines subqueries using configured Boolean logic.</p> <p>Parameter description:</p> <ul style="list-style-type: none"><li>• <b>must</b>: Clauses that must match for documents to be included in the results.</li><li>• <b>filter</b>: It is similar to <b>must</b>, but do not contribute to the relevance score.</li><li>• <b>should</b>: Clauses that should match, but are not required. They are like nice-to-haves.</li><li>• <b>must_not</b>: Clauses that must not match for documents to be included in the results.</li></ul> |
| bool.must | Yes       | Map  | <p>Clauses that must match for documents to be included in the results. Parameter description:</p> <ul style="list-style-type: none"><li>• <b>vector</b>: query vector</li><li>• <b>my_vector</b>: vector field</li><li>• <b>topk</b>: number of results to return</li></ul>                                                                                                                                                                                                                                                                                              |



| Parameter   | Mandatory | Type | Description                                                                                                                                     |
|-------------|-----------|------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| bool.filter | Yes       | Map  | Clauses that must match, but do not contribute to the relevance score. Standard OpenSearch query filters are supported, such as term and range. |

## Script Score Query

Script\_score query enables custom similarity calculations for vector searches by executing a user-defined script. It works as follows:

Pre-filtering works with any query. **script\_score** then calculates vector similarity on the pre-filtered results, and ranks the results. This query method does not use vector indexes. Its performance depends on the size of the intermediate result set after the pre-filtering. If the pre-filtering condition is set to **match\_all**, a brute-force search is performed on all data.

The following is an example:

```
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 3-82** script\_score query parameters

| Parameter    | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------|-----------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| script_score | Yes       | Map              | A root parameter for the script_score query.<br>Parameter description: <ul style="list-style-type: none"><li>• <b>query</b>: pre-filtering criteria. When it is set to <b>match_all</b>, a brute-force search is performed on all data.</li><li>• <b>script</b>: custom script that calculates similarity scores.</li></ul>                                                                                          |
| source       | Yes       | String           | Script name. The value is fixed to <b>vector_score</b> , indicating that a built-in script is used for calculating similarity.                                                                                                                                                                                                                                                                                       |
| lang         | Yes       | String           | Script language type. The value is fixed to <b>vector</b> .                                                                                                                                                                                                                                                                                                                                                          |
| field        | Yes       | String           | Queried vector field, for example, <b>my_vector</b> .                                                                                                                                                                                                                                                                                                                                                                |
| vector       | Yes       | Array/<br>String | Query vector value. It is used to calculate the similarity between indexed vectors and the query vector.<br>The value can be an array (for example, [1, 1]) or Base64-encoded value (for example, <b>AAABAAACAAAD</b> ).                                                                                                                                                                                             |
| metric       | Yes       | String           | Vector distance metric, which measures the similarity or distance between vectors.<br>Value range: <ul style="list-style-type: none"><li>• <b>euclidean</b> (default): Euclidean distance</li><li>• <b>inner_product</b>: inner product distance</li><li>• <b>cosine</b>: cosine distance</li><li>• <b>hamming</b>: Hamming distance, which can be used only when <b>dim_type</b> is set to <b>binary</b>.</li></ul> |

## Rescore Query

Rescore query rescores and reranks the top results returned by an initial query to improve recall.

When the GRAPH\_PQ or IVF\_GRAPH\_PQ indexing algorithm is used, query results are ranked based on the asymmetric distance calculated by PQ. Rescore query then rescores and reranks the initial search results to improve recall.

The following is an example of rescore query on a PQ index named **my\_index**:

```
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 3-83** Rescore query parameters

| Parameter   | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                               |
|-------------|-----------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rescore     | Yes       | Map     | Defines rescoring parameters.<br>Key parameters: <ul style="list-style-type: none"><li>• <code>window_size</code>: rescoring/reranking window size.</li><li>• <code>vector_rescore</code>: other vector rescoring settings.</li></ul>                                                                                     |
| window_size | Yes       | Integer | Rescoring/reranking window size. The vector search returns the top k results, but only the first <i>window_size</i> results are rescored and reranked. A larger value indicates a larger reranking scope and hence a higher recall rate, but it also leads to higher computational overhead.<br>Default value: <b>100</b> |
| field       | Yes       | String  | Queried vector field, for example, <b>my_vector</b> .                                                                                                                                                                                                                                                                     |

| Parameter | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------|-----------|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| vector    | Yes       | Array/<br>String | Query vector value. It is used to calculate the similarity between indexed vectors and the query vector.<br><br>The value can be an array (for example, [1, 1]) or Base64-encoded value (for example, <b>AAABAAACAAAD</b> ).                                                                                                                                                                                             |
| metric    | Yes       | String           | Vector distance metric, which measures the similarity or distance between vectors.<br><br>Value range: <ul style="list-style-type: none"><li>• <b>euclidean</b> (default): Euclidean distance</li><li>• <b>inner_product</b>: inner product distance</li><li>• <b>cosine</b>: cosine distance</li><li>• <b>hamming</b>: Hamming distance, which can be used only when <b>dim_type</b> is set to <b>binary</b>.</li></ul> |

## Painless Syntax Extension

Painless syntax extension allows the use of vector distance or similarity calculation functions in custom scripts. CSS extension supports several vector distance/similarity calculation functions, which users can use readily in custom Painless scripts to build flexible rescoring formulas.

The following is an example:

```
POST my_index/_search
{
 "size": 10,
 "query": {
 "script_score": {
 "query": {
 "match_all": {}
 },
 "script": {
 "source": "1 / (1 + euclidean(params.vector, doc[params.field]))",
 "params": {
 "field": "my_vector",
 "vector": [1, 2]
 }
 }
 }
 }
}
```

**Table 3-84** Supported vector distance/similarity calculation functions

| Function Signature               | Description                                                                                                                                                                                                     |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| euclidean(Float[], DocValues)    | Euclidean distance                                                                                                                                                                                              |
| cosine(Float[], DocValues)       | Cosine similarity                                                                                                                                                                                               |
| innerproduct(Float[], DocValues) | Inner product                                                                                                                                                                                                   |
| hamming(String, DocValues)       | Hamming distance Only vectors whose <b>dim_type</b> is <b>binary</b> are supported. The input query vector must be a Base64-encoded character string.<br>Only OpenSearch 1.3.6 clusters support this function.  |
| hammings(String, DocValues)      | Hamming distance Only vectors whose <b>dim_type</b> is <b>binary</b> are supported. The input query vector must be a Base64-encoded character string.<br>Only OpenSearch 2.19.0 clusters support this function. |

### 3.8.2.6 Optimizing Vector Cluster Performance

This topic explains how to optimize the performance of a CSS vector database from two aspects—write and query.

#### Optimizing Write Performance

Writing vector data incurs three major overheads: replica synchronization, index refresh, and segment merging. When index data is written in real time, frequent index refresh operations generate a large number of small segments. This triggers frequent vector index build and merge operations, which consume excessive CPU/IO resources. You can try the following solutions to optimize write performance.

Solution 1: temporarily disable replicas

- Description  
Temporarily disable replicas during data ingestion and enable them after data ingestion is complete. Use this solution when importing historical data in batches or performing a full update (for example, when initializing a vector database).
- Operation  
Set the number of replicas:

```
PUT my_index/_settings
{
 "number_of_replicas": 0
}
```

- **Result**  
Write performance is enhanced by avoiding real-time vector index building on replica nodes.

Solution 2: adjust the refresh interval

- **Description**  
Set the index refresh interval to 120s or longer to reduce the number of small segments generated during frequent index refreshes and also reduce the vector index building overhead caused by segment merges. You can also disable automatic index refresh by setting the refresh interval to -1. Use this solution in high-throughput write scenarios (for example, when writing vectorized log data).

- **Operation**  
Set **refresh\_interval**.

```
PUT my_index/_settings
{
 "refresh_interval": "120s"
}
```

- **Result**  
Index refreshes occur less frequently. This reduces the number of small segments and also the overhead of segment merges, leading to enhanced write performance.

Solution 3: increase indexing threads

- **Description**  
Increasing the number of threads for vector index building accelerates the indexing process. However, too many such threads will compete for query resources. Use this solution when there are sufficient CPU resources but the write latency is high—such as in GPU-accelerated environments.

- **Operation**  
The default value of **native.vector.index\_threads** is 4. Change this value as needed.

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

- **Result**  
Vector index building is accelerated, and the performance of concurrent writes is enhanced.

## Optimizing Query Performance

Query performance is affected by the following factors: the number of segments, the memory circuit breaker mechanism, and field recall. An excessively large number of segments impacts search efficiency; when off-heap memory becomes insufficient, vector index data is frequently swapped in and out of the memory; recalling all fields increases the load during the fetch phase. You can optimize query performance by addressing these factors.

Solution 1: perform force merge

- **Description**  
After batch data ingestion, perform the force merge operation to forcibly merge segments, thus reducing the number of segments. Typically, you should perform this operation after data ingestion and before data query (for example, after a scheduled batch ingestion).
- **Operation**  
Perform the force merge operation:  

```
POST my_index/_forcemerge?max_num_segments=1
```
- **Result**  
Multiple segments are merged into a single segment. This reduces the file scanning overhead and accelerates the query speed.

Solution 2: adjust the upper limit of the segment size

- **Description**  
During batch writes, the maximum size of segments generated by the system is 5 GB. You can increase this upper limit to reduce the number of segments generated after automatic merging. Typically, you should perform this operation before batch data ingestion starts.
- **Operation**  
Increase the maximum segment size:  

```
PUT my_index/_settings
{
 "index.merge.policy.max_merged_segment": "10gb"
}
```
- **Result**  
Increasing the maximum segment size helps to reduce the number of segments and thus accelerate query performance.

Solution 3: adjust the circuit breaker limit for off-heap memory

- **Description**  
When the off-heap memory required by vector indexes exceeds the circuit breaker limit, the index cache manager frequently swaps in and out index data from the cache, which slows down queries. You can raise the circuit breaker limit to reduce circuit breaking (indicated by `CircuitBreakingException` recorded in the log) resulted from insufficient memory.
- **Operation**  
The default circuit breaker limit for off-heap memory is 80%. You can adjust this limit as required.  

```
PUT _cluster/settings
{
 "persistent": {
 "native.cache.circuit_breaker.cpu.limit": "85%"
 }
}
```
- **Result**  
It is less likely for vector index data to be swapped out from the memory, and query jitter is reduced.

Solution 4: optimize field recall

- Description

If the query result needs to return only a few fields that are either keywords or numeric values, you can use the `docvalue_fields` parameter to fetch them. Use this method if only numeric or enumerated metadata (such as product IDs and class labels) needs to be fetched. It can significantly reduce overhead during the fetch phase.

- Operation

Use the `docvalue_fields` parameter to fetch only specific fields:

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

- Result

There is no need to parse the entire `_source` document. Column-oriented storage (docvalues) reduces the overhead during the fetch phase and improves query performance.

## Setting Cache Timeout

When the cluster's memory resources are insufficient, data is frequently updated, or high data freshness is required, you can enable automatic cache expiration to have inactive data cleared from the cache. This helps to optimize system performance, ensure data consistency, and improve query stability. Use this approach where data updates frequently or memory resources are stretching thin.

Run the following command to set cache timeout:

```
PUT _cluster/settings
{
 "persistent": {
 "native.cache.expiry.enabled": "true",
 "native.cache.expiry.time": "30m"
 }
}
```



**Table 3-85** Parameter description

| Parameter                   | Type    | Description                                                                                                                                                                                                                                                                         |
|-----------------------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| native.cache.expiry.enabled | Boolean | Whether to enable automatic cache expiration.<br>Value range: <ul style="list-style-type: none"><li><b>true</b>: Enable automatic cache expiration. Inactive data in the cache will be cleared.</li><li><b>false</b> (default value): Disable automatic cache expiration.</li></ul> |
| native.cache.expiry.time    | String  | Timeout of inactive cache items.<br>This parameter takes effect only when <b>native.cache.expiry.enabled=true</b> .<br>Value: a time string, for example, <b>24h</b> (24 hours) or <b>30m</b> (30 minutes).<br>Default value: <b>24h</b> .                                          |

### 3.8.2.7 Managing the Vector Index Cache

The CSS vector search engine is built on C++. It uses off-heap memory to enhance performance and efficiency. To better manage and optimize the vector index cache, the CSS vector database provides a number of dedicated APIs that can be used to flexibly monitor and tune cache usage to ensure stable query performance.

### Checking Cache Statistics

Use the following API to check the off-heap memory usage and return key metrics such as the current off-heap memory usage, cache hits, and data loading times. These metrics help you understand the cache status, so you can adjust the cache based on its usage. In the implementation of the vector search plug-in, each segment constructs and stores an index file. During query, all these files are loaded to the off-heap memory, and the plug-in uses a cache mechanism to manage the off-heap memory.

GET /\_vector/stats

Example response:

```
{
 "_nodes" : { # Node information
 "total" : 1, # Total number of nodes
 "successful" : 1, # Number of successful nodes
 "failed" : 0 # Number of failed nodes
 },
 "cluster_name" : "css-d3a7", # Cluster name
 "cpu_circuit_breaker_triggered" : false, # Whether circuit breaking is triggered
 "nodes" : {
 "cAHmVUZTR9ON7t6jxcDCkg" : { # Node UUID
 "cpu_cache_capacity_reached" : false, # Whether the off-heap memory usage of the current node
 reaches the upper limit
 "cpu_eviction_count" : 0, # Number of segment-level cache swap-outs on the current node
 }
 }
}
```

```
"cpu_hit_count" : 0, # Number of segment-level cache hits on the current node
"cpu_load_exception_count" : 0, # Number of segment-level index loading failures on the current
node
"cpu_load_success_count" : 0, # Number of segment-level index loading successes on the current
node
"cpu_miss_count" : 0, # Number of segment-level cache misses on the current node
"cpu_query_memory_usage" : 0, # Off-heap memory usage on the current node, in KB
"cpu_total_load_time" : 0 # Total segment-level index loading time on the current node, in ms
}
}
```

## Preloading Vector Indexes

Use the following API to preload specified vector indexes to the off-heap memory. This ensures that these indexes can be quickly accessed in subsequent queries, with high query efficiency and low latency.

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

where, {index\_name} specifies an index.

Example response:

```
{
 "_shards" : {
 "total" : 1,
 "successful" : 1,
 "failed" : 0
 }
}
```

## Clearing the Cache

When the cache is full, the system automatically swaps its data. However, frequent cache swapping may affect query performance. You can use the following API to manually clear the vector index cache. This way, you can reclaim off-heap memory resources and ensure hot data query performance.

- Clear the full cache:  

```
PUT /_vector/clear/cache
```
- Clear specified indexes from the cache:  

```
PUT /_vector/clear/cache/{index_name}
```

Example response:

```
{
 "acknowledged" : "true"
}
```

### 3.8.2.8 Using Nested Fields for Vector Search

Nested fields allow multiple vectorized records to be stored in a single document. For example, in an RAG scenario, documents usually need to be segmented by paragraph or by a fixed length, and then vectorized into multiple semantic vectors. By means of nested fields, these vectors can be written into a same Elasticsearch document. For a document that contains multiple vector records, if the query vector matches any of them, the document is returned.

## Constraints

Only OpenSearch 2.19.0 clusters support this feature.

## Creating a Vector Index

Create a vector index with nested fields. The index contains an **id** field whose type is **keyword**, and an **embedding** field whose type is **nested**. The embedding field contains two subfields: **chunk** and **emb**. The **chunk** subfield is of the **keyword** type, and the **emb** subfield is of the **vector** type.

```
PUT my_index
{
 "settings": {
 "index.vector": true
 },
 "mappings": {
 "properties": {
 "id": {
 "type": "keyword"
 },
 "embedding": {
 "type": "nested",
 "properties": {
 "chunk": {
 "type": "keyword"
 },
 "emb": {
 "type": "vector",
 "dimension": 2,
 "indexing": true,
 "algorithm": "GRAPH",
 "metric": "euclidean"
 }
 }
 }
 }
 }
}
```

## Importing Vector Data

Use the bulk operation to write data in arrays. Each document contains two vector records.

```
POST my_index/_bulk
{"index":{"id":1,"embedding":[{"chunk":1,"emb": [1, 1]}, {"chunk":2,"emb": [2, 2]}}}
{"index":{"id":2,"embedding":[{"chunk":1,"emb": [2, 2]}, {"chunk":2,"emb": [3, 3]}}}
{"index":{"id":3,"embedding":[{"chunk":1,"emb": [3, 3]}, {"chunk":2,"emb": [4, 4]}}}
```

## Vector Search

The nested query is required for nested fields. To perform such a query, you need to set the path parameter to specify the nested path, and set **score\_mode** to **max**, indicating the maximum similarity between all vectors in the document and the query vector.

- Standard query

Query the top 10 documents that are most similar to vector [1, 1].

```
GET my_index/_search
{
 "_source": {"excludes": ["embedding"]},
 "query": {
 "nested": {
```

```
"path": "embedding",
"score_mode": "max",
"query": {
 "vector": {
 "embedding.emb": {
 "vector": [1, 1],
 "topk": 10
 }
 }
}
```

An example of the query result:

```
{
 "took" : 2,
 "timed_out" : false,
 "_shards" : {
 "total" : 1,
 "successful" : 1,
 "skipped" : 0,
 "failed" : 0
 },
 "hits" : {
 "total" : {
 "value" : 3,
 "relation" : "eq"
 },
 "max_score" : 1.0,
 "hits" : [
 {
 "_index" : "my_index",
 "_type" : "_doc",
 "_id" : "Hc4Vc5QBSxCnghau22AE",
 "_score" : 1.0,
 "_source" : {
 "id" : 1
 }
 },
 {
 "_index" : "my_index",
 "_type" : "_doc",
 "_id" : "Hs4Vc5QBSxCnghau22AE",
 "_score" : 0.33333334,
 "_source" : {
 "id" : 2
 }
 },
 {
 "_index" : "my_index",
 "_type" : "_doc",
 "_id" : "H84Vc5QBSxCnghau22AE",
 "_score" : 0.11111111,
 "_source" : {
 "id" : 3
 }
 }
]
 }
}
```

- Pre-filtering query

First retrieve documents whose ID is ["2", "3"], and then return the top 10 documents that are most similar to the query vector [1, 1].

```
GET my_index/_search
{
 "query": {
 "nested": {
```

```
"path": "embedding",
"score_mode": "max",
"query": {
 "vector": {
 "embedding.emb": {
 "vector": [1, 1],
 "topk": 10,
 "filter": {
 "terms": {"id": ["2", "3"]}
 }
 }
 }
}
```

An example of the query result:

```
{
 "took" : 3,
 "timed_out" : false,
 "_shards" : {
 "total" : 1,
 "successful" : 1,
 "skipped" : 0,
 "failed" : 0
 },
 "hits" : {
 "total" : {
 "value" : 2,
 "relation" : "eq"
 },
 "max_score" : 0.33333334,
 "hits" : [
 {
 "_index" : "my_index",
 "_type" : "_doc",
 "_id" : "3t0ZypcB-Tff59gMTZO2",
 "_score" : 0.33333334,
 "_source" : {
 "id" : 2,
 "embedding" : [
 {
 "chunk" : 1,
 "emb" : [
 2,
 2
]
 },
 {
 "chunk" : 2,
 "emb" : [
 3,
 3
]
 }
]
 }
 },
 {
 "_index" : "my_index",
 "_type" : "_doc",
 "_id" : "390ZypcB-Tff59gMTZO2",
 "_score" : 0.11111111,
 "_source" : {
 "id" : 3,
 "embedding" : [
 {
 "chunk" : 1,
 "emb" : [
```

```
3,
3
]
},
{
 "chunk" : 2,
 "emb" : [
 4,
 4
]
}
]
}
}
]
}
}
```

### 3.8.2.9 Sample Python Code for Vector Search

OpenSearch provides standard REST APIs and clients developed using Java and Python.

This section provides a sample of Python code for creating vector indexes, and importing and querying vector data. It shows how to use the client to implement vector search.

#### Prerequisites

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

```
OpenSearch 2.19.0
pip install opensearch-py==2.8.0

OpenSearch 1.3.6
pip install opensearch-py==1.1.0
```

#### Sample Code

```
from opensearchpy import OpenSearch

Create a client.
def get_client(hosts: list, user: str = None, password: str = None):
 if user and password:
 return OpenSearch(hosts, http_auth=(user, password), verify_certs=False, ssl_show_warn=False)
 else:
 return OpenSearch(hosts)

Create an index table.
def create(client: OpenSearch, index: str):
 # Index mapping information
 index_mapping = {
 "settings": {
 "index": {
 "vector": "true", # Enable the vector feature.
 "number_of_shards": 1, # Set the number of index shards as needed.
 "number_of_replicas": 0 # Set the number of index replicas as needed.
 }
 },
 "mappings": {
 "properties": {
 "my_vector": {
 "type": "vector",
 "dimension": 2,

```

```
 "indexing": True,
 "algorithm": "GRAPH",
 "metric": "euclidean"
 }
 # Other fields can be added if necessary.
}
}
}
res = client.indices.create(index=index, body=index_mapping)
print("create index result: ", res)

Write data.
def write(client: OpenSearch, index: str, vecs: list, bulk_size=500):
 for i in range(0, len(vecs), bulk_size):
 actions = ""
 for vec in vecs[i: i + bulk_size]:
 actions += '{"index": {"_index": "%s"}}\n' % index
 actions += '{"my_vector": %s}\n' % str(vec)
 client.bulk(body=actions, request_timeout=3600)
 client.indices.refresh(index=index, request_timeout=3600)
 print("write index success!")

Query a vector index.
def search(client: OpenSearch, index: str, query: list, size: int):
 # Query statement. Select an appropriate query method.
 query_body = {
 "size": size,
 "query": {
 "vector": {
 "my_vector": {
 "vector": query,
 "topk": size
 }
 }
 }
 }
 res = client.search(index=index, body=query_body)
 print("search index result: ", res)

Delete an index.
def delete(client: OpenSearch, index: str):
 res = client.indices.delete(index=index)
 print("delete index result: ", res)

if __name__ == '__main__':
 os_client = get_client(hosts=['http://xx.xx.xx.xx:9200'])

 # For a security-mode cluster that uses HTTPS, run the following:
 # os_client = get_client(hosts=['https://xx.xx.xx.xx:9200', 'https://xx.xx.xx.xx:9200'], user='xxxxx',
 # password='xxxxx')

 # For a security-mode cluster with HTTPS disabled, run the following:
 # os_client = get_client(hosts=['http://xx.xx.xx.xx:9200', 'http://xx.xx.xx.xx:9200'], user='xxxxx',
 # password='xxxxx')

 # Test the index name.
 index_name = "my_index"

 # Create an index.
 create(os_client, index=index_name)

 # Write data.
 data = [[1.0, 1.0], [2.0, 2.0], [3.0, 3.0]]
 write(os_client, index=index_name, vecs=data)

 # Query an index.
 query_vector = [1.0, 1.0]
 search(os_client, index=index_name, query=query_vector, size=3)
```

```
Delete an index.
delete(os_client, index=index_name)
```

### 3.8.2.10 Sample Java Code for Vector Search

OpenSearch provides standard REST APIs and clients developed using Java and Python.

This section provides a sample of Java code for creating vector indexes, and importing and querying vector data. It shows how to use the client to implement vector search.

#### Prerequisites

Add the following Maven dependency based on the actual cluster version. OpenSearch 1.3.6 is used in this example.

```
// OpenSearch 2.19.0
<dependency>
 <groupId>org.opensearch.client</groupId>
 <artifactId>opensearch-rest-high-level-client</artifactId>
 <version>2.19.0</version>
</dependency>

// OpenSearch 1.3.6
<dependency>
 <groupId>org.opensearch.client</groupId>
 <artifactId>opensearch-rest-high-level-client</artifactId>
 <version>1.3.6</version>
</dependency>
```

#### Sample Code

```
package org.example;

import org.apache.http.HttpEntity;
import org.apache.http.HttpHost;
import org.apache.http.HttpStatus;
import org.apache.http.auth.AuthScope;
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.client.CredentialsProvider;
import org.apache.http.conn.ssl.NoopHostnameVerifier;
import org.apache.http.impl.client.BasicCredentialsProvider;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.opensearch.action.admin.indices.delete.DeleteIndexRequest;
import org.opensearch.action.admin.indices.refresh.RefreshRequest;
import org.opensearch.action.bulk.BulkRequest;
import org.opensearch.action.bulk.BulkResponse;
import org.opensearch.action.index.IndexRequest;
import org.opensearch.action.search.SearchResponse;
import org.opensearch.action.support.master.AcknowledgedResponse;
import org.opensearch.client.Request;
import org.opensearch.client.RequestOptions;
import org.opensearch.client.Response;
import org.opensearch.client.RestClient;
import org.opensearch.client.RestClientBuilder;
import org.opensearch.client.RestHighLevelClient;
import org.opensearch.client.indices.CreateIndexRequest;
import org.opensearch.client.indices.CreateIndexResponse;
import org.opensearch.common.settings.Settings;
import org.opensearch.common.xcontent.DeprecationHandler;
import org.opensearch.common.xcontent.NamedXContentRegistry;
import org.opensearch.common.xcontent.XContentParser;
import org.opensearch.common.xcontent.XContentType;
```



```
import javax.net.ssl.SSLContext;
import javax.net.ssl.TrustManager;
import javax.net.ssl.X509TrustManager;
import java.io.IOException;
import java.security.KeyManagementException;
import java.security.NoSuchAlgorithmException;
import java.security.SecureRandom;
import java.security.cert.X509Certificate;
import java.util.Arrays;
import java.util.List;

public class ClientExampleOS {
 private final RestHighLevelClient client;

 public ClientExampleOS(RestHighLevelClient client) {
 this.client = client;
 }

 //Create a client for accessing non-security clusters.
 public static RestHighLevelClient getClient(List<String> hosts, int port, String scheme) {
 HttpHost[] httpHosts = hosts.stream().map(p -> new HttpHost(p, port,
scheme)).toArray(HttpHost[]::new);
 return new RestHighLevelClient(RestClient.builder(httpHosts));
 }

 //Create a client for accessing security clusters.
 public static RestHighLevelClient getClient(List<String> hosts, int port, String scheme, String user, String
password) {
 final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
 credentialsProvider.setCredentials(AuthScope.ANY, new UsernamePasswordCredentials(user,
password));
 SSLContext sc = null;
 try {
 sc = SSLContext.getInstance("SSL");
 sc.init(null, trustAllCerts, new SecureRandom());
 } catch (KeyManagementException | NoSuchAlgorithmException e) {
 e.printStackTrace();
 }
 HttpHost[] httpHosts = hosts.stream().map(p -> new HttpHost(p, port,
scheme)).toArray(HttpHost[]::new);
 final SSLIOSessionStrategy sessionStrategy = new SSLIOSessionStrategy(sc,
NoopHostnameVerifier.INSTANCE);
 RestClientBuilder builder = RestClient.builder(httpHosts).setHttpClientConfigCallback(httpClientBuilder
-> {
 httpClientBuilder.disableAuthCaching();
 httpClientBuilder.setSSLStrategy(sessionStrategy);
 return httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
 });
 return new RestHighLevelClient(builder);
 }

 public static TrustManager[] trustAllCerts = new TrustManager[] {
 new X509TrustManager() {
 @Override
 public void checkClientTrusted(X509Certificate[] chain, String authType) {
 }
 @Override
 public void checkServerTrusted(X509Certificate[] chain, String authType) {
 }
 @Override
 public X509Certificate[] getAcceptedIssuers() {
 return null;
 }
 }
 };

 # Create an index.
 public void create(String index) throws IOException {
 CreateIndexRequest request = new CreateIndexRequest(index);
```

```
request.settings(Settings.builder()
 .put("index.vector", true) // Enable the vector feature.
 .put("index.number_of_shards", 1) //Set the number of index shards as needed.
 .put("index.number_of_replicas", 0) //Set the number of index replicas as needed.
);
String mapping =
 "{" +
 " \"properties\": {" +
 " \"my_vector\": {" +
 " \"type\": \"vector\", \" + // Set this field as the vector type.
 " \"indexing\": \"true\", \" + // Enable index acceleration.
 " \"dimension\": \"2\", \" + // Vector index
 " \"metric\": \"euclidean\", \" + // Similarity metric
 " \"algorithm\": \"GRAPH\" \" + // Index algorithm
 " } \" +
 " } \" +
 "}";
request.mapping(mapping, XContentType.JSON);
CreateIndexResponse response = client.indices().create(request, RequestOptions.DEFAULT);
if (response.isAcknowledged()) {
 System.out.println("create " + index + " success");
}
}

// Write data. You are advised to keep the batch size under 500 vectors.
public void write(String index, List<float[]> vectors) throws IOException {
 BulkRequest request = new BulkRequest();
 for (float[] vec : vectors) {
 request.add(new IndexRequest(index).source("my_vector", vec));
 }

 BulkResponse response = client.bulk(request, RequestOptions.DEFAULT);
 if (response.hasFailures()) {
 System.out.println(response.buildFailureMessage());
 } else {
 System.out.println("write bulk to " + index + " success");
 }

 // Optional. Elasticsearch will refresh it by default.
 client.indices().refresh(new RefreshRequest(index), RequestOptions.DEFAULT);
}

// Query vectors.
public void search(String index, float[] query, int size) throws IOException {
 String queryFormat =
 "{\n" +
 " \"size\": %d,\n" +
 " \"query\": {\n" +
 " \"vector\": {\n" +
 " \"my_vector\": {\n" + // Query the vector field name.
 " \"vector\": %s,\n" +
 " \"topk\": %d\n" +
 " }\n" +
 " }\n" +
 " }\n" +
 "}";
 String body = String.format(queryFormat, size, Arrays.toString(query), size);
 Request request = new Request("POST", index + "/" + "_search");
 request.setJsonEntity(body);
 Response response = client.getLowLevelClient().performRequest(request);
 if (response.getStatusLine().getStatusCode() != HttpStatus.SC_OK) {
 System.out.println(response.getEntity()); //Handle the error based on service requirements.
 return;
 }
 // Process the normally returned result based on service requirements.
 HttpEntity entity = response.getEntity();
 XContentType xContentType = XContentType.fromMediaTypeOrFormat("application/json");
 XContentParser parser = xContentType.xContent().createParser(NamedXContentRegistry.EMPTY,
 DeprecationHandler.IGNORE_DEPRECATIONS, entity.getContent());
```

```
SearchResponse searchResponse = SearchResponse.fromXContent(parser);
System.out.println(searchResponse);
}

// Delete an index.
public void delete(String index) throws IOException {
 DeleteIndexRequest request = new DeleteIndexRequest(index);
 AcknowledgedResponse response = client.indices().delete(request, RequestOptions.DEFAULT);
 if (response.isAcknowledged()) {
 System.out.println("delete " + index + " success");
 }
}

public void close() throws IOException {
 client.close();
}

public static void main(String[] args) throws IOException {
 // For a non-security mode cluster, run the following:
 RestHighLevelClient client = getClient(Arrays.asList("xx.xx.xx.xx"), 9200, "http");

 /*
 * For a security-mode cluster with HTTPS enabled, run the following:
 * RestHighLevelClient client = getClient(Arrays.asList("xx.xx.xx.xx", "xx.xx.xx.xx"), 9200, "https",
 "user_name", "password");
 * For a security-mode cluster with HTTPS disabled, run the following:
 * RestHighLevelClient client = getClient(Arrays.asList("xx.xx.xx.xx", "xx.xx.xx.xx"), 9200, "http",
 "user_name", "password");
 */

 ClientExampleOS example = new ClientExampleOS(client);
 String indexName = "my_index";

 //Create an index.
 example.create(indexName);

 //Write data.
 List<float[]> data = Arrays.asList(new float[]{1.0f, 1.0f}, new float[]{2.0f, 2.0f}, new float[]{3.0f, 3.0f});
 example.write(indexName, data);

 //Query an index.
 float[] queryVector = new float[]{1.0f, 1.0f};
 example.search(indexName, queryVector, 3);

 //Delete an index.
 example.delete(indexName);


 //Close the client.
 example.close();
}
}
```

### 3.8.3 Configuring Decoupled Storage and Compute for an OpenSearch Cluster

As data grows rapidly, traditional monolithic storage architectures struggle to balance performance and cost efficiency. By decoupling storage and compute, CSS's OpenSearch clusters allow separate storage of hot data (frequently accessed data) and cold data (rarely accessed historical data). This ensures real-time query performance for hot data while reducing long-term storage costs. The solution is ideal when cold data queries can tolerate a relatively low performance standard, such as in log analytics and auditing scenarios.

## How the Feature Works

Figure 3-14 Decoupled storage and compute



The core of decoupled storage and compute is storage tiering based on data temperature (how frequent data is accessed).

- Hot data: frequently accessed data stored on high-speed SSDs, searchable in milliseconds.
- Cold data: rarely accessed data stored in OBS (invisible to users), using Standard storage.

In CSS, indexes may go through the following stages during their entire lifecycle.

- Hot index (open): Writes are allowed. Query latency is in milliseconds.
- Frozen index (freeze): Rarely accessed data is stored in OBS. The index becomes read-only, and query latency is in seconds or even minutes.
- Deleted index: Frozen indexes are deleted periodically to reclaim storage resources.

## Impact on Billing

Upon freezing, indexes are dumped to OBS, and the storage class is Standard. OBS buckets that store frozen indexes are invisible to users—users cannot view or manage them on the OBS console. Note that storing indexes in OBS buckets incurs additional fees, billed on a pay-per-use basis according to consumption of Standard storage capacity. For details, see [Cloud Search Service Price Calculator](#).

## Constraints

- Only OpenSearch 1.3.6 and 2.19.0 clusters support decoupled storage and compute.
- While an index is being frozen, the system sets the index to the read-only state. Even after the data in the index is dumped to OBS, the index remains read-only and no data can be written into the index.
- During index freezing, the data in it can still be queried. After the freezing is complete, the index is closed and then re-opened. During this transition period, the index becomes unsearchable, and the cluster may be in the **red** state temporarily. The index becomes searchable again after being re-opened.

- After an index is frozen, its data is dumped to an OBS bucket, and any index data is deleted from local disks. A frozen index has an increased query latency. During an aggregation operation, the latency becomes even longer because the query is complex and a large amount of data needs to be retrieved.
- A frozen index with data already dumped to OBS cannot be unfrozen. That is, a read-only index cannot be rolled back to writable.
- In the case of two clusters configured with read/write splitting, the secondary cluster cannot set storage-compute decoupling for indexes synchronized from the primary cluster.
- When an index is frozen, its data is dumped to an OBS bucket. This process will consume network bandwidth. Make sure data transmission between your cluster and OBS does not exhaust the maximum bandwidth and QPS supported by OBS. For details, see [OBS Use Constraints](#). If these constraints are violated, the performance of queries that involve OBS will deteriorate. For example, the speed of restoring shards and querying data will be slow.

## Logging In to OpenSearch Dashboards

Log in to OpenSearch Dashboards and go to the command execution page. OpenSearch clusters support multiple access methods. This topic uses OpenSearch Dashboards as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
4. In the left navigation pane, choose **Dev Tools**.

## Freezing an Index

Freezing an index means to move cold data from solid state drives (SSDs) to OBS, freeing up SSD storage for hot, frequently accessed data. A frozen index becomes read-only.

1. Freeze an index.

Run the following command to dump the data of a specified index to OBS and set its storage class to Standard:

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

where, **index\_name** indicates the name of the index to be frozen.

Example response:

```
{
 "freeze_uuid": "pdsRgUtSTymVDWR_HoTGFw"
}
```

where, **freeze\_uuid** indicates the ID of the index freezing task. After an index freezing request is submitted, an asynchronous task is started, and an asynchronous task ID is returned, which you can use to query the progress of the index freezing task.

2. Check the progress of the index freezing task.

Run the following command to check the index freezing progress by specifying the task ID:

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

Example response:

```
{
 "stage" : "STARTED",
 "shards_stats" : {
 "INIT" : 0,
 "FAILURE" : 0,
 "DONE" : 0,
 "STARTED" : 1,
 "ABORTED" : 0
 },
 "indices" : {
 "data1" : [
 {
 "uuid" : "7OS-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
 }
 }
]
 }
}
```

**Table 3-86** Response parameters

| Parameter    | Description                                                                                                                                                                                                                                                                                                                         |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| stage        | Status of the index freezing task.<br>Value range: <ul style="list-style-type: none"><li>• <b>INIT</b>: The task has just started or is being initialized.</li><li>• <b>FAILURE</b>: failed</li><li>• <b>DONE</b>: complete</li><li>• <b>STARTED</b>: started</li><li>• <b>ABORTED</b>: canceled. This field is reserved.</li></ul> |
| shards_stats | Number of shards in each state.                                                                                                                                                                                                                                                                                                     |

| Parameter | Description                                                                |
|-----------|----------------------------------------------------------------------------|
| indices   | Status details of each index, as described in <a href="#">Table 3-87</a> . |

**Table 3-87** indices parameters

| Parameter               | Description                                                                                                          |
|-------------------------|----------------------------------------------------------------------------------------------------------------------|
| uuid                    | ID of the index freezing task.                                                                                       |
| index                   | Index and shard information.                                                                                         |
| start_ms                | Task start time.                                                                                                     |
| end_ms                  | Task end time. If the task is still ongoing, <b>-1</b> is displayed.                                                 |
| total_time              | How long the task has been ongoing.                                                                                  |
| total_time_in_millis    | How long the task has been ongoing, in milliseconds.                                                                 |
| stage                   | Shard status.                                                                                                        |
| failure                 | Task failure cause. If no failure has occurred, <b>null</b> is returned.                                             |
| size.total_bytes        | Total bytes to be frozen.                                                                                            |
| size.finished_bytes     | Number of bytes that have been frozen.                                                                               |
| size.percent            | Freezing progress in terms of bytes that have been frozen vs total bytes to freeze.                                  |
| file.total_bytes        | Total number of documents to be frozen.                                                                              |
| file.finished_bytes     | Number of documents that have been frozen.                                                                           |
| file.percent            | Freezing progress in terms of the number of documents that have been frozen vs. total number of documents to freeze. |
| rate_limit.paused_times | Number of times freezing has been or was suspended due to rate limiting.                                             |
| rate_limit.paused_nanos | Duration of freezing task suspension due to rate limiting, in nanoseconds.                                           |

When index freezing is complete, the settings parameters will be returned, as described in [Table 3-88](#).

**Table 3-88** settings parameters

| Parameter             | Description                                                              |
|-----------------------|--------------------------------------------------------------------------|
| index.frozen_low_cost | Frozen index. The value is fixed to <b>true</b> .                        |
| index.blocks.write    | Writes disallowed to a frozen index. The value is fixed to <b>true</b> . |
| index.store.type      | Index storage type. The value is fixed to <b>obs</b> .                   |

### 3. Query the list of frozen indexes.

#### NOTE

Here, only some of the key parameters for querying the frozen index list are described. For more information about the API, see [CAT indices](#).

Run the following command to query the frozen index list based on index freezing status:

```
GET _cat/freeze_indices?stage={STAGE}
```

STAGE indicates the index freezing status. The value can be:

- start: Freezing has started and is ongoing.
- done: Freezing has completed.
- unfreeze: Freezing has not started.
- Empty or any other value: Freezing is ongoing or has completed.

Example response:

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

## Optimizing Cache Settings for Frozen Indexes

After an index is frozen, its data is dumped to OBS. To reduce direct data retrieval from OBS and thus improve query performance, some data is cached in the cluster. Data that is requested for the first time is directly retrieved from OBS. The retrieved data is then cached in the cluster memory. In response to subsequent queries, the system searches the cache first. CSS allows you to query cache statistics for frozen indexes that are stored in OBS buckets. You can also reset the cache status and modify cache configuration.

### 1. Query cache statistics for frozen indexes stored in OBS buckets.

- Run the following command to query cache statistics for frozen indexes on all nodes:

```
GET _frozen_stats
```

- Run the following command to query cache statistics for frozen indexes on specified nodes:

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

where, **node\_id** indicates a cluster node ID.

The following is an example of the returned information for querying frozen index cache statistics on all nodes:

```
{
 "_nodes" : {
```



```
"total" : 1, // Total number of nodes
"successful" : 1, // Successful nodes
"failed" : 0 // Failed nodes
},
"cluster_name" : "css-zzz1", //Cluster name
"nodes" : {
 "7uwKO38RRoaON37YsXhCYw" : {
 "name" : "css-zzz1-ess-esn-2-1", //Node name
 "transport_address" : "10.0.0.247:9300", //Node transport address
 "host" : "10.0.0.247", //Node host
 "ip" : "10.0.0.247", //Node IP address
 "block_cache" : {
 "default" : {
 "type" : "memory", //Cache type. memory indicates in-memory cache.
 "block_cache_capacity" : 8192, //Cache capacity
 "block_cache_blocksize" : 8192, //Single block size in the cache, in bytes. In the example, the
 block size is 8 KB.
 "block_cache_size" : 12, //Cache capacity used
 "block_cache_hit" : 14, //Number of cache hits
 "block_cache_miss" : 0, //Number of cache misses
 "block_cache_eviction" : 0, //Number of cache evictions
 "block_cache_store_fail" : 0 //Number of cache storage failures, which occur when the
 cache is full.
 }
 }
 },
 "obs_stats" : {
 "list" : {
 "obs_list_count" : 17, //Number of times the OBS list API was called.
 "obs_list_ms" : 265, //Total length of time spent calling the OBS list API.
 "obs_list_avg_ms" : 15 //Average time spent calling the OBS list API.
 },
 "get_meta" : {
 "obs_get_meta_count" : 79, //Number of times the OBS get metadata API was called.
 "obs_get_meta_ms" : 183, //Total length of time spent calling the OBS get metadata API.
 "obs_get_meta_avg_ms" : 2 //Average time spent calling the OBS get metadata API.
 },
 "get_obj" : {
 "obs_get_obj_count" : 12, //Number of times the OBS get object API was called.
 "obs_get_obj_ms" : 123, //Total length of time spent calling the OBS get object API.
 "obs_get_obj_avg_ms" : 10 //Average time spent calling the OBS get object API.
 },
 "put_obj" : {
 "obs_put_obj_count" : 12, //Number of times the OBS put object API was called.
 "obs_put_obj_ms" : 2451, //Total length of time spent calling the OBS put object API.
 "obs_put_obj_avg_ms" : 204 //Average time spent calling the OBS put object API.
 },
 "obs_op_total" : {
 "obs_op_total_ms" : 3022, //Total length of time spent calling OBS APIs.
 "obs_op_total_count" : 120, //Total number of times calling OBS APIs.
 "obs_op_avg_ms" : 25 //Average time spent calling OBS APIs.
 }
 },
 "reader_cache" : {
 "hit_count" : 0,
 "miss_count" : 1,
 "load_success_count" : 1,
 "load_exception_count" : 0,
 "total_load_time" : 291194714,
 "eviction_count" : 0
 }
}
}
```

2. Reset the cache status for frozen indexes stored in OBS buckets.

Run the following command to reset the cache status for frozen indexes:

```
POST _frozen_stats/reset
```

This command is used to debug performance issues. If you reset the cache status and then run the cache query command, you can check the accurate cache command status. It is not advisable to use this command during runtime.

Example response:

```
{
 "_nodes" : {
 "total" : 1,
 "successful" : 1,
 "failed" : 0
 },
 "cluster_name" : "Es-0325-007_01",
 "nodes" : {
 "mqTdk2YRSPyOSXfesREFSg" : {
 "result" : "ok"
 }
 }
}
```

3. Modify cache configuration for frozen indexes.

Elasticsearch accesses different types of files using different methods. The cache system supports multi-level caches and uses blocks of different sizes to cache different types of files. For example, typically, a large number of small blocks are used to cache .fdx and .tip files, while a small number of large blocks are used to cache .fdt files. You can modify the cache configuration based on service needs. [Table 3-89](#) describes the parameters that can be configured.

**Table 3-89** Cache configuration parameters

| Parameter                           | Type  | Description                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------|-------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| low_cost.obs.blockcache.names       | Array | A list of names of multi-level caches. The cache system supports multi-level caches for data of different access granularities. There is a default cache named <b>default</b> . If you configure this parameter, the value must include at least the <b>default</b> cache in addition to other custom cache names. Default value: <b>default</b>                                                                            |
| low_cost.obs.blockcache.<NAME>.type | ENUM  | Storage type of the cache named <NAME>. Value range: <ul style="list-style-type: none"><li>• <b>memory</b>: The cache uses in-memory storage, which means it consumes memory resources.</li><li>• <b>file</b>: The cache uses disk storage. Its capacity is limited by available disk space. In this case, you are advised to use ultra-high I/O disks to improve cache performance.</li></ul> Default value: <b>memory</b> |

| Parameter                                            | Type    | Description                                                                                                                                                                                                                                                        |
|------------------------------------------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| low_cost.obs.blockcache.<NAME>.blockshift            | Integer | Block size in the cache named <NAME> (in bytes).<br>Formula: $2^{\text{blockshift}}$ . For example, if the value is 16, the block size is $2^{16}$ bytes, that is, 65536 bytes or 64 KB.<br>Value: number of left-shifts<br>Default value: 13 (equivalent to 8 KB) |
| low_cost.obs.blockcache.<NAME>.bank.count            | Integer | Number of partitions in the cache named <NAME>. Each partition is an independent unit in the cache.<br>Default value: <b>1</b>                                                                                                                                     |
| low_cost.obs.blockcache.<NAME>.number.blocks.perbank | Integer | Number of blocks in each partition of the cache named <NAME>.<br>Default value: <b>8192</b>                                                                                                                                                                        |
| low_cost.obs.blockcache.<NAME>.exclude.file.types    | Array   | List of file name extensions for files that cannot be cached in the cache named <NAME>.<br>If a file's extension is not in <b>exclude.file.types</b> or <NAME>.file.types (if configured), the default cache policy applies.                                       |
| low_cost.obs.blockcache.<NAME>.file.types            | Array   | List of file name extensions for files that can be cached in the cache named <NAME>.<br>If a file's extension is not in <b>exclude.file.types</b> or <NAME>.file.types (if configured), the default cache policy applies.                                          |
| index.frozen.obs.max_bytes_per_sec                   | String  | Maximum rate at which the documents of frozen indexes are uploaded to OBS.<br>Format: <value><unit>. The unit can be B (byte), KB, MB, or GB.<br>The change takes effect immediately.<br>Default value: <b>150 MB</b>                                              |

| Parameter                                         | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| low_cost.obs.index.upload.threshold.use.multipart | String | <p>Multipart upload threshold. In the process of dumping frozen indexes to OBS, if the size of a single document exceeds this threshold, a multipart upload is performed.</p> <p>Format: &lt;value&gt;&lt;unit&gt;. The unit can be B (byte), KB, MB, or GB.</p> <p>Multipart upload is supported by OBS.</p> <p>0 or not configured: Multipart upload may be disabled.</p> <p>Default value: 1 GB</p> |

The following is a typical cache configuration example. A two-level cache system is used: **default** and **large**. The **default** cache has a total of 30 x 4096 64-KB blocks. It is used to cache non-.fdt files. The **large** cache has 5 x 1000 2-MB 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"]
```

## Enhancing Query Performance for Frozen Indexes



### WARNING

Only OpenSearch 1.3.6 and 2.19.0 clusters support query performance enhancement for frozen indexes.

When frozen indexes are queried on the **Discover** page of OpenSearch Dashboards for the first time, all data needs to be retrieved from OBS because there is no data in the cache. If a large number of documents need to be returned, it takes a long time to retrieve the corresponding time fields and document metadata from OBS. By caching this part of data in the cluster, query performance can be improved significantly. This is how CSS improves query performance for frozen indexes. Local cache settings are preset. You can review them and modify them as needed.

1. Modify local cache settings for frozen indexes.

**Table 3-90** Local cache settings

| Configuration Item                    | Type    | Scope | Can Be Changed Dynamically | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------------|---------|-------|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| low_cost.local_cache.max.capacity     | Integer | node  | Yes                        | <p>Maximum number of cold data caches that can be made available on each node. (Each shard corresponds to a cache.)</p> <p>Value range: 10–5000</p> <p>Default value: <b>500</b></p> <ul style="list-style-type: none"><li>• If the heap memory usage remains high, decrease this value.</li><li>• If the value of <b>load_overflow_count</b> keeps increasing rapidly, increase this value.</li></ul>                                                                                                     |
| index.low_cost.local_cache.threshold  | Integer | index | Yes                        | <p>Threshold for enabling local caching of cold data.</p> <ul style="list-style-type: none"><li>• If in an index, the percentage of <b>date</b> fields is lower than this value, cold data of the <b>date</b> type will be cached locally. Otherwise, the local cache will not be used.</li><li>• If date fields account for the vast majority of the total data volume of the current index, you should not use this setting.</li></ul> <p>Value range: 0–100</p> <p>Default value: 50</p> <p>Unit: %</p> |
| index.low_cost.local_cache.evict_time | String  | index | Yes                        | <p>Retention duration for cold data in the local cache. The value is determined by <b>index.frozen_date</b> (time of index freezing). If <b>index.frozen_date</b> is unavailable, the value is determined by the index creation time.</p> <p>You are advised to adjust the retention duration based on your disk usage.</p> <p>Value range: 1 to 365 days</p> <p>Default value: <b>30d</b></p> <p>Unit: days</p>                                                                                           |

The following provides some examples, where **index\_name** indicates the modified index.

- 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 {index_name}/_settings
{
 "index.low_cost.local_cache.threshold": 20
}
```

- Run the following command to modify

**index.low\_cost.local\_cache.evict\_time:**

```
PUT {index_name}/_settings
{
 "index.low_cost.local_cache.evict_time": "7d"
}
```

## 2. Query local cache statistics for frozen indexes.

- Query local cache statistics for frozen indexes on all nodes:

```
GET /_frozen_stats/local_cache
```

- Query local cache statistics for frozen indexes on specified nodes:

```
GET /_frozen_stats/local_cache/{node_id}
```

where, **node\_id** indicates a cluster node ID.

Example response:

```
{
 "_nodes" : {
 "total" : 1,
 "successful" : 1,
 "failed" : 0
 },
 "cluster_name" : "test",
 "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.
 "get_hit_count" : 562, //Total number of hits in the local cold data cache.
 "get_miss_count" : 0, //Total number of local cold data cache misses.
 "get_total_ns" : 43849200, //Total duration for retrieving data from the local cold data
cache.
 "get_avg_ns" : 78023 //Average duration for retrieving data from the local cold
data cache.
 },
 "load_stats" : {
 "load_count" : 2, //Number of times cold data was loaded from the local cache
 "load_total_ms" : 29, //Total duration for loading cold data from the local cache
 "load_avg_ms" : 14, //Average duration for loading cold data from the local cache
 "load_fail_count" : 0, //Number of failures for loading cold data from the local
cache
 }
 }
 }
 }
}
```

```
 "load_overflow_count" : 0 //Number of times the local cold data cache exceeds the
 cache pool size.
 },
 "reload_stats" : {
 "reload_count" : 0, //Number of times the local cold data cache was regenerated.
 "reload_total_ms" : 0, //Total duration for regenerating the local cold data cache.
 "reload_avg_ms" : 0, //Average duration for regenerating the local cold data cache.
 "reload_fail_count" : 0 //Number of failures in regenerating the local cold data
 cache.
 },
 "init_stats" : {
 "init_count" : 0, //Number of times the local cold data cache was initialized.
 "init_total_ms" : 0, //Total duration for initializing the local cold data cache.
 "init_avg_ms" : 0, //Average duration for initializing the local cold data cache.
 "init_fail_count" : 0 //Number of failures in initializing the local cold data cache.
 }
}
}
```

## Querying the Real-Time Rates of OBS in Handling Cold Data

### WARNING

Only OpenSearch 1.3.6 and 2.19.0 clusters support this feature.

To help you understand how the storage-compute decoupling plug-in is working with OBS, an API for collecting statistics on the real-time rates of OBS has been added, and the real-time rates are recorded in the index **.freeze\_obs\_rate-YYYY.mm.dd**.

Calculation method: The average OBS operation rates in the last 5 seconds are calculated every 5 seconds.

The system index **.freeze\_obs\_rate-YYYY.mm.dd** records statistics on OBS real-time operation rates, helping you understand relevant trends about the OBS resources that store cold data. The default retention period of the index is 30 days.

#### 1. Querying the real-time rates of OBS in handling cold data.

- Run the following command to query the real-time OBS rates on all nodes:

```
GET _frozen_stats/obs_rate
```

- Run the following command to query the real-time OBS rates on specified nodes:

```
GET _frozen_stats/obs_rate/{node_id}
```

where, **node\_id** indicates a cluster node ID.

Example response:

```
{
 "_nodes" : {
 "total" : 1,
 "successful" : 1,
 "failed" : 0
 },
 "cluster_name" : "test",
 "nodes" : {
 "dfIDvcSwTJ-fkillT2zE3A" : {
 "name" : "node-1",
 "transport_address" : "127.0.0.1:9300",
```

```
"host" : "127.0.0.1",
"ip" : "127.0.0.1",
"update_time" : 1671777600482, // Time when the current statistics were updated.
"obs_rate" : {
 "list_op_rate" : 0.0, // Rate of OBS list operations. Unit: times/s.
 "get_meta_op_rate" : 0.0, // Rate of OBS get meta operations. Unit: times/s.
 "get_obj_op_rate" : 0.0, // Rate of OBS get operations. Unit: times/s.
 "put_op_rate" : 0.0, // Rate of OBS put operations. Unit: times/s.
 "obs_total_op_rate" : 0.0, // Rate of all OBS operations. Unit: times/s.
 "obs_upload_rate" : "0.0 MB/s", // OBS data upload rate. Unit: MB/s.
 "obs_download_rate" : "0.0 MB/s" // OBS data download rate. Unit: MB/s.
}
}
```

2. Modify the retention period of the **.freeze\_obs\_rate-YYYY.mm.dd** index that stores the OBS real-time rates. The default retention period of this index is 30 days.

Run the following command to change the index retention period to seven days:

```
PUT _cluster/settings
{
 "persistent": {
 "low_cost.obs_rate_index.evict_time": "7d"
 }
}
```

**Table 3-91** Configuration items

| Configurati<br>ion Item                    | Typ<br>e   | scop<br>e | Can Be<br>Change<br>d<br>Dynam<br>ically | Description                                                                                                                                                                                                 |
|--------------------------------------------|------------|-----------|------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| low_cost.ob<br>s_rate_inde<br>x.evict_time | Strin<br>g | nod<br>e  | Yes                                      | The retention period of the <b>.freeze_obs_rate-YYYY.mm.dd</b> index. <ul style="list-style-type: none"><li>• Value range: 1 to 365 days</li><li>• Default value: <b>30d</b></li><li>• Unit: days</li></ul> |

## Related Documents

- To ensure acceptable performance for cold data queries (for example, when historical data is queried periodically), use a tiered storage solution. For details, see [Switching Between Hot and Cold Storage for an OpenSearch Cluster](#).
- Decoupled storage and compute can be automated via index lifecycle management policies. For details, see [Decoupling Index Storage and Compute in an OpenSearch Cluster Through Index Lifecycle Management](#).



## 3.8.4 Configuring Flow Control for an OpenSearch Cluster

Configure flow control policies for your OpenSearch cluster in both the inbound and outbound directions, ensuring cluster stability by safeguarding against abnormal traffic.

An OpenSearch cluster can become overloaded due to traffic surges, malicious requests, and internal resource competition, which can even lead to node failures. Through policies like client request throttling, write backpressure, and traffic pattern analysis, flow control ensures proper resource allocation, thereby protecting clusters from overload. It covers the following scenarios:

- High-concurrency write handling: mitigates the risk of out-of-memory (OOM) exceptions under heavy write loads.
- Security defense: controls access by IP address using both blacklists and whitelists.
- Emergency response: blocks malicious or abnormal traffic in one click.
- Performance optimization: optimizes flow control thresholds and policies based on collected statistics.

### How the Feature Works

**Table 3-92** Flow control policies

| Policy                  | What It Does                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Details                                             |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|
| HTTP/HTTPS flow control | <p>Controls client access traffic using blacklists and whitelists, an upper limit on concurrent connections, and a rate limit on new connection attempts.</p> <ul style="list-style-type: none"><li>• Blacklists and whitelists: A whitelist takes precedence over a blacklist. If an IP address is on both lists, it is allowed. Requests sent through blacklisted connections will be ignored.</li><li>• Concurrent connection limit: Limits the total number of HTTP connections per second to prevent overload.</li><li>• New connection limit: Limits the number of new connections that can be set up per second. The warmup_period parameter protects against connection floods, allowing traffic to grow gradually and steadily.</li></ul> | <a href="#">Configuring HTTP/HTTPS Flow Control</a> |

| Policy                                   | What It Does                                                                                                                                                                                                                                                                                                                    | Details                                                                                                                                                    |
|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Memory flow control                      | When the heap memory usage exceeds a pre-defined threshold (for example, 80%), the system stops receiving large requests, and garbage collection (GC) is triggered to reclaim memory.<br><br>Write traffic is throttled by setting the backpressure factor (in_flight_factor) and the maximum delay for request handling (max). | <a href="#">Configuring Memory-based Flow Control</a>                                                                                                      |
| One-click traffic blocking               | When enabled, the system immediately disconnects all client connections, but not those used for OpenSearch Dashboards access or O&M and monitoring APIs, in order to restore the cluster.                                                                                                                                       | <a href="#">Configuring One-Click Traffic Blocking</a>                                                                                                     |
| Request statistics sampling and analysis | Records request metrics (such as bulk writes and queries) by client IP address, and exposes them via a statistics API to evaluate the cluster load and proactively identify abnormal traffic patterns.                                                                                                                          | <a href="#">Configuring Request Statistics Sampling and Analysis</a>                                                                                       |
| Access logging                           | Records the URLs and bodies of HTTP/HTTPS requests for cluster load and client request analysis.<br><br>Access logs can also be saved to files (that is, persisted to disk) to facilitate troubleshooting and performance analysis.                                                                                             | <ul style="list-style-type: none"><li>• <a href="#">Configuring Access Logging</a></li><li>• <a href="#">Configuring Access Logging in Files</a></li></ul> |

## Constraints

- Only OpenSearch 2.19.0 supports flow control.
- Enabling flow control may compromise cluster and node performance. For example, memory-based flow control may cause query failures on OpenSearch Dashboards.
- Configuring too many paths for memory-based flow control has a significant impact on performance.
- Enabling access logging will hurt cluster performance.
- After flow control policies are enabled, excess or non-compliant requests are rejected directly.

## Logging In to OpenSearch Dashboards

Log in to OpenSearch Dashboards and go to the command execution page. OpenSearch clusters support multiple access methods. This topic uses OpenSearch Dashboards as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
4. In the left navigation pane, choose **Dev Tools**.

## Configuring HTTP/HTTPS Flow Control

Control client access traffic using blacklists and whitelists, an upper limit on concurrent connections, and a rate limit on new connection attempts.

1. Enable HTTP/HTTPS flow control.

```
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
 }
}
```

**Table 3-93** Parameters for configuring HTTP/HTTPS flow control

| Parameter                | Type         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.http.enabled | Boolean      | Whether to enable HTTP/HTTPS flow control. Enabling it may impact node performance. <ul style="list-style-type: none"><li>• true: Enable HTTP/HTTPS flow control.</li><li>• false (default): Disable HTTP/HTTPS flow control.</li></ul>                                                                                                                                                                                                                                                                                             |
| flowcontrol.http.allow   | List<String> | A whitelist of client IP addresses or CIDR blocks that are allowed to access the cluster, supporting: <ul style="list-style-type: none"><li>• Individual IP addresses, for example, <b>192.18.0.1</b>.</li><li>• CIDR blocks, for example, <b>192.168.0.0/24</b>.</li><li>• Multiple IP addresses or CIDR blocks separated by commas (,), for example, <b>192.168.0.1/24, 192.168.2.1/24</b>.</li></ul> The default is an empty list, indicating no whitelist.<br>Setting this parameter to <b>null</b> restores the default value. |

| Parameter                   | Type         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.http.deny       | List<String> | <p>A blacklist of client IP addresses or CIDR blocks that are not allowed to access the cluster. A whitelist takes precedence over a blacklist. A blacklist supports the following:</p> <ul style="list-style-type: none"><li>• Individual IP addresses, for example, <b>192.18.0.1</b>.</li><li>• CIDR blocks, for example, <b>192.168.0.0/24</b>.</li><li>• Multiple IP addresses or CIDR blocks separated by commas (,), for example, <b>192.168.0.1/24, 192.168.2.1/24</b>.</li></ul> <p>The default is an empty list, indicating no blacklist.</p> <p>Setting this parameter to <b>null</b> restores the default value.</p> |
| flowcontrol.http.concurrent | Integer      | <p>Maximum number of concurrent HTTP/HTTPS connections that can be handled by each node per second. The default value is the number of available node vCPUs multiplied by 600.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| flowcontrol.http.newconnect | Integer      | <p>Maximum number of new HTTP/HTTPS connections that can be created per second per node.</p> <p>The default value is the number of available node vCPUs multiplied by 200.</p> <p>Setting this parameter to <b>null</b> restores the default value.</p>                                                                                                                                                                                                                                                                                                                                                                          |

| Parameter                      | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.http.warmup_period | Integer | <p>A grace period during which a system gradually ramp up from accepting zero HTTP/HTTPS requests to its full, maximum capacity.</p> <p>Value range: 0–10000</p> <p>Unit: ms</p> <p>The default value is 0, indicating that there is no such warm-up period, and the maximum rate can be reached instantly.</p> <p>Setting this parameter to <b>null</b> restores the default value.</p> <p>For example, if <b>flowcontrol.http.newconnect</b> is set to <b>100</b> and <b>flowcontrol.http.warmup_period</b> is set to <b>5000ms</b>, it takes 5 seconds for the system to reach 100 new connections per second.</p> |

## 2. Disable HTTP/HTTPS flow control.

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.http.enabled": false
 }
}
```

## Configuring Memory-based Flow Control

Enable write throttling to mitigate the risk of OOM exceptions when the heap memory usage of a node exceeds a predefined threshold.

### 1. Enable memory-based flow control.

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.memory.enabled": true,
 "flowcontrol.memory.heap_limit": "80%"
 }
}
```

**Table 3-94** Memory-based flow control parameters

| Parameter                     | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.memory.enabled    | Boolean | <p>Whether to enable memory-based flow control. When enabled, node heap memory usage is monitored and a threshold is set, and writes are throttled when this threshold is reached.</p> <ul style="list-style-type: none"><li>• true: Enable memory-based flow control.</li><li>• false (default): Disable memory-based flow control.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| flowcontrol.memory.heap_limit | String  | <p>Node heap memory usage threshold. When this threshold is exceeded, a write backpressure mechanism is triggered.</p> <p>Value range: 10%–100%</p> <p>The default value 90% of <b>flowcontrol.memory.heap_limit</b> is a conservative threshold. When the heap memory usage exceeds 90%, the system stops accepting client requests that exceed 64 KB, until heap memory usage decreases. Once the heap memory usage decreases to 85%, client data equivalent to <b>5% x maximum heap memory capacity</b> can be read. If the heap memory usage stays above 90% for a long time, client requests cannot be processed. In this case, garbage collection is triggered until the heap memory usage drops below this threshold.</p> <p>Generally, you can set this threshold to 80% or less to ensure that cluster nodes have reserved some heap memory for operations besides data writing, such as segment merges.</p> <p>Setting this parameter to <b>null</b> restores the default value.</p> |

| Parameter                            | Type      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.holding.in_flight_factor | Float     | <p>Backpressure factor, which controls the sensitivity of memory-based backpressure. A larger value indicates more powerful write throttling.</p> <p>Value range: <math>\geq 0.5</math></p> <p>The default value is <b>1.0</b>.</p> <p>Setting this parameter to <b>null</b> restores the default value.</p>                                                                                                                                                                                                                                                                                                                          |
| flowcontrol.holding.max              | TimeValue | <p>Maximum request handling delay allowed before requests are handled according to the policy defined by <b>flowcontrol.holding.max_strategy</b>.</p> <p>Value range: <math>\geq 15s</math></p> <p>Unit: second</p> <p>Default value: 60s</p> <p>Setting this parameter to <b>null</b> restores the default value.</p>                                                                                                                                                                                                                                                                                                                |
| flowcontrol.holding.max_strategy     | String    | <p>Handling policy for requests delayed longer than <b>flowcontrol.holding.max</b>.</p> <ul style="list-style-type: none"><li>• <b>keep</b> (default): Keep the backpressure status while waiting for the memory usage to drop. The server determines when to execute the requests based on real-time memory usage.</li><li>• <b>soft</b>: Forcibly execute the requests, but the <b>inFlight</b> circuit breaker gets to decide whether to reject them.</li><li>• <b>hard</b>: Reject the requests immediately and disconnect client connections.</li></ul> <p>Setting this parameter to <b>null</b> restores the default value.</p> |

| Parameter                        | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.memory.once_free_max | String  | Maximum memory that can be made available at a time for a re-enabled request queue. This parameter can be configured to prevent cluster overload caused by a flood of incoming requests.<br><br>Value range: 1%–50%<br><br>The default value is 5%.<br><br>Setting this parameter to <b>null</b> restores the default value.                                                                                                                                               |
| flowcontrol.memory.nudges_gc     | Boolean | Whether to trigger garbage collection (GC) to reclaim memory when the write pressure is too high. (The backpressure connection pool is checked every second. The write pressure is considered high if all existing connections are blocked and new write requests cannot be accepted.) <ul style="list-style-type: none"><li>• true (default): Trigger GC.</li><li>• false: Not to trigger GC.</li></ul> Setting this parameter to <b>null</b> restores the default value. |

## 2. Disable memory-based flow control.

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.memory.enabled": false
 }
}
```

## Configuring One-Click Traffic Blocking

In case of emergencies, the system immediately disconnects all client connections (excluding those used for OpenSearch Dashboards access or O&M and monitoring APIs) to restore clusters.

## 1. Enable one-click traffic blocking.

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.break.enabled": true
 }
}
```



**Table 3-95** Parameters for configuring one-click traffic blocking

| Parameter                 | Type    | Description                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.break.enabled | Boolean | Whether to enable one-click traffic blocking (similar to a circuit breaker). When enabled, the system immediately disconnects all client connections, but not those used for OpenSearch Dashboards access or O&M and monitoring APIs. <ul style="list-style-type: none"><li>true: Enable one-click disconnection.</li><li>false (default): Disable one-click disconnection.</li></ul> |

2. Disable one-click disconnection:

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.break.enabled": false
 }
}
```

## Configuring Request Statistics Sampling and Analysis

Collect request metrics by client IP address to help identify abnormal traffic patterns.

1. Enable request statistics sampling.

```
PUT _cluster/settings
{
 "transient": {
 "flowcontrol.log.access.enabled": true
 }
}
```

**Table 3-96** Parameters for request statistics sampling

| Parameter                      | Type    | Description                                                                                                                                                                                                                                                                                                                |
|--------------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.log.access.enabled | Boolean | Whether to enable request statistics sampling, that is, whether to collect request metrics (such as bulk writes and search/msearch requests) by client IP address. <ul style="list-style-type: none"><li>true: Enable request statistics sampling.</li><li>false (default): Disable request statistics sampling.</li></ul> |

| Parameter                    | Type    | Description                                                                                                                                                           |
|------------------------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.log.access.count | Integer | Maximum number of client IP addresses sampled.<br>Value range: 0–100<br>Default value: <b>10</b><br>Setting this parameter to <b>null</b> restores the default value. |

## 2. Check the sampled statistics.

- Check the flow control status of all nodes.  
GET /\_nodes/stats/filter/v2
- Check the flow control details of all nodes.  
GET /\_nodes/stats/filter/v2?detail
- Check the flow control status of a specified node.  
GET /\_nodes/{nodeId}/stats/filter/v2

**{nodeId}** indicates the node ID.

Example response:

```
{
 "_nodes" : {
 "total" : 1,
 "successful" : 1,
 "failed" : 0
 },
 "cluster_name" : "css-xxxx",
 "nodes" : {
 "d3qnVlpPTtSoadkV0LQEKA" : {
 "name" : "css-xxxx-ess-esn-1-1",
 "host" : "192.168.x.x",
 "timestamp" : 1672236425112,
 "flow_control" : {
 "http" : {
 "current_connect" : 52,
 "rejected_concurrent" : 0,
 "rejected_rate" : 0,
 "rejected_black" : 0,
 "rejected_breaker" : 0
 }
 },
 "access_items" : [
 {
 "remote_address" : "10.0.0.x",
 "search_count" : 0,
 "bulk_count" : 0,
 "other_count" : 4
 }
],
 "holding_requests" : 0
 }
 }
}
```

**Table 3-97** Response parameters

| Parameter           | Description                                                                                                                                                                                                                                                  |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| current_connect     | Number of HTTP connections to a node, which is recorded regardless of whether flow control is enabled. This value is equivalent to the <b>current_open</b> value of <b>GET /_nodes/stats/http</b> API. It shows the current client connections of each node. |
| rejected_concurrent | Number of concurrent connections rejected during flow control.<br>This metric is available only when <b>flowcontrol.http.enabled</b> is set to <b>true</b> . The count will not be cleared when flow control is disabled.                                    |
| rejected_rate       | Number of new connections rejected during flow control.<br>This metric is available only when <b>flowcontrol.http.enabled</b> is set to <b>true</b> . The count will not be cleared when flow control is disabled.                                           |
| rejected_black      | Number of new connections rejected by a preconfigured blacklist during flow control.<br>This metric is available only when <b>flowcontrol.http.enabled</b> is set to <b>true</b> . The count will not be cleared when flow control is disabled.              |
| rejected_breaker    | Number of new connections rejected during one-click traffic blocking.<br>This metric is available only when <b>flowcontrol.break.enabled</b> is set to <b>true</b> . The count will not be cleared when one-click traffic blocking is disabled.              |
| access_items        | IP addresses of clients that recently accessed the cluster.<br>The number of client IP addresses sampled is determined by <b>flowcontrol.log.access.count</b> .                                                                                              |
| remote_address      | IP addresses and the number of requests.                                                                                                                                                                                                                     |
| search_count        | Number of times a client accessed a database using <b>_search</b> and <b>_msearch</b> .                                                                                                                                                                      |
| bulk_count          | Number of times a client accessed a database using <b>_bulk</b> .                                                                                                                                                                                            |
| other_count         | Number of times a client accessed a database using other request methods.                                                                                                                                                                                    |
| holding_requests    | Number of connections to the current node where writes are halted due to flow control.                                                                                                                                                                       |

## 3. Disable request statistics sampling.

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.log.access.enabled": false
 }
}
```

## Configuring Access Logging

When access logging is enabled, the system records the URLs and bodies of HTTP/HTTPS requests for cluster load and request analysis.

## 1. Enable access logging.

- Enable access logging for all nodes in a cluster.

```
PUT /_access_log?duration_limit=30s&capacity_limit=1mb
```

- Enable access logging for a specified node in a cluster.

```
PUT /_access_log/{nodeId}?duration_limit=30s&capacity_limit=1mb
```

**{nodeId}** indicates the node ID.

**Table 3-98** Parameters for enabling access logging

| Parameter      | Type   | Description                                                                                                                                                                                                                                                                                                                                                           |
|----------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| duration_limit | String | Maximum duration of access log records. When this duration is reached, access logging stops.<br>Value range: 10 to 120<br>Unit: s<br>The default value is <b>30</b> .<br>Setting this parameter to <b>null</b> restores the default value.<br>Access logging stops when either <b>duration_limit</b> or <b>capacity_limit</b> is reached.                             |
| capacity_limit | String | Maximum memory capacity for recording access logs. When the size of an access log reaches this limit, access logging stops.<br>Value range: 1 to 5<br>Unit: MB<br>The default value is <b>1</b> .<br>Setting this parameter to <b>null</b> restores the default value.<br>Access logging stops when either <b>duration_limit</b> or <b>capacity_limit</b> is reached. |

## 2. Check access logs.

- Check the access logs of all nodes in a cluster.  
GET /\_access\_log
- Check the access logs of a specified cluster node.  
GET /\_access\_log/{nodeId}

**{nodeId}** indicates the node ID.

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 3-99** Response parameters

| Parameter      | Description                                                                      |
|----------------|----------------------------------------------------------------------------------|
| name           | Node name                                                                        |
| host           | Node IP address                                                                  |
| count          | Number of node access requests in a statistical period                           |
| access         | Details about node access requests in a statistical period                       |
| time           | Request time                                                                     |
| remote_address | Source IP address and port number in the request                                 |
| url            | Original URL of the request                                                      |
| method         | Request method                                                                   |
| content        | Request content. If the value is an empty string (""), there is no request body. |

3. Delete access logs. Logs are recorded in the memory. After reading the logs, you are advised to delete them to reclaim resources.

Delete access logs for all nodes:

```
DELETE /_access_log
```

## Configuring Access Logging in Files

Access logs can be persisted to disk for troubleshooting and analysis. Use this function sparingly, as it can impact cluster performance. Remember to disable it immediately after resolving the issue.

1. Enable access logging in files.

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.log.file.enabled": true
 }
}
```

**Table 3-100** Enable access logging in files

| Parameter                    | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| flowcontrol.log.file.enabled | Boolean | <p>Whether to record access logs in files. When enabled, the log of each access request is recorded in files.</p> <p>The log file name is <b>Cluster name_access_log.log</b>. You can check this file only through the log backup function.</p> <ul style="list-style-type: none"><li>• true: Record access logs in files.</li><li>• false (default): Not to record access logs in files.</li></ul> |

2. Disable access logging in files.

```
PUT /_cluster/settings
{
 "persistent": {
 "flowcontrol.log.file.enabled": false
 }
}
```

## 3.8.5 Configuring Large Query Isolation for an OpenSearch Cluster

Large query isolation can be configured to manage queries that have high memory usage or take too long to complete. This helps improve the stability of OpenSearch clusters and prevent out-of-memory (OOM) exceptions.

As business grows, your OpenSearch clusters may face mounting query pressure. Some complex queries may occupy excessive node memory, triggering frequent garbage collection or even OOM exceptions, which may compromise cluster performance and stability. Large query isolation enables effective management of memory-intensive, time-consuming query requests, ensuring cluster stability. Large query isolation includes the following:

- Isolating large queries: Manages memory-intensive/time-consuming queries separately to avoid impacting other queries.
- Query cancelation based on a heap memory usage threshold: Cancels a large query in the isolation pool when the node heap memory usage reaches a predefined threshold.
- Global query timeout: Automatically cancels tasks when they last longer than a predefined timeout. This timeout applies globally.

## How the Feature Works

- Isolation pool management: Large queries are placed in an isolation pool, where they are managed separately from other normal queries. Tasks in the isolation pool may be canceled based on preset memory or duration thresholds.
- Query cancelation policies:
  - fair: Selects which query to cancel by considering both memory usage and query duration.
  - mem-first: Cancels the query that has the highest memory usage.
  - time-first: Cancels the query that has lasted the longest.
- Native cancel API: OpenSearch's native cancel API can be used to cancel tasks, ensuring compatibility.

## Constraints

- Only OpenSearch 2.19.0 supports large query isolation.
- Large query isolation is enabled by default, whereas global query timeout is disabled by default. If you enable them, the configuration will take effect immediately.

## Logging In to OpenSearch Dashboards

Log in to OpenSearch Dashboards and go to the command execution page. OpenSearch clusters support multiple access methods. This topic uses OpenSearch Dashboards as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
4. In the left navigation pane, choose **Dev Tools**.

## Configuring Large Query Isolation

Large query isolation places large queries in an isolation pool, where they may be canceled based on preset memory or duration thresholds. Large query isolation is enabled by default. Any change takes effect immediately.

1. Run the following command to enable large query isolation:

```
PUT _cluster/settings
{
 "persistent": {
 "search.isolator.enabled": true
 }
}
```

```
}
}
```

**Table 3-101** Enabling large query isolation

| Parameter               | Type    | Description                                                                                                                                                                                                                                                                          |
|-------------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.enabled | Boolean | Whether to enable large query isolation. When enabled, large queries are managed separately from other normal queries. <ul style="list-style-type: none"><li>• <b>true</b> (default): Enable large query isolation.</li><li>• <b>false</b>: Disable large query isolation.</li></ul> |

2. Run the following commands to configure thresholds for defining a large query task:

```
PUT _cluster/settings
{
 "persistent": {
 "search.isolator.memory.task.limit": "50MB",
 "search.isolator.time.management": "10s"
 }
}
```

**Table 3-102** Parameters for configuring large query isolation thresholds

| Parameter                         | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.memory.task.limit | String | <p>Large query memory threshold: When a query requests more memory than specified by this threshold, it is placed into an isolation pool.</p> <p>Value range: <b>0b</b> to the maximum heap memory of a node.</p> <p>Unit: MB, GB, or other units supported by OpenSearch.</p> <p>The default value is <b>50 MB</b>.</p> <p><b>NOTE</b></p> <p>Run the following command to query the node heap memory usage and upper limit of a cluster:</p> <pre>GET _cat/nodes?<br/>&amp;h=id,ip,port,r,ramPercent,ramCurrent,heapMax,heapCurrent</pre> |



| Parameter                       | Type   | Description                                                                                                                                                                                                                                                                     |
|---------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.time.management | String | Large query duration threshold:<br>When a query has lasted longer than specified by this threshold, it is placed into an isolation pool.<br>Value range: $\geq 0\text{ms}$<br>Unit: s (second), ms (millisecond), m (minute), or h (hour).<br>The default value is <b>10s</b> . |

3. Configure the resource usage thresholds for triggering query cancelation.

```
PUT _cluster/settings
{
 "persistent": {
 "search.isolator.memory.pool.limit": "50%",
 "search.isolator.memory.heap.limit": "90%",
 "search.isolator.count.limit": 1000
 }
}
```

**Table 3-103** Parameters for configuring query cancelation thresholds

| Parameter                         | Type   | Description                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.memory.pool.limit | String | Isolation pool memory usage threshold, which indicates the maximum node heap memory usage allowed. When the total memory requested by all large queries in the isolation pool exceeds this threshold, one of the large queries is automatically canceled based on a predefined policy.<br>Value range: 0.0–100.0%<br>The default value is <b>50</b> . |
| search.isolator.memory.heap.limit | String | Heap memory usage threshold, which indicates the actual node heap memory usage, contributed by both writes and queries. When the node heap memory usage exceeds this threshold, one of the large queries is automatically canceled based on a predefined policy.<br>Value range: 0.0–100.0%<br>The default value is <b>90%</b> .                      |

| Parameter                   | Type    | Description                                                                                                                                                                                                                     |
|-----------------------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.count.limit | Integer | The maximum number of queries in the isolation pool. When this threshold is reached, query cancelation is triggered, and no new queries will be accepted.<br>Value range: <b>10–50000</b><br>The default value is <b>1000</b> . |

 **NOTE**

In addition to **search.isolator.memory.pool.limit** and **search.isolator.count.limit**, 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.

4. Configure a policy for selecting queries to cancel in the isolation pool.

```
PUT _cluster/settings
{
 "persistent": {
 "search.isolator.strategy": "fair",
 "search.isolator.strategy.ratio": "0.5%"
 }
}
```

**Table 3-104** Parameters for configuring a query cancelation policy

| Parameter                | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.strategy | String | Policy for selecting which query to cancel when query isolation is triggered. <ul style="list-style-type: none"><li>• fair (default): Selects which query to cancel by considering both memory usage and query duration. If the difference between the memory used by two queries <math>\leq</math> <b>maximum node heap memory x fair policy threshold</b>, the query that has lasted the longest is canceled; on the contrary, if it is greater than that, the most memory-intensive query is canceled instead.</li><li>• mem-first: Cancels the query that has the highest memory usage.</li><li>• time-first: Cancels the query that has lasted the longest.</li></ul> The large query isolation pool is checked every second until the heap memory is within the safe range. |

| Parameter                      | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.strategy.ratio | String | <p>Fair policy threshold, which is the ratio of the memory difference between two queries in the isolation pool to the maximum node heap memory. If this threshold is not reached, the query that has lasted the longest will be canceled. Otherwise, the query that occupies the most memory will be canceled.</p> <p>This parameter is valid only when <b>search.isolator.strategy</b> is set to <b>fair</b>.</p> <p>Value range: 0.0–100.0%</p> <p>The default value is <b>1</b>.</p> |

## Configuring Global Query Timeout

A global query timeout, as the name indicates, applies to all queries. Global query timeout is disabled by default. Any change takes effect immediately.

Run the following command to enable and configure a global query timeout:

```
PUT _cluster/settings
{
 "persistent": {
 "search.isolator.time.enabled": true,
 "search.isolator.time.limit": "110s"
 }
}
```

**Table 3-105** Parameter description

| Parameter                    | Type    | Description                                                                                                                                                                                                                                                                                  |
|------------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.time.enabled | Boolean | <p>Whether to enable a global query timeout. When enabled, queries are automatically canceled when they last longer than a predefined timeout.</p> <ul style="list-style-type: none"><li>true: Enable global query timeout.</li><li>false (default): Disable global query timeout.</li></ul> |
| search.isolator.time.limit   | String  | <p>The value of the global query timeout.</p> <p>Value range: <math>\geq 0</math> ms</p> <p>The default value is 120s.</p>                                                                                                                                                                   |

### Configuring the Maximum Number of Log Records Retained for Query Cancellation

Run the following command to set the maximum number of log records retained for canceled queries:

```
PUT _cluster/settings
{
 "persistent": {
 "search.isolator.log.count": "100"
 }
}
```

| Parameter                 | Data Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| search.isolator.log.count | Integer   | <p>Maximum number of log records retained for canceled queries. When this limit is exceeded, excess log records are discarded.</p> <p>This parameter is valid only when <b>search.isolator.enabled</b> is set to <b>true</b>.</p> <p>Value range: 0–5000</p> <p>The default value is <b>100</b>.</p> <p>You can use the following API to check query cancellation logs:</p> <ul style="list-style-type: none"><li>• Query all nodes:<br/>GET /_isolator_metrics</li><li>• Query a single node:<br/>GET /_isolator_metrics/{nodeId}</li><li>• Query details about canceled queries on all nodes:<br/>GET /_isolator_metrics?detailed</li><li>• Query details about canceled queries on a single node:<br/>GET /_isolator_metrics/{nodeId}?detailed</li></ul> <p>In the commands above, <b>nodeId</b> indicates the node ID.</p> |

### 3.8.6 Configuring Enhanced Aggregation for an OpenSearch Cluster

CSS OpenSearch clusters enhance aggregation performance in the face of large data volumes by leveraging vectorization and optimized clustering, enabling faster analytics and decision-making in complex situations.

#### How the Feature Works

Enhanced aggregation works by pre-sorting and physically clustering data using carefully selected sorting and clustering keys, thereby minimizing data scanning and computational overheads during aggregation.

- **Sorting key:** A field (such as timestamps) used to physically order documents on disk, ensuring that documents with identical or similar sorting key values are stored contiguously.
- **Clustering key:** A subset of fields from the sorting key that groups related documents into contiguous physical blocks. This way, aggregations can process contiguous data blocks, instead of scattered documents.

**Table 3-106** Common scenarios for enhanced aggregation

| Scenario                                               | Description                                                                                                                                                                                                                                 |
|--------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Aggregation of low-cardinality fields                  | Aggregates fields that have a small number of unique values. In this case, grouping is used. One example is counting the number of orders per city. A clustering operation can easily achieve this purpose.                                 |
| Aggregation of high-cardinality fields                 | Aggregates fields that have a large number of unique values. Histogram aggregation is typically used. One example is counting the hourly visits. A clustering key can be used to accelerate data aggregation by a specified scope or range. |
| Hybrid aggregation of low- and high-cardinality fields | Groups and aggregates low-cardinality fields first (for example, orders per city), and then creates a histogram using high-cardinality fields (for example, timestamps). Multi-level clustering improves the efficiency of hybrid queries.  |

## Constraints

Only OpenSearch 2.19.0 clusters support this feature.

## Aggregation of Low-Cardinality Fields

Generally, low-cardinality fields use grouping as a way of aggregation. With appropriate sorting keys, grouping prepares the data for batch vectorization.

For example, to aggregate two low-cardinality fields **city** and **product**, perform the following steps:

1. Run the following command to set the index **testindex**:

```
PUT testindex
{
 "mappings": {
 "properties": {
 "date": {
 "type": "date",
 "format": "yyyy-MM-dd"
 },
 "city": {
 "type": "keyword"
 },
 "product": {
 "type": "keyword"
 },
 "trade": {
 "type": "double"
 }
 }
 }
}
```

```
}
}
},
"settings": {
 "index": {
 "search": {
 "turbo": {
 "enabled": "true"
 }
 },
 },
 "sort": {
 "field": [
 "city",
 "product"
]
 },
 "cluster": {
 "field": [
 "city",
 "product"
]
 }
}
}
```

**Table 3-107** Parameters for low-cardinality field aggregation

| Parameter                  | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.search.turbo.enabled | Boolean          | Whether to enable enhanced aggregation. Normally, enhanced aggregation must be enabled where aggregations are used.<br>The value can be: <ul style="list-style-type: none"><li>● <b>false</b> (default): Disable enhanced aggregation.</li><li>● <b>true</b>: Enable enhanced aggregation.</li></ul>                                                                                                                                |
| index.sort.field           | Array of strings | Specify sorting keys. Sorting keys are fields used to sequence or rank documents.<br>You can specify one or multiple fields as sorting keys. When multiple fields are specified, they will apply in the sequence in which they are specified. Documents are first ranked by the first field, then the initial result set is ranked by the second field, and so on.<br>Value range: The value must be fields contained in the index. |

| Parameter           | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.cluster.field | Array of strings | <p>Specify clustering keys. Clustering keys determine which documents are collected into the same clusters.</p> <p>During an aggregation operation, documents in the same cluster can be processed in batches, significantly enhancing aggregation performance.</p> <p>Constraint: The clustering keys must be a subset of the sorting keys.</p> <p>Value range: The value must be fields contained in the index.</p> |

2. Run the following command to import sample data.

```
PUT /_bulk
{ "index": { "_index": "testindex", "_id": "1" } }
{ "date": "2025-01-01", "city": "cityA", "product": "books", "trade": 3000.0 }
{ "index": { "_index": "testindex", "_id": "2" } }
{ "date": "2025-01-02", "city": "cityA", "product": "books", "trade": 1000.0 }
{ "index": { "_index": "testindex", "_id": "3" } }
{ "date": "2025-01-01", "city": "cityA", "product": "bottles", "trade": 100.0 }
{ "index": { "_index": "testindex", "_id": "4" } }
{ "date": "2025-01-02", "city": "cityA", "product": "bottles", "trade": 300.0 }
{ "index": { "_index": "testindex", "_id": "5" } }
{ "date": "2025-01-01", "city": "cityB", "product": "books", "trade": 7000.0 }
{ "index": { "_index": "testindex", "_id": "6" } }
{ "date": "2025-01-02", "city": "cityB", "product": "books", "trade": 1000.0 }
```

3. Run the following query command to aggregate low-cardinality fields.

```
POST testindex/_search
{
 "size": 0,
 "aggs": {
 "groupby_city": {
 "terms": {
 "field": "city"
 },
 "aggs": {
 "groupby_product": {
 "terms": {
 "field": "product"
 },
 "aggs": {
 "avg_trade": {
 "avg": {
 "field": "trade"
 }
 }
 }
 }
 }
 }
 }
}
```

## Aggregation of High-Cardinality Fields

High-cardinality fields commonly use histogram aggregation, which facilitates data processing per range or scope.

For example, to aggregate the typical high-cardinality field **date**, perform the following steps:

1. Run the following command to set the index **testindex**:

```
PUT testindex
{
 "mappings": {
 "properties": {
 "timestamp": {
 "type": "date",
 "format": "yyyy-MM-dd"
 },
 "score": {
 "type": "double"
 }
 }
 },
 "settings": {
 "index": {
 "search": {
 "turbo": {
 "enabled": "true"
 }
 },
 "sort": {
 "field": [
 "timestamp"
]
 }
 }
 }
}
```

**Table 3-108** Parameters for high-cardinality field aggregation

| Parameter                  | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.search.turbo.enabled | Boolean          | Whether to enable enhanced aggregation. Normally, enhanced aggregation must be enabled where aggregations are used.<br>The value can be: <ul style="list-style-type: none"><li>● <b>false</b> (default): Disable enhanced aggregation.</li><li>● <b>true</b>: Enable enhanced aggregation.</li></ul>                                                                                                                                |
| index.sort.field           | Array of strings | Specify sorting keys. Sorting keys are fields used to sequence or rank documents.<br>You can specify one or multiple fields as sorting keys. When multiple fields are specified, they will apply in the sequence in which they are specified. Documents are first ranked by the first field, then the initial result set is ranked by the second field, and so on.<br>Value range: The value must be fields contained in the index. |

2. Run the following command to import sample data.

```
PUT /_bulk
{ "index": { "_index": "testindex", "_id": "1" } }
```



```
{ "date": "2025-01-01", "score": "12.0" }
{ "index": { "_index": "testindex", "_id": "2" } }
{ "date": "2025-01-02", "score": "24.0" }
{ "index": { "_index": "testindex", "_id": "3" } }
{ "date": "2025-01-01", "score": "53.0" }
{ "index": { "_index": "testindex", "_id": "4" } }
{ "date": "2025-01-02", "score": "22.0" }
{ "index": { "_index": "testindex", "_id": "5" } }
{ "date": "2025-01-01", "score": "99.0" }
{ "index": { "_index": "testindex", "_id": "6" } }
{ "date": "2025-01-02", "score": "26.0" }
```

3. Run the following query command to aggregate high-cardinality fields.

This query groups the **date** field using a histogram and then calculates the average score.

```
POST testindex/_search?pretty
{
 "size": 0,
 "aggs": {
 "groupbytime": {
 "date_histogram": {
 "field": "date",
 "calendar_interval": "day"
 },
 "aggs": {
 "avg_score": {
 "avg": {
 "field": "score"
 }
 }
 }
 }
 }
}
```

## Hybrid Aggregation of Low- and High-Cardinality Fields

Where low-cardinality and high-cardinality fields are mixed, first groups and aggregates low-cardinality fields, and then aggregates high-cardinality fields using histograms.

For example, to first group the low-cardinality field **city**, then group the low-cardinality field **product**, and then group the high-cardinality field **date** into a histogram, perform the following steps:

1. Run the following command to set the index **testindex**:

```
PUT testindex
{
 "mappings": {
 "properties": {
 "date": {
 "type": "date",
 "format": "yyyy-MM-dd"
 },
 "city": {
 "type": "keyword"
 },
 "product": {
 "type": "keyword"
 },
 "trade": {
 "type": "double"
 }
 }
 },
 "settings": {
```

```
"index": {
 "search": {
 "turbo": {
 "enabled": "true"
 }
 },
 "sort": {
 "field": [
 "city",
 "product",
 "date"
]
 },
 "cluster": {
 "field": [
 "city",
 "product"
]
 }
}
```

**Table 3-109** Parameters for hybrid aggregation of low- and high-cardinality fields

| Parameter                  | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.search.turbo.enabled | Boolean          | Whether to enable enhanced aggregation. Normally, enhanced aggregation must be enabled where aggregations are used.<br>The value can be: <ul style="list-style-type: none"><li><b>false</b> (default): Disable enhanced aggregation.</li><li><b>true</b>: Enable enhanced aggregation.</li></ul>                                                                                                                                                                                                                                                           |
| index.sort.field           | Array of strings | Specify sorting keys. Sorting keys are fields used to sequence or rank documents.<br>You can specify one or multiple fields as sorting keys. When multiple fields are specified, they will apply in the sequence in which they are specified. Documents are first ranked by the first field, then the initial result set is ranked by the second field, and so on.<br>Constraint: High-cardinality fields must be among the sorting keys, and must follow the last low-cardinality field.<br>Value range: The value must be fields contained in the index. |

| Parameter           | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.cluster.field | Array of strings | <p>Specify clustering keys. Clustering keys determine which documents are collected into the same clusters.</p> <p>During an aggregation operation, documents in the same cluster can be processed in batches, significantly enhancing aggregation performance.</p> <p>Constraint: The clustering keys must be a subset of the sorting keys.</p> <p>Value range: The value must be fields contained in the index.</p> |

2. Run the following command to import sample data.

```
PUT /_bulk
{ "index": { "_index": "testindex", "_id": "1" } }
{ "date": "2025-01-01", "city": "cityA", "product": "books", "trade": 3000.0 }
{ "index": { "_index": "testindex", "_id": "2" } }
{ "date": "2025-01-02", "city": "cityA", "product": "books", "trade": 1000.0 }
{ "index": { "_index": "testindex", "_id": "3" } }
{ "date": "2025-01-01", "city": "cityA", "product": "bottles", "trade": 100.0 }
{ "index": { "_index": "testindex", "_id": "4" } }
{ "date": "2025-01-02", "city": "cityA", "product": "bottles", "trade": 300.0 }
{ "index": { "_index": "testindex", "_id": "5" } }
{ "date": "2025-01-01", "city": "cityB", "product": "books", "trade": 7000.0 }
{ "index": { "_index": "testindex", "_id": "6" } }
{ "date": "2025-01-02", "city": "cityB", "product": "books", "trade": 1000.0 }
```

3. Run the following query command to perform a hybrid aggregation of low- and high-cardinality fields.

```
POST testindex/_search
{
 "size": 0,
 "aggs": {
 "groupby_region": {
 "terms": {
 "field": "city"
 },
 "aggs": {
 "groupby_host": {
 "terms": {
 "field": "product"
 },
 "aggs": {
 "groupby_timestamp": {
 "date_histogram": {
 "field": "date",
 "interval": "day"
 },
 "aggs": {
 "avg_score": {
 "avg": {
 "field": "trade"
 }
 }
 }
 }
 }
 }
 }
 }
 }
}
```

```
}
}
```

### 3.8.7 Configuring Read/Write Splitting Between Two OpenSearch Clusters


Read/write splitting directs writes to the leader cluster and queries to the follower cluster, separating the loads. If the leader becomes unavailable, the follower automatically upgrades to leader to ensure service continuity. You may use this feature where high availability and high query performance are prioritized.

#### How the Feature Works

Read/write splitting between a pair of leader and follower clusters works as follows:

- **Data synchronization:** The leader synchronizes data changes to the follower through a REST API. Two synchronization methods are supported:
  - **Exact match:** synchronization of a specified index
  - **Match by wildcard:** batch synchronization of indexes matching a wildcard expression
- **Leader/follower switchover:** If the leader becomes unavailable, the follower automatically upgrades to leader to take over write and query loads.

**Figure 3-15** How read/write splitting works



- **Read and write split when both clusters are available (left):** The leader handles writes, and the follower handles queries.
- **Automatic failover when the leader fails (right):** If the leader becomes unavailable, the follower automatically upgrades to leader to ensure service continuity.

#### Constraints

- Only OpenSearch 2.19.0 clusters support read/write splitting.
- The leader and follower clusters must use the same software version.

#### Prerequisites

Two clusters of the same version have been created. One functions as the leader cluster, and the other the follower cluster. The follower cluster must be able to access the REST API (default port: 9200) of the leader cluster.

## Logging In to OpenSearch Dashboards

Log in to OpenSearch Dashboards and go to the command execution page. OpenSearch clusters support multiple access methods. This topic uses OpenSearch Dashboards as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
4. In the left navigation pane, choose **Dev Tools**.

## Connecting the Leader and Follower Clusters

1. Run the following command to configure information about the leader cluster in the follower cluster:

```
PUT /_cluster/settings
{
 "persistent" : {
 "cluster" : {
 "remote.rest" : {
 "leader1" : {
 "seeds" : [
 "http://10.0.0.1:9200",
 "http://10.0.0.2:9200",
 "http://10.0.0.3:9200"
],
 "username": "elastic",
 "password": "*****"
 }
 }
 }
 }
}
```

**Table 3-110** Request body parameters

| Parameter      | Description                                                                                                                       |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <i>leader1</i> | Name of the leader cluster configuration task, which is user-defined and will be used for configuring read/write splitting later. |
| seeds          | Address for accessing the leader cluster. When HTTPS is enabled for the cluster, the URL schema must use HTTPS.                   |
| username       | Username of the leader cluster. This parameter is required only when security mode is enabled for the leader cluster.             |
| password       | Password of the leader cluster. This parameter is required only when security mode is enabled for the leader cluster.             |

Example response:

```
{
 "acknowledged" : true, //Whether the operation is successful
 "persistent" : {
 "cluster" : {
 "remote" : {
 "rest" : {
 "leader1" : {
 "seeds" : [
 "http://10.0.0.1:9200",
 "http://10.0.0.2:9200",
 "http://10.0.0.3:9200"
],
 "username": "elastic",
 "password": "*****"
 }
 }
 }
 }
 },
 "transient" : { }
}
```

2. After the configuration is complete, run the following command in the follower cluster to check the connection between the follower and leader clusters:

```
GET _remote/rest/info
```

Example response:

```
{
 "leader1" : {
 "connected" : true //The two clusters are connected.
 }
}
```

## Index Synchronization

There are two ways to synchronize indexes: synchronization of a specified index and synchronization of indexes matching a wildcard expression.

During synchronization, indexes in the follower cluster become read-only. The synchronization is performed periodically. The default synchronization interval is 30 seconds. For how to change it, see [Changing the Synchronization Interval](#).

---

### CAUTION

The following index configuration items cannot be modified: **number\_of\_shards**, **version.created**, **uuid**, **creation\_date**, and **soft\_deletes.enabled**.

---

Method 1: synchronizing a specified index

- Run the following command in the follower cluster to synchronize a single index from the leader cluster to the follower cluster without modifying index settings:

```
PUT start_remote_sync
{
 "remote_cluster": "leader1",
 "remote_index": "data1_leader",
 "local_index": "data1_follower"
}
```

- Run the following command in the follower cluster to synchronize a single index from the leader cluster to the follower cluster while modifying some of the index settings—enabling synchronization of index settings:

```
PUT start_remote_sync
{
 "remote_cluster": "leader1",
 "remote_index": "data1_leader",
 "local_index": "data1_follower",
 "settings": {
 "number_of_replicas": 4
 },
 "settings_sync_enable": true,
 "settings_sync_patterns": ["*"],
 "settings_sync_exclude_patterns": ["index.routing.allocation.*"],
 "alias_sync_enable": true,
 "state_sync_enable": true
}
```

**Table 3-111** Request body parameters

| Parameter                      | Description                                                                                                                                                                                                                                         |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| remote_cluster                 | Name of the leader cluster configuration task, which was set in <a href="#">Connecting the Leader and Follower Clusters</a> . <b>leader1</b> was set in our example.                                                                                |
| remote_index                   | Name of the index to be synchronized in the leader cluster                                                                                                                                                                                          |
| local_index                    | Index name in the follower cluster                                                                                                                                                                                                                  |
| settings                       | Index settings to be synchronized                                                                                                                                                                                                                   |
| settings_sync_enable           | Whether to enable synchronization of index settings in the leader cluster. The default value is <b>false</b> .                                                                                                                                      |
| settings_sync_patterns         | Prefix of leader cluster index settings to be synchronized. The default value is *. This parameter takes effect when <b>settings_sync_enable</b> is set to <b>true</b> . The index settings configured in <b>settings</b> will not be synchronized. |
| settings_sync_exclude_patterns | Prefix of leader cluster index settings not to be synchronized. The default value is empty. This parameter is valid only when <b>settings_sync_enable</b> is set to <b>true</b> .                                                                   |
| alias_sync_enable              | Whether to enable index alias synchronization in the leader cluster. The default value is <b>false</b> .                                                                                                                                            |
| state_sync_enable              | Whether to enable index status synchronization in the leader cluster. The default value is <b>false</b> .                                                                                                                                           |

#### Method 2: synchronizing all indexes matching a wildcard expression

- Run the following command in the follower cluster to create a pattern-matching index synchronization policy, which synchronizes matched indexes from the leader cluster to the follower cluster:

```
PUT auto_sync/pattern/${PATTERN}
{
 "remote_cluster": "leader1",
 "remote_index_patterns": "log*",
 "local_index_pattern": "{{remote_index}}-sync",
}
```

```
"apply_exist_index": true
}
```

- Run the following command in the follower cluster to create a pattern-matching index synchronization policy, which synchronizes matched indexes from the leader cluster to the follower cluster, with some of the index settings modified—enabling synchronization of index settings:

```
PUT auto_sync/pattern/${PATTERN}
{
 "remote_cluster": "leader1",
 "remote_index_patterns": "log*",
 "local_index_pattern": "{{remote_index}}-sync",
 "apply_exist_index": true,
 "settings": {
 "number_of_replicas": 4
 },
 "settings_sync_enable": true,
 "settings_sync_patterns": ["*"],
 "settings_sync_exclude_patterns": ["index.routing.allocation.*"],
 "alias_sync_enable": true,
 "state_sync_enable": true
}
```

**Table 3-112** Request body parameters

| Parameter                      | Description                                                                                                                                                                                                                                         |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PATTERN                        | Name of the pattern for index matching.                                                                                                                                                                                                             |
| remote_cluster                 | Name of the leader cluster configuration task, which was set in <a href="#">Connecting the Leader and Follower Clusters</a> . In our example, <b>leader1</b> is used.                                                                               |
| remote_index_patterns          | Pattern for matching indexes to be synchronized in the leader cluster. The wildcard (*) is supported.                                                                                                                                               |
| local_index_pattern            | Index pattern in the follower cluster. The index template can be replaced. For example, if this parameter is set to <b>{{remote_index}}-sync</b> , the index <b>log1</b> changes to <b>log1-sync</b> after synchronization.                         |
| apply_exist_index              | Whether to synchronize existing indexes in the leader cluster. The default value is <b>true</b> .                                                                                                                                                   |
| settings                       | Index settings to be synchronized                                                                                                                                                                                                                   |
| settings_sync_enable           | Whether to enable synchronization of index settings in the leader cluster. The default value is <b>false</b> .                                                                                                                                      |
| settings_sync_patterns         | Prefix of leader cluster index settings to be synchronized. The default value is *. This parameter takes effect when <b>settings_sync_enable</b> is set to <b>true</b> . The index settings configured in <b>settings</b> will not be synchronized. |
| settings_sync_exclude_patterns | Prefix of leader cluster index settings not to be synchronized. The default value is empty. This parameter is valid only when <b>settings_sync_enable</b> is set to <b>true</b> .                                                                   |



| Parameter         | Description                                                                                               |
|-------------------|-----------------------------------------------------------------------------------------------------------|
| alias_sync_enable | Whether to enable index alias synchronization in the leader cluster. The default value is <b>false</b> .  |
| state_sync_enable | Whether to enable index status synchronization in the leader cluster. The default value is <b>false</b> . |

## Stopping Index Synchronization

Run the following command in the follower cluster to stop synchronization tasks for specified indexes. Subsequent changes to the indexes in the leader cluster will not be synchronized to the follower cluster. The read-only state of the indexes in the follower cluster will be cancelled, so that new data can be written into them.

```
PUT /log*/stop_remote_sync
```

In this command, **log\*** indicates the index name. You can specify multiple index names (separated by commas) or use a wildcard. In this example, synchronization tasks for all indexes that start with **log** are stopped.

## Querying and Deleting Created Patterns

1. Run the following command in the follower cluster to query created patterns:

- Query the list of patterns.

```
GET auto_sync/pattern
```

- Query a specified pattern by name.

```
GET auto_sync/pattern/{PATTERN}
```

The following is an example of the response:

```
{
 "patterns" : [
 {
 "name" : "pattern1",
 "pattern" : {
 "remote_cluster" : "leader",
 "remote_index_patterns" : [
 "log*"
],
 "local_index_pattern" : "{{remote_index}}-sync",
 "settings" : { }
 }
 }
]
}
```

2. Run the following command in the follower cluster to delete a specified pattern by name:

```
DELETE auto_sync/pattern/{PATTERN}
```

## Enabling Forcible Synchronization

By default, the plug-in determines whether to synchronize metadata based on whether the number of documents in the index of the leader cluster changes. If the leader cluster only updates documents and the number of documents remains unchanged, the plug-in does not synchronize the updates to the follower cluster. The configuration can be modified. After forcible synchronization is enabled, the

index metadata of the leader cluster is forcibly synchronized to the follower cluster in each synchronization cycle.

The following is an example of enabling forcible synchronization:

```
PUT _cluster/settings
{
 "persistent": {
 "remote_sync.force_synchronize": true
 }
}
```

## Changing the Synchronization Interval

The default synchronization interval between the leader and follower clusters is 30 seconds and can be changed.

The example request below changes the synchronization interval to 2 seconds:

```
PUT {index_name}/_settings
{
 "index.remote_sync.sync_interval": "2s"
}
```

## Changing the Synchronization Speed

You can change the data synchronization speed between the leader and follower clusters by configuring cluster-level settings.

The following is an example request: set the block size to 2 MB, the number of blocks to 20, and the maximum synchronization traffic to 100 MB per second.

```
PUT _cluster/settings
{
 "persistent": {
 "remote_sync.chunk_size": "2MB",
 "remote_sync.max_concurrent_file_chunks": 20,
 "remote_sync.max_bytes_per_sec": "100MB"
 }
}
```

**Table 3-113** Request body parameters

| Parameter                              | Description                                                                                                |
|----------------------------------------|------------------------------------------------------------------------------------------------------------|
| remote_sync.chunk_size                 | Block size for index synchronization.<br>Default value: 1 MB<br>Format: a string of characters             |
| remote_sync.max_concurrent_file_chunks | Number of blocks for concurrent index synchronization.<br>Default value: <b>10</b><br>Format: a number     |
| remote_sync.max_bytes_per_sec          | Maximum data synchronization traffic per second.<br>Default value: 40 MB<br>Format: a string of characters |

## Querying Index Synchronization Status

Obtain the auto synchronization status of a specified index.

An example request is as follows:

```
GET {index_name}/sync_stats
```

The following is an example of the response:

```
{
 "indices" : {
 "data1_follower" : {
 "shards" : {
 "0" : [
 {
 "primary" : false,
 "total_synced_times" : 27,
 "total_empty_times" : 25,
 "total_synced_files" : 4,
 "total_synced_bytes" : 3580,
 "total_paused_nanos" : 0,
 "total_paused_times" : 0,
 "current" : {
 "files_count" : 0,
 "finished_files_count" : 0,
 "bytes" : 0,
 "finished_bytes" : 0
 }
 },
 {
 "primary" : true,
 "total_synced_times" : 28,
 "total_empty_times" : 26,
 "total_synced_files" : 20,
 "total_synced_bytes" : 17547,
 "total_paused_nanos" : 0,
 "total_paused_times" : 0,
 "current" : {
 "files_count" : 0,
 "finished_files_count" : 0,
 "bytes" : 0,
 "finished_bytes" : 0
 }
 }
]
 }
 }
 }
}
```

## Switching the Roles of the Leader and Follower Clusters

When the leader cluster becomes faulty, perform a leader/follower switchover to have the follower cluster take over services. The steps are as follows:

1. Determine the index synchronization method between the leader and follower clusters. Check whether pattern-matching index synchronization policies have been configured in the follower cluster. For the command to use, see [Querying and Deleting Created Patterns](#).
  - If there are no such policies, synchronization is performed for specified indexes between the leader and follower clusters. In this case, go to [3](#).
  - If there are such policies, index synchronization between the leader and follower clusters is based on index patterns. In this case, go to [2](#).

2. Delete pattern-matching index synchronization policies in the follower cluster. For the command to use, see [Querying and Deleting Created Patterns](#).
3. Perform [Stopping Index Synchronization](#) in the follower cluster. Then redirect read and write traffic to it. If the leader and follower clusters synchronize indexes based on index patterns, use a wildcard to match indexes when running the command that stops index synchronization.
4. After the leader cluster recovers, configure information about the follower cluster in the leader cluster, and connect the leader and follower clusters again. For details, see [Connecting the Leader and Follower Clusters](#).
5. Under the leader cluster, perform [Index Synchronization](#) to synchronize data from the follower cluster to the leader cluster, and then perform a leader/follower switchover to switch back.

### 3.8.8 Switching Between Hot and Cold Storage for an OpenSearch Cluster


In an OpenSearch cluster, switching between hot and cold data storage means to allocate data to nodes of different performance standards based on data temperature (that is, how often data is accessed). The goal is to achieve optimal storage costs and query performance.

- Hot data: frequently accessed data, low latency required, stored using high-performance hardware (such as SSDs) to ensure fast write and query performance.
- Cold data: rarely accessed data, cost-optimized storage (such as HDDs).

If your cluster stores data used for different purposes, such as real-time analytics, log analytics, and monitoring data archives, you can switch between hot and cold storage for specified indexes to balance performance and costs.

## How the Feature Works

**Figure 3-16** How cold/hot storage switchover works



The key is to allocate index data storage by node labels and index allocation policies.

- Node labels:
  - Data node (hot): Stores real-time data by default and supports high-concurrency read/write requests.
  - Cold data node (cold): Stores historical data. Cold data nodes use less expensive hardware and deliver lower query performance than data nodes.
- Data allocation:
 

You can configure an index template or directly configure specific indexes to allocate data to data nodes or cold data nodes.

Procedure: Enable cold data nodes when creating a cluster, and configure an index template or specific index settings to allocate index data storage. The cluster automatically allocates data based on your settings.

## Constraints

- You cannot add cold data nodes to existing clusters that did not have such nodes enabled upon cluster creation.
- In comparison with data nodes, cold data nodes deliver lower query performance. Determine which type of node to use based on service needs.

## Switching Over Between Hot and Cold Storage

**Step 1** Check whether cold data nodes are enabled in the target cluster.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, check whether the **Node Information** area contains cold data node information.

**Figure 3-17** Cold data node information

| Node                        |              |
|-----------------------------|--------------|
| Node Specifications         | ess.s   8 GB |
| Node Storage Type           | 40 G         |
| Nodes                       | 1            |
| Cold Data Node Spec         | ess.s   8 GB |
| Cold Data Node Storage Type | 40 G         |
| Cold Data Node              | 3            |

- If there is information about cold data nodes, the cluster has cold data nodes. Go to the next step.
- Otherwise, the cluster does not have cold data nodes, in which case, you cannot switch between cold and hot data storage.

**Step 2** Log in to OpenSearch Dashboards.

1. On the cluster information page, click **Dashboards** in the upper-right corner to log in to OpenSearch Dashboards.
2. In the left navigation pane, choose **Dev Tools**.

**Step 3** On OpenSearch Dashboards, configure index allocation policies.

You can configure an index template or directly configure specific indexes to allocate data to data nodes or cold data nodes.

- **Configuring an index template**

Configure an index template to allocate indexes to cold or hot data nodes. For example, run the following command to set a template to store indexes that start with **myindex** to cold data nodes:

```
PUT _template/test
{
 "order": 1,
 "index_patterns": "myindex*",
 "settings": {
 "refresh_interval": "30s",
 "number_of_shards": "3",
```

```
"number_of_replicas": "0",
"routing.allocation.require.box_type": "cold"
}
```

Parameter description:

- **index\_patterns**: rule for matching index names (for example, **myindex\***).
- **box\_type**: node type for matching indexes. **cold** indicates cold data nodes, and **hot** indicates regular data nodes.

- **Configuring specific indexes**

To change the node type for an existing index, run the following command:

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

Parameter description:

- **myindex**: index name.
- **box\_type**: node type for matching indexes. **cold** indicates cold data nodes, and **hot** indicates regular data nodes.

#### Step 4 Verify the switchover between hot and cold storage.

Run the following command to check the distribution of index shards:

```
GET _cat/shards/myindex?v
```

When the data volume is large, the switchover may take a long time. There may be an intermediate state where the data of an index resides on both cold data nodes and data nodes.

As shown in the following figure, all shards of the **myindex** index are stored on the cold data node **css-e668-ess-cold-esn-1-1**.

| index   | shard | prirep | state   | docs     | store  | ip             | node                      |
|---------|-------|--------|---------|----------|--------|----------------|---------------------------|
| myindex | 1     | p      | STARTED | 14085446 | 17.8gb | 192.168.91.188 | css-e668-ess-cold-esn-1-1 |
| myindex | 2     | p      | STARTED | 14094005 | 17.9gb | 192.168.91.188 | css-e668-ess-cold-esn-1-1 |
| myindex | 0     | p      | STARTED | 14094742 | 17.8gb | 192.168.91.188 | css-e668-ess-cold-esn-1-1 |

#### Step 5 Roll back the cold-hot switchover configuration.

To cancel the cold-hot switchover configuration, run the following command:

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

After the rollback, index data will be evenly and randomly allocated to both cold data nodes and data nodes.

----End

## Related Operations

- You can scale cold data nodes by adding or reducing nodes or their storage capacity. For details, see [Scaling an OpenSearch Cluster](#).
- If your cluster does not have cold data nodes but you wish to cut storage costs, you can try using decoupled storage and compute. For details, see [Configuring Decoupled Storage and Compute for an OpenSearch Cluster](#).

## 3.8.9 Configuring an Index Recycle Bin for an OpenSearch Cluster

By default, when indexes are deleted in an OpenSearch cluster, the index data is deleted permanently. To prevent data loss caused by accidental deletion, CSS provides an index recycle bin. When enabled, deleted indexes are temporarily stored in the recycle bin, allowing for recovery before they are permanently removed. This feature improves data reliability and operational security.

### How the Feature Works

The index recycle bin works as follows:

- Data recycling mechanism: When an index is deleted, it is first moved to the recycle bin. From there, it will be automatically cleared when a predefined retention duration expires. Indexes in the recycle bin are still part of the cluster metadata and can be restored using an API.
- Status management
  - When the DELETE API is used to delete an index, the index is closed, and the cluster status temporarily changes to Red.
  - When the restore API is used to restore an index, the index is reopened and shards are initialized, which may also cause the cluster status to temporarily change to Red.

### Constraints

- Only OpenSearch 2.19.0 supports the index recycle bin.
- Indexes in the recycle bin are still part of the cluster metadata. Before they are removed from the recycle bin, same-name indexes cannot be created for the cluster.

### Logging In to OpenSearch Dashboards

Log in to OpenSearch Dashboards and go to the command execution page. OpenSearch clusters support multiple access methods. This topic uses OpenSearch Dashboards as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
4. In the left navigation pane, choose **Dev Tools**.

### Enabling the Index Recycle Bin

Run the following command to enable the index recycle bin:

```
PUT _cluster/settings
{
 "persistent": {
 "index.trash.enabled": true
 }
}
```



**Table 3-114** Parameters for enabling the index recycle bin

| Parameter               | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index.trash.enabled     | Boolean | Whether to enable the index recycle bin. When the index recycle bin is enabled, deleting indexes will first move them to the recycle bin. The indexes will be permanently deleted only when they are cleared from the recycle bin. <ul style="list-style-type: none"><li>• <b>true</b>: Enable the index recycle bin.</li><li>• <b>false</b> (default): Disable the index recycle bin.</li></ul> |
| indices.trash.keep.time | String  | Index retention duration in the recycle bin. Indexes will be automatically removed from the recycle bin when this duration expires.<br><br>Examples of supported time formats: <b>1d</b> (one day), <b>7d</b> (seven days), <b>1w</b> (one week), <b>1h</b> (one hour). The minimum duration is <b>1d</b> .<br><br>Default value: <b>1d</b> .                                                    |

 **NOTE**

The following command deletes an index:

```
DELETE {INDEX_NAME}
```

**INDEX\_NAME** indicates the name of the index to be deleted. Wildcards can be used to specify indexes.

## Viewing Indexes in the Recycle Bin

Run the following command to view indexes in the index recycle bin:

```
GET _cat/trash?v=true&s=index
```

**Table 3-115** Parameter description

| Parameter | Description                                                                                                                                                                                                                                               |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| v         | Whether to display the table header when the return format is a table. <ul style="list-style-type: none"><li>• <b>true</b>: display the table header.</li><li>• <b>false</b>: not display the table header.</li></ul> The default value is <b>false</b> . |
| s         | Index sorting fields. Options include <b>index</b> , <b>uuid</b> , <b>pri</b> , <b>rep</b> , <b>trash.ts</b> , and <b>delete.time</b> .                                                                                                                   |

| Parameter | Description                                                                                                                                                                                        |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| format    | Return format of the command. The default format is table. Other options include json, yaml, coBr, and smile. coBr and smile are binary formats.                                                   |
| h         | Set the column names to be displayed. By default, all column names are displayed. To specify multiple column names, separate them using commas (,), for example, <b>h=index,uuid,delete.time</b> . |

The following is an example of the output.

```
index uuid pri rep trash.ts delete.time
index1 CMD3FCLzTOyTg4RUekWNNA 1 1 1714465116615 23.6h
index1 6ATijuu6SfqamVI-WMyOKg 1 1 1714466233898 23.9h
```

**Table 3-116** Parameters in the output

| Column      | Description                                                                                                                 |
|-------------|-----------------------------------------------------------------------------------------------------------------------------|
| index       | Index name                                                                                                                  |
| uuid        | Index UUID                                                                                                                  |
| pri         | Number of shards of an index                                                                                                |
| rep         | Number of replicas of an index                                                                                              |
| trash.ts    | Time when an index was moved to the recycle bin                                                                             |
| delete.time | Remaining retention duration of an index in the recycle bin. When the value changes to 0, the index is permanently deleted. |

## Restoring an Index from the Recycle Bin

Run the following command to restore an index from the recycle bin:

```
POST /trash/recover/{INDEX_NAME}
```

**INDEX\_NAME** indicates the name of the index to be restored. Wildcards can be used to specify indexes.

## Emptying the Recycle Bin

Run the following command to empty the recycle bin:

```
POST trash/empty
```


Emptying the recycle bin will permanently delete all data in it. Please exercise caution.

## 3.9 Configuring Networking for an OpenSearch Cluster

### 3.9.1 Configuring Routes for an OpenSearch Cluster

By default, a CSS OpenSearch cluster cannot access a target service—such as a client program or a Lightweight Directory Access Protocol (LDAP) service—that is located in a different VPC. To enable it, you need to configure a route for the cluster.

**Figure 3-18** Configuring a route for an OpenSearch cluster



1. Establish a VPC peering connection to connect two VPCs.
2. On CSS, configure a route for the OpenSearch cluster that points to the client.
3. Verify network connectivity after starting services.

Configuring routes to enable cluster communication offers the following advantages:

- Improved cross-VPC communication performance: A route table specifies the traffic forwarding paths, so that traffic does not have to pass through the public network.
- Secure communication: Private networks can communicate through VPC peering connections, which is more secure than transmitting data over the public network.
- Flexible access control policies: You can customize access permissions for specific subnets.
- Support for a hybrid cloud architecture: Services can communicate across VPCs in different environments.

#### Constraints

- By default, routes can be configured only for OpenSearch clusters whose image version is no earlier than x.x.x\_25.3.0\_x.x.x.

- Adding cluster routes will alter network connections, possibly causing disruptions in the network. Exercise caution.

## Prerequisites

- The OpenSearch cluster is Available, and port 9200 is allowed in the inbound direction by the cluster's security group.
- The client is running properly, and the server's security group allows port 9200 in the outbound direction.

## Establishing a VPC Peering Connection

Before configuring routes for a CSS cluster, establish a connection between the CSS cluster's VPC and that of the client, that is, a VPC peering connection.

- If the CSS cluster and the client belong to the same account, see [Creating a VPC Peering Connection to Connect Two VPCs in the Same Account](#).
- If the CSS cluster and the client belong to different accounts, see [Creating a VPC Peering Connection to Connect Two VPCs in Different Accounts](#).

Set the peer VPC to the VPC where the CSS cluster is located and the local VPC to the VPC where the client is located.

When adding routes for a VPC peering connection, configure the destination addresses based on service requirements.

If the CSS cluster and the client share the same security group, there is no need to configure security group rules for the instances in the VPC at either end of the VPC peering connection.

## Configuring a Route for an OpenSearch Cluster

Configure a route for the OpenSearch cluster that points to the client, enabling the cluster to access the client.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Overview** tab. In the **Configuration** area, click **Add Route** next to **Cluster Route**.

**Figure 3-19** Adding a route



5. In the displayed dialog box, configure the route information.

**Table 3-117** Adding a route

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP Address  | Enter the IP address (or a CIDR block) of the client or server that the cluster needs to access, for example, <b>10.10.1.0</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Subnet Mask | <p>Enter the subnet mask used with the IP address. The subnet mask must align with the network part of the IP address. That is:</p> <ul style="list-style-type: none"><li>• The network bits (the 1s) of the subnet mask must encompass those of the IP address.</li><li>• The host bits (the 0s) of the subnet mask must accommodate the IP address's host bits.</li></ul> <p>The subnet mask should not cover too wide a range, avoiding unnecessary IP address exposure or waste of the address space. 255.255.255.0 (that is, /24) is recommended, which applies to most internal networking scenarios.</p> |

6. Click **OK** to add the route.

After the route is added, click **View Route** to check the route information. You can click **Add Route** to add another route; or click **Delete Route** to delete one—In the **Delete Route** dialog box, select a route, enter **DELETE** manually, and then click **OK**.

## Testing Connectivity

Connectivity between the OpenSearch cluster and client cannot be verified right away. After the route is configured, start services. If the search service can be accessed, the route is configured successfully.

## 3.9.2 Configuring Public Network Access for an OpenSearch Cluster

When public network access is enabled for an OpenSearch cluster, the cluster is automatically assigned a public IP address with dedicated dynamic BGP bandwidth, making it accessible from the Internet via HTTPS. You can configure access control from the public network by IP addresses or IP address ranges.

When enabling public network access for an OpenSearch cluster, a shared load balancer is typically used for load balancing. If your workloads require quicker access, you are advised to use a dedicated load balancer to connect your cluster. For details, see [Configuring a Dedicated Load Balancer for an OpenSearch Cluster](#).

## Constraints

- Enabling public network access for a CSS cluster may incur fees, as it will need to use EIP and bandwidth resources.

- To enable public network access for an OpenSearch cluster, two conditions must be met: the security mode and HTTPS access are both enabled.
- Public network access and the VPC Endpoint service share a load balancer. If you configure a whitelist for public network access, this whitelist is deployed to the shared load balancer. As such, it will control access not only from the public network but also from private IP addresses through VPCEP. In this case, you need to add IP address **198.19.128.0/17** to the public network access whitelist to allow traffic through VPCEP.

## Enabling Public Network Access

To enable public network access for an existing cluster, perform the following steps:

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, check whether **Security Mode** and **HTTPS Access** are enabled in the **Configuration** area.
  - If they are enabled, go to the next step to enable public network access.
  - If either one is disabled, public network access cannot be enabled for the cluster.

**Figure 3-20** Checking whether public network access can be enabled for the cluster

|                       |                                                |
|-----------------------|------------------------------------------------|
| Security Mode         | Enabled                                        |
| Reset Password        | <a href="#">Reset</a>                          |
| Cluster Route         | <a href="#">Add Route</a>                      |
| Enterprise Project    | <a href="#">default</a>                        |
| Public Network Access | -- <a href="#">Enable</a>                      |
| HTTPS Access          | Enabled   <a href="#">Download Certificate</a> |

5. Click **Enable** next to **Public Network Access**. In the displayed dialog box, configure the necessary settings.

**Table 3-118** Enabling public network access


| Parameter | Description                                                                         |
|-----------|-------------------------------------------------------------------------------------|
| Bandwidth | Cluster bandwidth for public network access.<br>Value range: 1 Mbit/s to 200 Mbit/s |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure Whitelist | <p>Control public network access to the cluster using a whitelist.</p> <ul style="list-style-type: none"><li>• If a whitelist is configured, only IP addresses that are on this whitelist can access the cluster over the public network.<br/>Click <b>+Add</b>. In the displayed text box, enter IP addresses or CIDR blocks that are allowed to access the cluster from the public network. Separate them using commas (,). Each value must be unique. An example of valid values: <b>192.168.1.1,10.0.0.0/24</b>. Examples of invalid values: <b>0.0.0.0</b>, <b>xx.xx.xx.x0</b>, <b>172.16.0.0-172.16.255.255</b>, non-standard formats (e.g., <b>192.168.1</b>), and duplicate values.</li><li>• If no whitelist is configured, all public IP addresses can access the cluster. However, this can be a security risk and should be avoided.</li></ul> |

6. Click **OK** to enable public network access.

After public network access is enabled, the public IP address, public network access control, and bandwidth information is displayed.

**Figure 3-21** Enabling public network access



## Managing Public Network Access

When public network access is enabled, you can check the public IP address, and modify the bandwidth and access control settings.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, manage public network access settings in the **Configuration** area.

**Figure 3-22** Managing public network access

- **Checking the public IP address**  
Record the IP address and port number displayed next to **Public Network Access**.
- **Modifying public network access control settings**  
Click **Modify** next to **Public Network Access Control**. In the displayed dialog box, add or remove IP addresses or CIDR blocks to or from the whitelist. Click **OK** to save the change.
- **Modifying public network bandwidth**  
Click **Modify** next to **Bandwidth**. In the displayed dialog box, change the bandwidth. Click **OK** to save the change.

## Disabling Public Network Access

If public network access is no longer required for a cluster, disable it to release resources.

### WARNING

After the public IP address is disassociated, the cluster can no longer be accessed from the Internet through this IP address. If you disable public network access for a cluster and then re-enable it, the public IP address for accessing the cluster may change. Exercise caution.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, find **Public Network Access** in the **Configuration** area, and click **Disable** next to it. In the displayed dialog box, enter **CONFIRM** and click **OK**.

After public network access is disabled, the public IP address and the **Public Network Access Control** and **Bandwidth** parameters disappear.

## Accessing a Cluster Through a Public IP Address

After public network access is enabled, the cluster is assigned a public IP address. You can use this IP address plus a port number to access this cluster.



For example, if the public IP address is **10.62.xxx.xxx** and the port number is **9200**, run the following cURL command to view indexes in the cluster.

```
curl -u username:password -k 'https://10.62.xxx.xxx:9200/_cat/indices'
```

where, **username** and **password** indicate the username and password of the HTTPS-enabled security-mode cluster.

### 3.9.3 Configuring VPC Endpoint Service for an OpenSearch Cluster

VPC Endpoint enables you to access resources across VPCs using a dedicated gateway, without exposing the network information of servers. A VPC endpoint can be accessed via an IPv4 address or private domain name.

- An IPv4 address is automatically allocated when VPC Endpoint is enabled.
- A private domain name is allocated only when you enable private domain names.

VPC Endpoint uses a shared load balancer for internal network access. If your workloads require faster access, we recommend that you use a dedicated load balancer to handle access to your cluster. For details, see [Configuring a Dedicated Load Balancer for an OpenSearch Cluster](#).

#### Impact on Billing

VPC endpoints, if created for the cluster, will incur extra fees, depending on the resource usage. For details, see [VPC Endpoint Billing](#).

#### Constraints


- If a shared VPC and a subnet within this shared VPC were selected earlier for the cluster on the **Network** page, **VPC Endpoint** cannot be enabled for the cluster.
- VPC endpoint creation requires specific permissions. For details, see [VPCEP Permissions](#).
- Public network access and the VPC Endpoint service share a load balancer. If you configure a whitelist for public network access, and because this whitelist is deployed to the shared load balancer, it will control not only access from the public network, but also access using private IP addresses through VPCEP. In this case, you need to add IP address **198.19.128.0/17** to the public network access whitelist to allow traffic through VPCEP.
- After VPCEP is enabled, access to CSS through a VPCEP IP address or private domain name from within the internal network is not controlled by any cluster security group rules. Rather, you need to configure a VPCEP whitelist to implement access control. For details, see [Configuring Access Control for an Interface VPC Endpoint](#).

#### Enabling VPC Endpoint

If VPC Endpoint was not enabled when the cluster was created, you can enable it as follows:

1. Log in to the [CSS management console](#).

2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, check the cluster's subnet information in the **Configuration** area. When you add a VPC endpoint, an IP address that belongs to the current subnet of the cluster is automatically assigned to it.

**Figure 3-23** Checking the current subnet

If you want to use another subnet, switch the subnet first, and then enable VPC Endpoint. For details, see [Can I Expand the Subnet for an Elasticsearch or OpenSearch Cluster?](#)


5. Choose **Cluster Access > VPC Endpoint**.
6. Toggle on **VPC Endpoint**. In the displayed dialog box, select relevant options.

**Table 3-119** Options for enabling VPC Endpoint

| Option                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create Private Domain Name | <p>Whether to create a private domain name for the VPC endpoint.</p> <ul style="list-style-type: none"><li>• Enable: The system automatically assigns a private domain name to the VPC endpoint. After cluster creation, you can check this private domain name on the <b>VPC Endpoint</b> tab of the cluster details page.</li><li>• Disable: No private domain name will be configured for the VPC endpoint. The cluster can only be accessed through an IP address assigned to the VPC endpoint.</li></ul> |

7. Click **OK** to enable VPC Endpoint. After the VPC endpoint is enabled, its information is displayed below. When its status changes to **Accepted**, the current cluster can be accessed through the VPC endpoint.

Figure 3-24 Checking the VPC endpoint status



Managing VPC Endpoints

After VPC Endpoint is enabled for a cluster, you can set access control, check VPC endpoint information, and deny access from specific VPC endpoints.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Cluster Access > VPC Endpoint**.
5. **Configure access control for a VPC endpoint.**
  - a. Click **Modify** on the right of **VPC Endpoint Whitelist**. In the displayed dialog box, add accounts that are allowed to access the cluster through the VPC endpoint. If no account is added or the account ID is set to \*, all users are allowed to access the cluster through the VPC endpoint.
    - Click **Add** to add accounts in **Account ID**. To obtain your authorized account ID, point to your username in the upper right corner, and choose **My Credentials**. Copy the value of **Account ID**.
    - Click **Delete** in the **Operation** column to delete an authorized account.
  - b. Click **OK**.
6. **Check VPC endpoint information.**

The VPC endpoint list shows the VPC endpoints created for the current cluster. You can obtain their **Status**, **Service Address**, and **Private Domain Name**.

Figure 3-25 Managing VPC endpoints

| VPC Endpoint ID | Status   | Service Address | Private Domain Name | Max. Connections | Owner      | Created          | Operation     |
|-----------------|----------|-----------------|---------------------|------------------|------------|------------------|---------------|
| 6c6e-03be-4...  | Accepted | 192...          | vpcep-03be-439...   | 3000             | c4b2-08... | Aug 30, 2024 ... | Accept Reject |

7. **Modify VPC endpoint status** to make the cluster accessible or inaccessible through specific endpoints.

- In the VPC endpoint list, select the target endpoint and click **Reject** in the **Operation** column. If the endpoint status changes to **Rejected**, it means the cluster is no longer accessible through this endpoint.
- Select an endpoint whose status is **Rejected**. Click **Accept** in the **Operation** column. If the endpoint status changes to **Accepted**, it means the cluster is accessible again through this endpoint.

## Disabling VPC Endpoint

If the cluster no longer requires cross-VPC access via VPC endpoints, disable VPC Endpoint to release resources.

### WARNING

After VPC Endpoint is disabled, the cluster is no longer accessible through a VPCEP IP address or private domain name. If you disable VPC Endpoint and then re-enable it, the VPCEP IP address or private domain name for accessing the cluster may change. When it happens, you may need to update the client connection. If you just need to temporarily disable VPC Endpoint (rather than permanently releasing its resources), do so by rejecting specific VPC endpoints. For details, see [7](#).

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Cluster Access > VPC Endpoint**.
5. Toggle off **VPC Endpoint**. In the displayed dialog box, enter **CONFIRM** and click **OK**.

## Accessing a Cluster Through a VPC Endpoint

1. Obtain the private domain name or IP address of a VPC endpoint.  
In the VPC endpoint list, check **Service Address** or **Private Domain Name**.

**Figure 3-26** Checking the service address or private domain name of a VPC endpoint

| VPC Endpoint ID  | Status   | Service Address | Private Domain Name  | Max. Connections | Owner        | Created          | Operation     |
|------------------|----------|-----------------|----------------------|------------------|--------------|------------------|---------------|
| 6c5e...03be-4... | Accepted | 192...          | vpcep-...03be-439... | 3000             | c4b2...08... | Aug 30, 2024 ... | Accept Reject |

2. On the client (e.g., an ECS), run the curl command to access the cluster.  
For example, run the following command to check the cluster's index information:
  - For a cluster with the security mode disabled:  

```
curl "http://<host>:9200/_cat/indices"
```
  - For a security-mode cluster that uses HTTP:  

```
curl -u <user>:<password> "http://<host>:9200/_cat/indices"
```
  - For a security-mode cluster that uses HTTPS:  

```
curl -u <user>:<password> -k "https://<host>:9200/_cat/indices"
```

where, **user** and **password** are the username and password used for access the cluster, and **host** indicates the VPC endpoint's private domain name or IP address.

### 3.9.4 Configuring a Dedicated Load Balancer for an OpenSearch Cluster

CSS integrates shared load balancers, through which you can enable access to a cluster from the public network as well as through the VPC Endpoint service. Dedicated load balancers provide higher performance and more diverse features than shared load balancers. This topic describes how to configure a dedicated load balancer for a cluster.

#### Scenarios

Advantages of connecting to a cluster through a dedicated load balancer:

- A non-security cluster can also utilize the capabilities of the Elastic Load Balance (ELB) service.
- You can use custom certificates for HTTPS two-way authentication.
- Seven-layer traffic monitoring and alarm configuration are supported, allowing you to keep close track of the cluster status.

There are eight different ELB service forms for clusters in different security modes to connect to a dedicated load balancer. [Table 3-120](#) describes the ELB capabilities for different cluster configurations. [Table 3-121](#) describes the configurations for different ELB service forms.

**Table 3-120** 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 mode    | No authentication                                 | Supported          | Supported              | Not supported              |
|                      | One-way authentication                            | Supported          | Supported              | Supported                  |
|                      | Two-way authentication                            |                    |                        |                            |
| Security mode + HTTP | Password authentication                           | Supported          | Supported              | Not supported              |

| Security Mode         | Service Form Provided by ELB for External Systems                                                        | ELB Load Balancing | ELB Traffic Monitoring | ELB Two-way Authentication |
|-----------------------|----------------------------------------------------------------------------------------------------------|--------------------|------------------------|----------------------------|
|                       | One-way authentication + Password authentication<br><br>Two-way authentication + Password authentication | Supported          | Supported              | Supported                  |
| Security mode + HTTPS | One-way authentication + Password authentication<br><br>Two-way authentication + Password authentication | Supported          | Supported              | Supported                  |

**Table 3-121** Configurations for different ELB service forms depending on the cluster

| Security Mode     | Service Form Provided by ELB for External Systems | ELB Listener      | ELB Listener  | ELB Listener           | Backend Server Group | Backend Server Group | Backend Server Group |
|-------------------|---------------------------------------------------|-------------------|---------------|------------------------|----------------------|----------------------|----------------------|
|                   |                                                   | Frontend Protocol | Frontend Port | SSL Authentication     | Backend Protocol     | Health Check Port    | Health Check Path    |
| Non-security mode | 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         | Service Form Provided by ELB for External Systems | ELB Listener      | ELB Listener  | ELB Listener           | Backend Server Group | Backend Server Group | Backend Server Group          |
|-----------------------|---------------------------------------------------|-------------------|---------------|------------------------|----------------------|----------------------|-------------------------------|
|                       |                                                   | Frontend Protocol | Frontend Port | SSL Authentication     | Backend Protocol     | Health Check Port    | Health Check Path             |
| 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                 |                               |

## Constraints

- You are not advised to connect a load balancer that has been associated with a public IP address to a non-security mode cluster. Allowing public network access through such a load balancer may cause security risks because a non-security mode cluster can be accessed using HTTP without security authentication.
- HTTPS-enabled security-mode clusters do not support HTTP-based frontend authentication. If the frontend uses HTTP, disable security mode for your cluster first. For details, see [Changing the Security Mode of an OpenSearch Cluster](#). Before changing the security mode, disable load balancing first. After the security mode is changed, enable load balancing again.

## Prerequisites

- A dedicated load balancer has been created. For details, see [Creating a Dedicated Load Balancer](#). This load balancer must meet the following requirements:
  - Its VPC is the same as that of the CSS cluster. The network between the two are connected.
  - **IP as a Backend** is enabled. This is necessary to connect a dedicated load balancer to a CSS cluster.
  - Determine whether to configure an EIP based on service needs. A public IP address is displayed for the load balancer connecting the CSS cluster only if an EIP is configured. This will enable public network access to the cluster through this load balancer.
- If the ELB listener uses HTTPS, upload a server certificate or CA certificate to the ELB console. For details, see [Adding a Certificate](#).
  - If one-way authentication is used, upload a server certificate.
  - If two-way authentication is used, upload a server certificate and a CA certificate.

## Connecting a Cluster to a Load Balancer

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Cluster Access** tab, and then click the **Load Balancing** tab. On the **OpenSearch** tab, toggle on **Load Balancing**. In the displayed dialog box, set the parameters.

**Table 3-122** Configuring load balancing

| Parameter     | Description                                                                                                                                                       |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Load Balancer | Select the dedicated load balancer you have created earlier.<br><br>To create a dedicated load balancer, see <a href="#">Creating a Dedicated Load Balancer</a> . |



| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>To configure a load balancer, you must have the permission to access ELB resources. By configuring an IAM agency, you can authorize CSS to access its ELB resources through an associated account.</p> <ul style="list-style-type: none"><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-elb-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>ELB Administrator</b> role or the <b>ELB FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<pre>"elb:loadbalancers:list", "elb:loadbalancers:get", "elb:certificates:list", "elb:healthmonitors:*", "elb:members:*", "elb:pools:*", "elb:listeners:*"</pre></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<pre>"iam:agencies:listAgencies", "iam:roles:listRoles", "iam:agencies:getAgency", "iam:agencies:createAgency", "iam:permissions:listRolesForAgency", "iam:permissions:grantRoleToAgency", "iam:permissions:listRolesForAgencyOnProject", "iam:permissions:revokeRoleFromAgency", "iam:roles:createRole"</pre></li><li>To use an IAM agency, the following minimum permissions are required:<pre>"iam:agencies:listAgencies", "iam:agencies:getAgency", "iam:permissions:listRolesForAgencyOnProject", "iam:permissions:listRolesForAgency"</pre></li></ul> |

**Figure 3-27** Enabling load balancing**Basic Configuration**

Load Balancer  [Create Load Balancer](#) [?](#)

Agency  [Create Agency](#) [?](#)

[Cancel](#) [OK](#)

- Click **OK** to enable load balancing.  
Load balancer information is displayed.

6. In the **Listener** area, click  to configure listener information.

**Table 3-123** Listener configuration

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Frontend Protocol  | Protocol used by the client and listener to distribute traffic.<br>Select <b>HTTP</b> or <b>HTTPS</b> .<br>Select this protocol based on your connectivity needs.                                                                                                                                                                                                                                                 |
| Frontend Port      | Port used by the client and listener to distribute traffic.<br>Set this parameter based on site requirements.                                                                                                                                                                                                                                                                                                     |
| SSL Authentication | Client authentication mode. Set this parameter only when <b>Frontend Protocol</b> is set to <b>HTTPS</b> .<br>Both one-way and two-way authentication are supported.<br>Select an authentication mode that suits your needs.                                                                                                                                                                                      |
| Server Certificate | The server certificate is used for SSL handshake. The certificate content and private key must be provided. It is required only when <b>Frontend Protocol</b> is set to <b>HTTPS</b> .<br>Select the server certificate created on ELB.                                                                                                                                                                           |
| CA Certificate     | Also called client CA public key certificate. It is used to verify the issuer of a client certificate. It is required only when <b>SSL Authentication</b> is set to <b>Two-way authentication</b> .<br>Select the CA certificate created on ELB.<br>When HTTPS two-way authentication is enabled, an HTTPS connection can be established only when the client can provide the certificate issued by a trusted CA. |

**Figure 3-28** Configuring a listener  
monitor configuration

Frontend Protocol ☐ HTTP ☒ HTTPS

Frontend Port

SSL Authentication ☐ One-way authentication ☒ Two-way authentication

Server Certificate   [Create CA Certificate](#) 

CA Certificate   [Create CA Certificate](#) 

7. (Optional) In the **Listener** area, click **Configure** next to **Access Control** to go to the listener list of the load balancer. Click **Configure** in the Access Control column of a listener to configure access control for that listener. For more information, see [What Is Access Control?](#).

Without access control policies, all IP addresses are allowed to access the CSS cluster through this load balancer, which may create security risks.

8. In the **Health Check** area, you can check the health check result for each node IP address.

**Table 3-124** Health check result description

| Health Check Result | Description                          |
|---------------------|--------------------------------------|
| Normal              | The node IP address is connected.    |
| Abnormal            | The node IP address is disconnected. |

9. If the cluster no longer needs a dedicated load balancer, disassociate it to release resources.

Choose **Load Balancing** > **OpenSearch**, toggle off **Load Balancing**. In the displayed dialog box, click **OK**.

---

 **CAUTION**

After the load balancer is disassociated, any listener or backend server group configurations will be permanently deleted.

---

## Accessing a Cluster Through a Load Balancer by Executing cURL Commands

1. In the left navigation pane on the CSS console, choose **Clusters**.
2. Log in to the [CSS management console](#).
3. In the navigation pane on the left, choose **Clusters > OpenSearch**.
4. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
5. Click the **Cluster Access** tab, and then click the **Load Balancing** tab. On the **OpenSearch** tab, record the private or public IP address or IPv6 address of the load balancer, as well as the frontend protocol/port of the listener.

### CAUTION

You are not advised to connect a load balancer that has been associated with a public IP address to a non-security mode cluster. Access from the public network using such a load balancer may cause security risks because a non-security mode cluster can be accessed using HTTP without security authentication.

6. Run the following cURL commands on an ECS to check whether the dedicated load balancer can connect to the cluster.

**Table 3-125** Commands for accessing different types of clusters

| Security Mode        | Service Form Provided by ELB for External Systems | cURL Command for Accessing a Cluster                                                       |
|----------------------|---------------------------------------------------|--------------------------------------------------------------------------------------------|
| Non-security mode    | No authentication                                 | <code>curl http://IP:port</code>                                                           |
|                      | One-way authentication                            | <code>curl --cacert ./ca.crt https://IP:port</code>                                        |
|                      | Two-way authentication                            | <code>curl --cacert ./ca.crt --cert ./client.crt --key ./client.key https://IP:port</code> |
| Security mode + HTTP | Password authentication                           | <code>curl http://IP:port -u user:pwd</code>                                               |
|                      | One-way authentication + Password authentication  | <code>curl --cacert ./ca.crt https://IP:port -u user:pwd</code>                            |

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

**Table 3-126** Variables

| Variable | Description                                                                                  |
|----------|----------------------------------------------------------------------------------------------|
| IP       | IP address of a load balancer instance.                                                      |
| port     | Frontend protocol and port configured for the listener.                                      |
| user     | Username of the cluster. This parameter is required only for a security-mode cluster.        |
| pwd      | Password of the username above. This parameter is required only for a security-mode cluster. |

If cluster information is returned, the connection is successful.

See also: [Sample Code for Accessing a Cluster Using HTTPS Two-way Authentication Through a Load Balancer \(Java\)](#).

## Sample Code for Accessing a Cluster Using HTTPS Two-way Authentication Through a Load Balancer (Java)

This section provides sample code that can be used on a Java client to access a cluster using HTTPS two-way authentication through a load balancer.

## Sample Code for ES Secured Client With Cer Demo

```
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 ES Secured Client With Cer Demo {
 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 by using 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 {
 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, 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;
}
}
```

## Sample Code for SecuredHttpClientConfigCallback

```
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)
credentials 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 Sample 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>
```

## 3.10 Backing up and Restoring the Data of an OpenSearch Cluster

### 3.10.1 Creating Snapshots to Back Up the Data of an OpenSearch Cluster

CSS allows you to use snapshots to back up and restore the data of an OpenSearch cluster. By storing a snapshot in an OBS bucket, you save a point-in-time copy of the cluster's data. By restoring this snapshot, you can restore the cluster to a previous state. There are two ways to create snapshots to back up a CSS cluster: automatic and manual.

- **Automatic snapshot creation:** Snapshots are automatically created daily or weekly based on a preset time schedule, ensuring ongoing data protection. By configuring an automatic snapshot policy, you reduce manual operations while improving backup reliability and efficiency.
- **Manual snapshot creation:** You create snapshots manually for special occasions, for example, prior to a high-risk operation (such as a cluster upgrade). This ensures you can restore data using snapshots in case anything should go wrong. Manual snapshots provide additional flexibility.

#### Impact on Billing

When you use snapshots for data backup, the generated snapshots are stored in OBS buckets, which will result in additional costs. For details, see [Object Storage Service Billing](#).

#### Constraints

- Creating cluster snapshots will increase CPU load and disk I/O. To minimize impact on system performance, you are advised to schedule it to happen during off-peak hours.
- When creating or restoring a snapshot, you can only perform certain operations, including accessing OpenSearch Dashboards, viewing metrics, and deleting other snapshots. However, you cannot perform the following operations: modifying configuration, restarting or deleting the cluster, deleting the snapshot that is being created or restored, and creating or restoring another snapshot. While a snapshot is being created or restored for a cluster, any automatic snapshot creation task initiated for that cluster will be canceled.
- When you create the first snapshot in a cluster, a full backup is performed on all data. Subsequent snapshots are all incremental, storing only the changes since the previous snapshot. Because each incremental snapshot relies on its predecessors, you must restore them in the correct sequence.

- In a cluster where storage and compute are decoupled, an index that is both frozen and closed cannot be backed up.

## Prerequisites

- The OBS bucket used for storing cluster snapshots has been created. The OBS bucket must be in the same region as the target OpenSearch cluster, and the storage class must be **Standard** or **Warm**.
- The cluster status is **Available**.

## Enabling Cluster Snapshots

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Cluster Snapshots** tab, and check whether cluster snapshots are enabled.
  - If they are enabled, skip this step.
  - If they are disabled, go to the next step.
5. On the **Cluster Snapshots** tab, click **Enable Snapshot**. In the displayed dialog box, configure the necessary settings.

**Figure 3-29** Enable Snapshot

### Cluster Snapshots

Create snapshots to save point-in-time copies of the cluster's data to OBS buckets, enabling cluster data protection and restoration.



**Table 3-127** Enabling cluster snapshots

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBS Bucket | <p>From the drop-down list, select an OBS bucket for storing snapshots.</p> <p>If no OBS buckets meet your requirements, click <b>Create Bucket</b> to go to the OBS console and create one. For details, see <a href="#">Creating a Bucket</a>.</p> <p>To grant an IAM user access to an OBS bucket, you need to grant the <b>GetBucketStoragePolicy</b>, <b>GetBucketLocation</b>, <b>ListBucket</b>, and <b>ListAllMyBuckets</b> permissions to that user.</p> |

| Parameter                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup Path                      | <p>Snapshot storage path in the OBS bucket.</p> <p>The backup path cannot:</p> <ul style="list-style-type: none"><li>• Contain the following characters: \:*?"&lt;&gt; '{}.</li><li>• Start with a slash (/).</li><li>• Start or end with a period (.).</li><li>• Contain more than two consecutive slashes (/) or periods (.).</li><li>• Exceed 512 characters.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                           |
| Maximum Backup Rate (per Second) | <p>The parameter sets the maximum backup rate per node. When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual backup rate may not reach the configured value, as it depends on factors such as OBS performance and disk I/O.</p> <p><b>Value format:</b> number + unit</p> <ul style="list-style-type: none"><li>• Number range: 0–9999</li><li>• Unit: KB, MB, GB, TB, PB, or B</li></ul> <p><b>Default value:</b> 40 MB</p> <p>The value 0MB means there is no limit on how fast data is backed up to snapshots. An overly high backup rate may lead to excessive resource usage, which may impact cluster stability. Configure this parameter carefully to maintain optimal performance.</p> |

| Parameter                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Maximum Recovery Rate (per Second) | <p>The parameter sets the maximum recovery rate per node. When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual recovery rate may not reach the configured value, as it depends on factors such as OBS performance and disk I/O.</p> <p><b>Value format:</b> number + unit</p> <ul style="list-style-type: none"><li>• Number range: 0–9999</li><li>• Unit: KB, MB, GB, TB, PB, or B</li></ul> <p>The default value is 0MB, indicating no limit.</p> <p>An overly high recovery rate may lead to excessive resource usage, which may impact cluster stability. Configure this parameter carefully to maintain optimal performance.</p> <p>For OpenSearch clusters, the recovery rate is also limited by the <b>indices.recovery.max_bytes_per_sec</b> parameter.</p> <ul style="list-style-type: none"><li>• If <b>Maximum Recovery Rate (per Second)</b> is less than <b>indices.recovery.max_bytes_per_sec</b>, the former takes effect.</li><li>• If <b>Maximum Recovery Rate (per Second)</b> is greater than <b>indices.recovery.max_bytes_per_sec</b>, the latter takes effect.</li></ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>• To check the value of <b>indices.recovery.max_bytes_per_sec</b>, run the following command:<br/>GET _cluster/settings</li><li>• To modify <b>indices.recovery.max_bytes_per_sec</b>, run the following command:<br/>PUT _cluster/settings<br/>{<br/>  "transient": {<br/>    "indices.recovery.max_bytes_per_sec": "100mb"<br/>  }<br/>}</li></ul> |

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IAM Agency | <p>Select an IAM agency to grant the current account the permission to access and use OBS. To store snapshots to an OBS bucket, you must have the required OBS access permissions.</p> <ul style="list-style-type: none"> <li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-obs-agency</b>.</li> <li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the OBS Administrator permissions deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control. <pre> "obs:bucket:GetBucketLocation", "obs:object:GetObjectVersion", "obs:object:GetObject", "obs:object:DeleteObject", "obs:bucket:HeadBucket", "obs:bucket:GetBucketStoragePolicy", "obs:object:DeleteObjectVersion", "obs:bucket:ListBucketVersions", "obs:bucket:ListBucket", "obs:object:PutObject" </pre> </li> <li>When OBS buckets use SSE-KMS encryption, the IAM agency must be granted KMS permissions. You can click <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b> to have the following custom policies created automatically. <pre> "kms:cmk:create", "kms:dek:create", "kms:cmk:get", "kms:dek:decrypt", "kms:cmk:list" </pre> </li> <li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required: <pre> "iam:agencies:listAgencies", "iam:roles:listRoles", "iam:agencies:getAgency", "iam:agencies:createAgency", "iam:permissions:listRolesForAgency", "iam:permissions:grantRoleToAgency", "iam:permissions:listRolesForAgencyOnProject", "iam:permissions:revokeRoleFromAgency", "iam:roles:createRole" </pre> </li> <li>To use an IAM agency, the following minimum permissions are required: <pre> "iam:agencies:listAgencies", "iam:agencies:getAgency", "iam:permissions:listRolesForAgencyOnProject", "iam:permissions:listRolesForAgency" </pre> </li> </ul> <p><b>WARNING</b><br/>The agency name can contain only letters (case-sensitive), digits, underscores (_), and hyphens (-). Otherwise, the backup will fail.</p> |

Determine whether to enable automatic snapshot creation based on service needs.

- To enable it, select **Automatic Snapshot Creation**, and configure necessary settings by referring to [Configuring Automatic Snapshot Creation](#).
- To disable it, unselect **Automatic Snapshot Creation**, in which case, you will need to manually create snapshots.

6. Click **OK** to enable cluster snapshots.  
The configuration information will be displayed on the **Cluster Snapshots** tab.

Configuring Automatic Snapshot Creation

You can enable automatic snapshot creation after cluster snapshots are already enabled.

1. On the **Cluster Snapshots** tab, click **Modify Settings**. In the displayed dialog box, select **Automatic Snapshot Creation**, and configure necessary settings.

Figure 3-30 Modifying Settings




Table 3-128 Configuring automatic snapshot creation

| Parameter            | Description                                                                                                                                                                                                                                                                                                                                                |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Snapshot Name Prefix | A snapshot name consists of the snapshot name prefix and timestamp, for example, <i>snapshot-1566921603720</i> .<br><br>The snapshot name prefix can contain 1 to 32 characters. It must start with a lowercase letter, and it can only contain lowercase letters, digits, hyphens (-), and underscores (_).<br><br>The default value is <b>snapshot</b> . |
| Time Zone            | Select a time zone for the backup start time.                                                                                                                                                                                                                                                                                                              |
| Backup Start Time    | Specify the start time of auto backup.<br><br>Select a value from the drop-down list. The interval can be <b>Daily</b> , <b>Hourly</b> , or weekly (by selecting a specific day of the week), and the backup time can be set to any hour from 00:00 to 23:00 (full hours only).                                                                            |

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Retained Snapshots | <p>Number of automatic snapshots to be retained. The value ranges from 1 to 90. The system automatically deletes excess snapshots every half hour. (The expiration deletion policy applies only to the snapshots that were automatically created at the same frequency as the current automatic snapshot creation policy.)</p> <p><b>NOTE</b></p> <p>If the snapshot creation interval is short or if the data size of indexes is large, the number of automatic snapshots retained may not reach the value set using this parameter.</p>                                                                                                                                                                                                                             |
| Index              | <p>Specify the name of the index to be backed up.</p> <ul style="list-style-type: none"><li>You can restore a specified index. To select multiple indexes, use commas (,) to separate them, for example, <b>index1,index2,index3</b>.</li><li>You can use an asterisk (*) to match multiple indexes. For example, <b>index*</b> indicates that all indexes with the prefix <b>index</b> will be backed up.</li></ul> <p>If you do not specify this parameter, all indexes in the cluster are backed up by default.</p> <p>The index name cannot contain spaces, uppercase letters, or special characters "<code>&lt; &gt;/?</code>".</p> <p><b>NOTE</b></p> <p>You can run the <b>GET /_cat/indices</b> command to query the names of all indexes in the cluster.</p> |

- Click **OK** to enable automatic snapshot creation.

After automatic snapshot creation is enabled, the **Automatic Snapshot Creation** parameter in the **Cluster Snapshots** area is shown as **Enabled**.

**Figure 3-31** Automatic snapshot creation enabled

#### Cluster Snapshots

Create snapshots to save point-in-time copies of the cluster's data to OBS buckets, enabling [Cluster snapshot user guide](#) 

OBS Bucket

12345678

IAM Agency

css\_obs\_agency

Backup Path

css\_repository/css-46e5-wangjing

Automatic Snapshot Creation

Enabled

Snapshots automatically created according to the snapshot policy as well as those manually created are displayed in the cluster snapshot task list.

When the snapshot status changes to **Available**, the snapshot is created. If snapshot creation fails, click on the right of the snapshot list to view the number of failed tasks and learn the failure causes. A maximum of 20 failed tasks can be displayed. When the snapshot function is disabled or the cluster is deleted, the failure records are also cleared.

3. Check whether data is successfully backed up.

In the cluster snapshot task list, click the snapshot name. The **View Details** dialog box is displayed. Check the shards and indexes that have been backed up to see if the backup is successful.

**Figure 3-32** View Details

### View Details



Created

Aug 08, 2025 17:00:09 GMT+08:00

Shards

Successful 47 | Failed 0 | Total 47

Backup Duration (s)

1

Index

|      |                    |
|------|--------------------|
| es_v | pref-              |
| 0000 | pref-              |
| 0000 | 010,es_write_pref- |
| 0000 | 034,es_write_pref- |

4. To disable automatic snapshot creation, click **Modify Settings**. In the displayed dialog box, deselect **Automatic Snapshot Creation**, and click **OK**. In the displayed dialog box, choose whether to delete snapshots created automatically. Then click **OK** to disable automatic snapshot creation.




### CAUTION

- When disabling automatic snapshot creation, if you select **Delete snapshots created automatically**, only existing snapshots will be deleted. Snapshots that are being created will still be created.
- If automatic snapshot creation is disabled without deleting automatically created snapshots, all such snapshots in the cluster cannot be manually deleted later once automatic snapshot creation is enabled again. Instead, they can only be deleted by the system based on the snapshot lifecycle policy. For example, if **Retained Snapshots** is set to **10** in the policy, the system automatically deletes excess snapshots (oldest first) on the next half hour when the number of existing snapshots exceeds this limit.

After automatic snapshot creation is disabled, the **Automatic Snapshot Creation** parameter in the **Cluster Snapshots** area is shown as **Disabled**.



Figure 3-33 Disabling automatic snapshot creation



Manually Creating a Snapshot

When snapshots are enabled for a cluster, you can manually create snapshots for it.

- 1. Under **Cluster Snapshot Tasks**, click **Manually Create Snapshot**. In the displayed dialog box, configure the snapshot policy.

Figure 3-34 Manually creating a snapshot




Table 3-129 Parameters for manually creating a snapshot

| Parameter     | Description                                                                                                                                                                                                                                                                                                                   |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Snapshot Name | <p>Set the snapshot name. Unlike automatically created snapshots, manually created snapshots do not have time information automatically added to their names.</p> <p>The value can contain 4 to 64 characters and must start with a letter. Only lowercase letters, digits, hyphens (-), and underscores (_) are allowed.</p> |

| Parameter            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Index                | <p>Specify the name of the index to be backed up.</p> <ul style="list-style-type: none"><li>You can restore a specified index. To select multiple indexes, use commas (,) to separate them, for example, <b>index1,index2,index3</b>.</li><li>You can use an asterisk (*) to match multiple indexes. For example, <b>index*</b> indicates that all indexes with the prefix <b>index</b> will be backed up.</li></ul> <p>If you do not specify this parameter, all indexes in the cluster are backed up by default.</p> <p>The index name cannot contain spaces, uppercase letters, or special characters "<code>&lt; &gt;/?</code>".</p> <p><b>NOTE</b><br/>You can run the <b>GET /_cat/indices</b> command to query the names of all indexes in the cluster.</p> |
| Snapshot Description | <p>Add a snapshot description.</p> <p>The value contains 0 to 256 characters and cannot contain special characters "<code>!&lt;&gt;=&amp;\</code>".</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

- Click **OK** to start creating snapshots.


Snapshots manually created as well as those automatically created according to the snapshot policy are displayed in the cluster snapshot task list.

When the snapshot status changes to **Available**, the snapshot is created. If snapshot creation fails, click on the right of the snapshot list to view the number of failed tasks and learn the failure causes. A maximum of 20 failed tasks can be displayed. When the snapshot function is disabled or the cluster is deleted, the failure records are also cleared.

- Check whether data is successfully backed up.

In the cluster snapshot task list, click the snapshot name. The **View Details** dialog box is displayed. Check the shards and indexes that have been backed up to see if the backup is successful.

Figure 3-35 View Details




Disabling Cluster Snapshots

If you no longer need the cluster snapshot function, you can disable it to reclaim storage resources.

On the **Cluster Snapshots** page, click **Disable Snapshot**. In the displayed dialog box, click **OK**. Disabling cluster snapshots does not automatically delete existing snapshots. Instead, you need to manually delete them on the OBS console.

Figure 3-36 Disable Snapshot



3.10.2 Restoring a Snapshot of an OpenSearch Cluster

Use a snapshot to restore a point-in-time copy of an OpenSearch cluster. You can restore it to the original cluster or to another eligible cluster. This process fetches the backed-up index data, metadata, and shard structure from OBS and reconstruct them in the selected destination cluster, ensuring data consistency. This capability is essential for operations like data rollback, cluster migration, and disaster recovery.

## Constraints

- Cluster snapshots will increase CPU usage and disk I/O. You are advised to take cluster snapshots during off-peak hours.
- If the source cluster is unavailable, you can restore its snapshots to a designated destination cluster.
- While you are creating a snapshot for a cluster or restoring it using one, you can perform certain operations, including scaling out the cluster (except when the cluster is the destination of a restoration operation), accessing OpenSearch Dashboards, viewing metrics, and deleting other snapshots. However, you cannot perform the following operations: restarting or deleting the cluster, deleting a snapshot that is being created or restored, and creating or restoring another snapshot. While a snapshot is being created or restored for a cluster, any automatic snapshot creation task initiated for that cluster will be canceled.
- When you create the first snapshot in a cluster, a full backup is performed on all data. Subsequent snapshots are all incremental, storing only the changes since the previous snapshot. Because each incremental snapshot relies on its predecessors, you must restore them in the correct sequence.
- Before snapshot restoration is completed, the indexes being restored to the destination cluster cannot be queried.
- If the two clusters have indexes with the same names but different shard structures, such indexes cannot be restored.
- The version of the destination cluster must be the same as (recommended) or later than that of the source cluster. Snapshots created for an Elasticsearch 7.x cluster can be restored to an OpenSearch 1.3.6 or 2.19.0 cluster.

## Prerequisites

In the source cluster's snapshot task list, there are snapshots whose status is **Available**. The status of the destination cluster is also **Available**.

## Restoring Data

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the source cluster. The cluster information page is displayed.
4. Click the **Cluster Snapshots** tab.
5. In the cluster snapshot task list, select a snapshot, and click **Restore** in the **Operation** column. In the displayed dialog box, configure necessary settings.

**Figure 3-37** Restoring a snapshot

| Created  | Operation                                      |
|---------------------------------------------------------------------------------------------|------------------------------------------------|
| Jul 31, 2025 20:00:07 GMT+08:00                                                             | <a href="#">Restore</a> <a href="#">Delete</a> |
| Jul 31, 2025 15:57:23 GMT+08:00                                                             | <a href="#">Restore</a> <a href="#">Delete</a> |

**Table 3-130** Restoration settings

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Index          | <p>Specify the name of the index you want to restore.</p> <p>Constraints:</p> <ul style="list-style-type: none"><li>• The value is a string of 0 to 1024 characters that cannot contain uppercase letters, spaces, or the following special characters: "<code>\&lt; &gt;/?</code>".</li><li>• When restoring an index whose name is prefixed with <b>.kibana</b>, the index name must be specified.</li><li>• The <b>.opendistro_security</b> index cannot be restored.</li></ul> <p>Value range:</p> <ul style="list-style-type: none"><li>• You can use an asterisk (*) to match multiple indexes. For example, <b>index*</b> indicates that all indexes with the prefix <b>index</b> will be restored. When an asterisk (*) is used for index matching, the <b>.opendistro_security</b> index and any system indexes whose name is prefixed with <b>.kibana</b> are filtered out by default.</li><li>• You can restore indexes by specifying their names, for example, <b>index1,index2,index3</b>.</li></ul> <p>Default value:</p> <p>By default, this parameter is left blank. That is, no index name is specified, and all indexes will be restored.</p> |
| Rename Pattern | <p>Index name matching rule. Enter a regular expression. Indexes that match the regular expression will be restored.</p> <p>The <b>Rename Pattern</b> and <b>Rename Replacement</b> take effect only when they are both configured at the same time.</p> <p>The value is a string of 0 to 1024 characters that cannot contain uppercase letters, spaces, or the following special characters: "<code>\&lt; &gt;/?</code>".</p> <p>The default value is <b>index_(.+)</b>, which indicates all indexes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

| Parameter                                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Rename Replacement                                 | <p>Rule for index renaming.</p> <p>The <b>Rename Pattern</b> and <b>Rename Replacement</b> take effect only when they are both configured at the same time.</p> <p>The value is a string of 0 to 1024 characters that cannot contain uppercase letters, spaces, or the following special characters: "<code>&lt;</code>"<code>&gt;</code>"<code>?</code>"<code>.</code>".</p> <p>The default value <b>restored_index_\$1</b> indicates that <b>restored_</b> will be added to the beginning of the names of all restored indexes.</p>                                                                                                                                                                                                                                                                                |
| Cluster                                            | <p>Select the cluster where you want to restore the data to. You can select the current cluster or another.</p> <ul style="list-style-type: none"><li>You must specify a cluster whose status is <b>Available</b>. If the status of the current cluster is <b>Unavailable</b>, you cannot restore the snapshot to the current cluster.</li><li>When you choose to restore data to another cluster, make sure the target cluster has the same or later version than the current cluster, and that the two clusters are in the same region. After a frozen index is restored using a snapshot, the index is stored on the local node and is read-only by default. To make it writable, run the following command to configure the index:<pre>PUT index_name/_settings {   "index.blocks.write": null }</pre></li></ul> |
| Overwrite same-name indexes in destination cluster | <p>Whether to overwrite same-name indexes in the destination cluster. By default, this option is unselected.</p> <p>Restoring a snapshot will overwrite existing data in the destination cluster. When there are same-name indexes in the destination cluster, you need to select this option in order to restore same-name, same-shard structure indexes. Indexes with a different shard structure cannot be restored.</p>                                                                                                                                                                                                                                                                                                                                                                                          |

6. Click **OK** to start restoring the snapshot.

If restoration succeeds, **Task Status** of the snapshot in the snapshot task list will change to **Restoration succeeded**, and the index data captured in the snapshot is reconstructed.

**Figure 3-38** Restoration succeeded

| Name/ID  | Snapshot Status  | Task Status            |
|-------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| snapshot-<br>a671bb49-...                                                                 |  Available       | --                                                                                                        |
| snapshot-<br>371bff61-..                                                                  |  Available       |  Restoration succeeded |

## Verifying Snapshot Restoration

After snapshot restoration, go to the destination cluster to check the restored index information.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
4. In the left navigation pane, choose **Dev Tools**.
5. Run the following command to check index information.

```
GET _cat/indices
```

### 3.10.3 Deleting an OpenSearch Cluster Snapshot

Delete snapshots that are no longer needed to reclaim storage spaces.

#### Constraints

- Deleting a snapshot will permanently delete its data.
- The deletion of automatically generated snapshots is subject to the following constraints:
  - While automatic snapshot creation is enabled, the system automatically deletes snapshots on the nearest half hour after their retention period expires.
  - After automatic snapshot creation is disabled, existing auto-created snapshots will not be deleted automatically, but they can be deleted manually.
  - If automatic snapshot creation is re-enabled, the system resumes automatic deletion of all auto-created snapshots upon expiration, including these created before automatic snapshot creation was previously disabled.

#### Manually Deleting a Snapshot

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Cluster Snapshots** tab.

5. In the cluster snapshot task list, select the snapshot you want to delete, and click **Delete** in the **Operation** column. In the displayed dialog box, enter **DELETE** and click **OK**.

## 3.11 Scaling an OpenSearch Cluster

### 3.11.1 Scaling Out/Up an OpenSearch Cluster

If an OpenSearch cluster struggles to maintain performance in the face of rapid data growth or sustained high memory usage, you may scale it horizontally by adding more nodes and more node types, or do so vertically by increasing the storage capacity of existing nodes.

**Table 3-131** Scaling scenarios

| Type             | Scenario                                                                                                                                 | Change Process                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Adding new nodes | If a cluster experiences rapid data growth or sustained high memory usage, you may add more nodes to it to improve cluster availability. | <p>Procedure for adding data nodes or cold data nodes:</p> <ol style="list-style-type: none"><li>1. New nodes are added and cluster configuration is modified.</li><li>2. A shard redistribution is automatically triggered to reallocate certain shards to newly added nodes.</li><li>3. After the redistribution is complete, the new nodes start to handle query and indexing requests.</li></ol> <p>Adding master or client nodes does not trigger data migration.</p> <p>During the capacity expansion, the system ensures that each shard has at least one available replica to ensure service continuity.</p> |




| Type                             | Scenario                                                                                                                                                                                                                                                                                                   | Change Process                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Adding new node types            | For a cluster that has no master or client nodes, as the load on its data plane increases, you may add master or client nodes to it to share the load of the data nodes.                                                                                                                                   | <p>Procedure for adding a new node type: New nodes are added and cluster configuration is modified. Adding master or client nodes does not trigger data migration, and therefore does not interrupt services.</p> <p>When client nodes are added, the cluster's address changes from the data node address to the client node address. You must update the client configuration to use this new address. Otherwise, the client nodes cannot function.</p> |
| Increasing node storage capacity | <ul style="list-style-type: none"><li>• If a cluster experiences rapid data growth and its storage capacity becomes insufficient, you can expand the cluster nodes' storage capacity.</li><li>• If the disk usage of a cluster remains high, you can expand the cluster nodes' storage capacity.</li></ul> | <ol style="list-style-type: none"><li>1. A request is sent to the EVS service to expand the disk capacities for all cluster nodes.</li><li>2. After disk capacity is expanded by EVS, the size of existing disk volumes is expanded on all cluster nodes.</li></ol> <p>The disk capacity expansion will not interrupt ongoing services.</p>                                                                                                               |

## Prerequisites

- The cluster status is **Available**, and there are no ongoing tasks.
- Your CSS resource quotas are sufficient for the capacity expansion you are about to perform. You can check available resources on the **Modify Configuration** page.

**Figure 3-39** Checking available resources

< | **Modify Configuration**



## Impact on Billing

For a pay-per-use cluster, you can see its new price when confirming the scale-out or scale-up on the console. After the scale-out or scale-up is complete, the new price will apply. For pricing details, see [Cloud Search Service Price Calculator](#).

For a yearly/monthly cluster, a new order will be placed after you confirm the scale-out or scale-up. Then you can check the new price.

## Constraints

- The node storage capacity of a cluster can only be increased—not decreased. Select an appropriate node storage capacity based on the data volume and projected data growth.
- For a yearly/monthly cluster, you cannot increase the number of nodes and node storage capacity at the same time. For a pay-per-use cluster, however, you can do that.
- The storage capacity of the master and client nodes in a cluster cannot be expanded.
- The storage capacity of data nodes that use local disks cannot be expanded.
- For the range of node quantities supported by each node type, see [Table 3-132](#).

**Table 3-132** Node quantity ranges

| Node Type       | Value Range                                                                                                           |
|-----------------|-----------------------------------------------------------------------------------------------------------------------|
| Data nodes      | <ul style="list-style-type: none"><li>• Without master nodes: 1 to 32</li><li>• With master nodes: 1 to 200</li></ul> |
| Master nodes    | 3, 5, 7, or 9 (must be an odd number from 3 to 9)                                                                     |
| Client nodes    | 1–32                                                                                                                  |
| Cold data nodes | 1–32                                                                                                                  |

## Change Impact

Before the change, learn about possible impacts and operation suggestions, and develop a plan to minimize these impacts.

Expanding the storage capacity of cluster nodes (vertical scaling) does not impact services, whereas adding new nodes or new node types (horizontal scaling) may have the following impacts:

- **Impact on performance**  
Adding new nodes does not interrupt services. However, after new nodes are added, data shards need to be redistributed to these nodes to balance the load, and this process will consume I/O performance. This is why you are advised to perform the operation during off-peak hours.  
To minimize this impact, it is advisable to adjust the data migration rate based on the cluster's traffic cycle: increase the data migration rate during

off-peak hours to shorten the task duration, and decrease it **before** peak hours arrive to ensure optimal cluster performance. The data migration rate is determined by the **indices.recovery.max\_bytes\_per\_sec** parameter. The default value of this parameter is the number of vCPUs multiplied by 32 MB. For example, for four vCPUs, the data migration rate is 128 MB. Set this parameter to a value between 40 MB and 1000 MB based on your service requirements.

```
PUT /_cluster/settings
```

```
{
 "transient": {
 "indices.recovery.max_bytes_per_sec": "1000MB"
 }
}
```

- Characteristics of this process

Once started, a scaling task cannot be stopped until it succeeds or fails.

## Duration

- It takes 10 to 30 minutes to add new nodes or new node types (horizontal scaling), depending on the cluster's scheduling capabilities.
- It takes 10 to 15 minutes to expand the storage capacity of cluster nodes (vertical scaling).

## Adding More Nodes or Increasing Node Storage Capacity


If a cluster experiences rapid data growth or sustained high memory usage, you may add more nodes to it or increase the cluster nodes' storage capacity. This helps to improve service and data availability and durability.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. Click the **Scale Cluster** tab.
5. Set the necessary parameters.

**Table 3-133** Adding more nodes or increasing node storage capacity

| Parameter | Description                    |
|-----------|--------------------------------|
| Action    | Select <b>Scale out</b> .      |
| Resources | Quantities of resources added. |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nodes     | <p>Increase the number of nodes and node storage capacity in the <b>Nodes</b> and <b>Node Storage Type</b> columns. You can change multiple node types at the same time.</p> <ul style="list-style-type: none"><li>For the range of node quantities supported by each node type, see <a href="#">Constraints</a>.</li><li>The value range of node storage capacity is determined by the node flavor you select. The value must be divisible by 20.</li></ul> <p><b>NOTE</b><br/>For a yearly/monthly cluster, you cannot increase the number of nodes and node storage capacity at the same time.</p> |

**Figure 3-40** Scaling out a cluster

6. Click **Next**.
7. Confirm the information and click **Submit**.
8. Click **Back to Cluster List** to go back to the **Clusters** page. The **Task Status** is **Scaling out**. When **Cluster Status** changes to **Available**, the cluster has been successfully scaled out.

## Adding New Node Types

For a cluster that has no master or client nodes, as the load on its data plane increases, you may add master or client nodes to it to share the load of the data nodes.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. On the **Modify Configuration** page, choose the **Add Master/Client Nodes** tab.  
If a cluster already has both master and client nodes, the **Add Master/Client Nodes** tab is unavailable.
5. On the Add Master/Client Nodes tab, configure the nodes.

**Table 3-134** Adding master or client nodes

| Parameter           | Description                                                                                                                                                                                                                                                                                                                          |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node types          | Select the type of nodes you want to add. <ul style="list-style-type: none"><li>Only one node type can be selected at a time. If you need to add both Master and Client nodes, you need to perform this task twice.</li><li>If the cluster already has Master or Client nodes, only the other node type is displayed here.</li></ul> |
| Node Specifications | Select node specifications based on site requirements.                                                                                                                                                                                                                                                                               |
| Nodes               | Set the number of nodes to add for this node type. For the value range, see <a href="#">Constraints</a> .                                                                                                                                                                                                                            |
| Node Storage Type   | Set the node storage type.<br>The storage capacity per master or client node is fixed to 40 GB.                                                                                                                                                                                                                                      |

**Figure 3-41** Adding master/client nodes

The screenshot shows a web interface for adding nodes to a cluster. At the top, there are two radio buttons: 'Master node' (which is selected) and 'Client node'. Below this are four tabs: 'Type', 'Node Specifications', 'Nodes', and 'Node Storage Type'. Under the 'Type' tab, 'Master Node' is selected. Under 'Node Specifications', a dropdown menu shows 'ess spec-4u1g(sold out)'. Under 'Nodes', there is a numeric input field with the value '3' and minus/plus buttons. Under 'Node Storage Type', a dropdown menu shows 'Common I/O' and a '40 GB' label.

- Click **Next**.
- Confirm the information and click **Submit**.

Return to the cluster list page. The **Task Status** of the cluster is **Scaling out**.

- If you added master nodes, you can tell they are added successfully when **Cluster Status** changes to **Available**.
- This is also true with client nodes. You can restart data nodes and cold data nodes to stop Cerebro and OpenSearch Dashboards processes on them.

## Related Documents

- You can also upgrade the node specifications and change the EVS disk type of an OpenSearch cluster. For details, see [Changing the Node Specifications of an OpenSearch Cluster](#).
- During a capacity expansion, if the resources in an AZ become insufficient, you can migrate the nodes in this AZ to another that has sufficient resources, and then perform the capacity expansion in the new AZ. For details, see [Switching AZs for an OpenSearch Cluster](#).

## 3.11.2 Scaling In an OpenSearch Cluster

If an OpenSearch cluster has excess capacity—due to off-peak traffic or reduced data volume—you can reduce its nodes to optimize costs.

**Table 3-135** Scaling scenarios

| Type                     | Scenario                                           | Change Process                                                                                                                                                                                                                                       |
|--------------------------|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Removing nodes randomly  | Randomly removes cluster nodes to optimize costs.  | 1. Migrate the shards on the to-be-removed nodes to the remaining nodes.<br><br>2. After data migration, bring the nodes offline and modify the cluster configuration.<br><br>Nodes are removed one at a time, so as to avoid interrupting services. |
| Removing specified nodes | Removes specified cluster nodes to optimize costs. |                                                                                                                                                                                                                                                      |

## Prerequisites

- The cluster status is **Available**, and there are no ongoing tasks.
- All mission-critical data has been backed up.

## Impact on Billing

For a pay-per-use cluster, you can see its new price when confirming the scale-in on the console. After the scale-in is complete, the new price will apply. For pricing details, see [Cloud Search Service Price Calculator](#).

For a yearly/monthly cluster, a refund request will be triggered after you confirm the scale-in. Then you can check the new price.

## Constraints

- During a scale-in, the data on the to-be-removed nodes needs to be migrated to the remaining nodes. The timeout for data migration per node is 48 hours. Scale-in will fail if this timeout expires. When the cluster has large quantities of data, you are advised to manually adjust the data migration rate and avoid performing the migration during peak hours.
- For a cluster without master nodes, scale-in is allowed only if the number of data nodes plus cold data nodes is at least three. During a scale-in, you can only remove less than half of the data nodes plus cold data nodes, and the remaining number must be greater than the maximum number of index replicas.
- For a cluster with master nodes, scale-in is allowed only if the cluster has at least two data nodes. During a scale-in, you can only remove less than half of the cluster's master nodes.
- After the scale-in, the cluster nodes' disk usage must be less than 80%.
- After the scale-in, there has to be at least one of each type of node in each AZ. For a cross-AZ cluster, the difference between the numbers of the same-type nodes in different AZs cannot exceed 1.
- For the range of node quantities supported by each node type, see [Table 2-137](#).

**Table 3-136** Node quantity ranges

| Node Type       | Value Range                                                                                                       |
|-----------------|-------------------------------------------------------------------------------------------------------------------|
| Data nodes      | <ul style="list-style-type: none"><li>Without master nodes: 1 to 32</li><li>With master nodes: 1 to 200</li></ul> |
| Master nodes    | 3, 5, 7, or 9 (must be an odd number from 3 to 9)                                                                 |
| Client nodes    | 1–32                                                                                                              |
| Cold data nodes | 1–32                                                                                                              |

## Change Impact

Before the change, learn about possible impacts and operation suggestions, and develop a plan to minimize these impacts.

- Impact on performance

During a scale-in, shards on the to-be-removed nodes are migrated to the remaining nodes. This process will consume I/O performance. This is why you are advised to perform the operation during off-peak hours.

To minimize this impact, it is advisable to adjust the data migration rate based on the cluster's traffic cycle: increase the data migration rate during off-peak hours to shorten the task duration, and decrease it **before** peak hours arrive to ensure optimal cluster performance. The data migration rate is determined by the `indices.recovery.max_bytes_per_sec` parameter. The default value of this parameter is the number of vCPUs multiplied by 32 MB. For example, for four vCPUs, the data migration rate is 128 MB. Set this parameter to a value between 40 MB and 1000 MB based on your service requirements.

```
PUT /_cluster/settings
{
 "transient": {
 "indices.recovery.max_bytes_per_sec": "1000MB"
 }
}
```

- Cluster load change

After a scale-in, the remaining nodes will need to handle all of the cluster's load. This may lead to higher CPU, memory, and disk I/O usage, impacting query and write performance. If shards are unevenly allocated, performance bottlenecks may occur. This is why before a scale-in, it is necessary to evaluate whether the remaining nodes have the capacity to handle the current cluster load.

- Characteristics of this process

Once started, a scaling task cannot be stopped until it succeeds or fails.

## Scale-in Duration

The following formula can be used to estimate how long a scale-in operation will take:

**Scale-in duration (min) = 5 (min) x Number of nodes to be removed + Data migration duration (min)**

where, 5 minutes indicates how long non-data migration operations (e.g., initialization) typically take per node. It is an empirical value.

**Data migration duration (min) = Total data size of the nodes to be removed (MB) ÷ [Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s)]**

where,

- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.
- The formulas above use estimates under ideal conditions. The actual migration speed depends on cluster load.

## Removing Nodes Randomly

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. Click the **Scale Cluster** tab.
5. Click **Scale in** to set parameters.


**Table 3-137** Removing nodes randomly

| Parameter | Description                      |
|-----------|----------------------------------|
| Action    | Select <b>Scale in</b> .         |
| Resources | Quantities of resources reduced. |



| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>When a node is deleted, NICs are released. This means you need to have VPC permissions. Select an IAM agency to grant the current account the permission to access and use VPC.</p> <ul style="list-style-type: none"><li>This parameter is available only when the new IAM plane is connected.</li><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-upgrade-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>VPC Administrator</b> role or the <b>VPC FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/><pre>"vpc:subnets:get",<br/>"vpc:ports:*"</pre></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</pre></li><li>To use an IAM agency, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</pre></li></ul> |
| Nodes     | <p>Reduce the number of nodes in the <b>Nodes</b> column. You can change multiple node types at the same time.</p> <p>For the range of node quantities supported by each node type, see <a href="#">Constraints</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Figure 3-42 Scaling in a cluster



6. Click **Next**.
7. Confirm the information and click **Submit**.
8. Click **Back to Cluster List** to go back to the **Clusters** page. **Task Status** is **Scaling in**. When **Cluster Status** changes to **Available**, the cluster has been successfully scaled in.

## Removing Specified Nodes

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. On the **Modify Configuration** page, click the **Scale In Nodes** tab.
5. Set scale-in parameters.

**Table 3-138** Removing specified nodes (Scale In Nodes)

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>When a node is deleted, NICs are released. This means you need to have VPC permissions. Select an IAM agency to grant the current account the permission to access and use VPC.</p> <ul style="list-style-type: none"><li>This parameter is available only when the new IAM plane is connected.</li><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-upgrade-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>VPC Administrator</b> role or the <b>VPC FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/><pre>"vpc:subnets:get",<br/>"vpc:ports:*"</pre></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</pre></li><li>To use an IAM agency, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</pre></li></ul> |
| Node Type | Expand the node type that needs be changed to show all nodes under it. Select the nodes you want to remove.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

6. Click **Next**.

7. Confirm the change information and click **Submit**. In the confirm dialog box, choose to migrate data, which helps to prevent data loss, and click **OK**.

During data migration, the system migrates all data from the to-be-removed nodes to the remaining nodes, and removes these nodes upon completion of the data migration. If the data on the to-be-removed nodes has replicas on other nodes, data migration can be skipped and the cluster change can be completed faster.

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

## Related Documents

- For an OpenSearch cluster, you can also optimize costs by changing node specifications and EVS disk types. For details, see [Changing the Node Specifications of an OpenSearch Cluster](#).
- If you want to reduce cluster nodes even though your cluster is not eligible for a scale-in operation, you can simply create a new cluster. Then you can migrate data to the new cluster using snapshots, and delete the old cluster after data is migrated.

## 3.12 Upgrading the Version of an OpenSearch Cluster

In CSS, you can upgrade the version of an OpenSearch cluster to use feature and performance enhancements or fix known issues.

**Table 3-139** Upgrade scenarios

| Upgrade Type         | Scenario                                                                                                                                                                                                                                                                                                                                                                                                         | Upgrade Process                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Same-version upgrade | Update the cluster's kernel patches. The cluster is upgraded to the latest image of the current version to fix known issues or optimize performance. For example, if the cluster version is <b>1.3.6(1.3.6_24.3.3_0102)</b> , upon a same-version upgrade, the cluster will be upgraded to the latest image <b>1.3.6(1.3.6_24.3.4_0109)</b> of version 1.3.6. (The version numbers used here are examples only.) | <ol style="list-style-type: none"><li>1. Bring a node offline, and migrate its data to other nodes.</li><li>2. Upgrade the node to the target version.</li><li>3. Add the node back into the cluster. The system automatically triggers a shard reallocation, migrating some of the shards to the new node.</li><li>4. Confirm that the node status is normal. Then upgrade the remaining nodes one by one.</li></ol> <p>Nodes in clusters are upgraded one by one to avoid interrupting services.</p> |

| Upgrade Type          | Scenario                                                                                                                                                                                                                                                                                                                                                                      | Upgrade Process |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Cross-version upgrade | Upgrade a cluster to the latest image of the target version to enhance functionality or incorporate versions. For example, if the cluster version is <b>1.3.6(1.3.6_24.3.3_1224)</b> , upon a cross-version upgrade, the cluster will be upgraded to the latest image <b>2.17.1(2.17.1_24.3.4_0109)</b> of version 2.17.1. (The version numbers used here are examples only.) |                 |

## Constraints

- A maximum of 20 clusters can be upgraded at the same time.
- During an upgrade, the data on the upgraded nodes needs to be migrated to other nodes. The timeout for data migration per node is 48 hours. Upgrade will fail if this timeout expires. When the cluster has large quantities of data, you are advised to manually adjust the data migration rate and avoid performing the migration during peak hours.
- Currently, OpenSearch 1.3.6 cannot be upgraded to OpenSearch 2.17.1.

## Upgrade Impact

Before upgrading a cluster, it is essential to assess the potential impacts and review operational recommendations. This enables proper scheduling of the upgrade time, ensuring a smooth upgrade process and minimizing service interruptions.

- Impact on performance

The nodes of a cluster are upgraded one at a time to ensure service continuity. However, data migration that occurs during the upgrade consumes I/O performance, and taking individual nodes offline still has some impact on the overall cluster performance.

To minimize this impact, it is advisable to adjust the data migration rate based on the cluster's traffic cycle: increase the data migration rate during off-peak hours to shorten the task duration, and decrease it **before** peak hours arrive to ensure optimal cluster performance. The data migration rate is determined by the **indices.recovery.max\_bytes\_per\_sec** parameter. The default value of this parameter is the number of vCPUs multiplied by 32 MB. For example, for four vCPUs, the data migration rate is 128 MB. Set this parameter to a value between 40 MB and 1000 MB based on your service requirements.

```
PUT /_cluster/settings
{
 "transient": {
 "indices.recovery.max_bytes_per_sec": "1000MB"
```

```
}
}
```

- **Impact on request handling**  
During the upgrade, nodes are upgraded one at a time. Requests sent to the node that is being upgraded may fail. To mitigate this impact, the following measures may be taken:
  - Use a VPC endpoint or a dedicated load balancer to handle access requests to your cluster, which makes sure that requests are automatically routed to available nodes.
  - Enable an exponential backoff & retry mechanism on the client (configure three retries).
  - Perform the upgrade during off-peak hours.
- **Rebuilding OpenSearch Dashboards and Cerebro**  
OpenSearch Dashboards and Cerebro will be rebuilt during the upgrade, making them temporarily unavailable. Additionally, due to cross-version compatibility issues, OpenSearch Dashboards may become unavailable during the upgrade. These problems will go away once the upgrade is completed.
- **Characteristics of the upgrade process**  
Once started, an upgrade task cannot be stopped until it succeeds or fails. An upgrade failure only impacts a single node, and does not interrupt services if there are data replicas. If necessary, you can restore a node that has failed to be upgraded by performing [Replacing Specified Nodes for an OpenSearch Cluster](#).

## Upgrade Duration

The following formula can be used to estimate how long it takes to upgrade a cluster:

**Upgrade duration (min) = 15 (min) x Total number of nodes to be upgraded + Data migration duration (min)**

where,

- 15 minutes indicates how long non-data migration operations (e.g., initialization and image upgrade) typically take per node. It is an empirical value.
- The total number of nodes is the sum of the number of data nodes, master nodes, client nodes, and cold data nodes in the cluster.

**Data migration duration (min) = Total data size (MB)/[Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s) x Concurrency]**

where,

- Typically, when concurrency is higher than 1, the actual data migration rate can only reach half of the throughput intended by the configured concurrency, depending on the cluster load.
- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.

The formulas above use estimates under ideal conditions. In practice, it is advisable to add 20% to 30% of redundancy.

## Pre-Upgrade Check

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

**Table 3-140** 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 <b>green</b> or <b>yellow</b> can work properly and have no unallocated primary shards.                                                                                                                                                                                 | The cluster is available and there are no ongoing tasks.                                                                                                                                                                                                                                                             |
| Node quantity  | System check | After an upgrade task is started, the system automatically checks node quantities. To ensure service continuity, there has to be at least two nodes of each type for a cluster in each of its AZs. For a cluster with master nodes, there has to be at least two data nodes. For a cluster without master nodes, the number of data nodes plus cold data nodes must be three at least. | <ul style="list-style-type: none"><li>• Number of nodes of each type in each AZ <math>\geq 2</math></li><li>• For a cluster with master nodes: Number of data nodes <math>\geq 2</math></li><li>• For a cluster without master nodes: Number of data nodes + Number of cold data nodes <math>\geq 3</math></li></ul> |
| 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 total disk capacity of the remaining nodes is sufficient for handling all the data of the cluster, and that the nodes' disk usage stays below 80%.                                       | After a single node is brought offline, the total disk capacity of the remaining nodes is sufficient for handling all the data of the cluster, and the nodes' disk usage stays below 80%.                                                                                                                            |

| Check Item    | Check Method | Description                                                                                                                                                                                                 | Normal Status                                                                           |
|---------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| Data replicas | System check | Check that the remaining data nodes and cold data nodes can handle the maximum number of primary and standby shards of indexes in the cluster. There should not be unallocated replicas during the upgrade. | Number of data nodes + Number of cold data nodes > Maximum number of index replicas + 1 |
| Data backup   | System check | Before the upgrade, back up data to prevent data loss caused by upgrade failures. When submitting an upgrade task, you can choose whether to check for full index snapshots.                                | 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 if custom plugins are installed in the source cluster. If a cluster has a custom plugin, upload all plugin packages of the target version on the plugin management page before the upgrade. During the upgrade, install the custom plugin in the new nodes. Otherwise, the custom plugins will be lost after the cluster is successfully upgraded. 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><b>NOTE</b></p> <p>If the uploaded plugin package is incorrect or incompatible, the plugin package cannot be automatically installed during the upgrade. As a result, the upgrade task fails. To restore a cluster, you can terminate the upgrade task and restore the node that fails to be upgraded by performing <a href="#">Replacing Specified Nodes for an OpenSearch Cluster</a>.</p> <p>After the upgrade is complete, the status of the custom plugin is reset to <b>Uploaded</b>.</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 <b>opensearch.yml</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Clusters' custom configurations are not lost after the upgrade.                        |

| Check Item              | Check Method            | Description                                                                                                                                                                                                                                                                                                                                          | Normal Status                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Non-standard operations | Manual check            | Check whether non-standard operations have been performed in the cluster. Non-standard operations refer to manual operations that are not recorded. These operations cannot be automatically passed on during the upgrade, for example, modification of the <b>opensearch_dashboards.yml</b> configuration file, system settings, and return routes. | 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 <b>opensearch_dashboards.yml</b> file, cannot be retained, and you must back up the file in advance. |
| 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.                                                                                                                                                                                                                                                                                                       |

| Check Item          | Check Method            | Description                                                                                                                                                                                                                                                                                                                | Normal Status                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Check Cluster Loads | System and manual check | <p>If the cluster is heavily loaded, there is a high probability that the upgrade will get stuck or fail. You are advised to check the cluster load before the upgrade and perform the upgrade only during off-peak hours.</p> <p>You can also choose to check the cluster load while configuring upgrade information.</p> | <ul style="list-style-type: none"><li>• <b>nodes.thread_pool.search.queue &lt; 1000</b>: Check that the maximum number of requests in the search queue is less than 1000.</li><li>• <b>nodes.thread_pool.write.queue &lt; 200</b>: Check that the maximum number of requests in the write queue is less than 200.</li><li>• <b>nodes.process.cpu.percent &lt; 90</b>: Check that the maximum CPU usage is less than 90%.</li><li>• <b>nodes.os.cpu.load_average/Number of vCPUs &lt; 80%</b>: Check that the number of running processes plus the number of processes waiting for CPUs is less than 80% of the total number of vCPUs.</li></ul> |

## Performing the Upgrade Task

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Cluster Snapshots** tab, and perform a full data backup. For details, see [Manually Creating a Snapshot](#).

When creating an upgrade task, you can choose to check whether the full index data has been backed up using snapshots. This helps to prevent data loss in case of an upgrade failure.

5. Click the **Version Upgrade** tab, and set upgrade parameters.

Table 3-141 Upgrade parameters

| Parameter    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Upgrade Type | Select an upgrade type. <ul style="list-style-type: none"><li>Same-version upgrade: upgrade kernel patches to the latest images within the current cluster version.</li><li>Cross-version upgrade: upgrade a cluster to the latest image of the target version.</li></ul>                                                                                                                                                                                                                                                                                                                                       |
| Target Image | Image of the target version. When you select an image, the image name and target version details are displayed.<br>The supported target versions are displayed in the drop-down list of <b>Target Image</b> . If no target image is available, possible causes are as follows: <ul style="list-style-type: none"><li>The current cluster is of the latest version.</li><li>The current cluster was created before 2023 and has vector indexes.</li><li>The new version images are unavailable in the current region.</li><li>The current cluster does not support the upgrade type you have selected.</li></ul> |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>After a node is upgraded, NICs need to be reattached to it. To do that, you must have the permission to access VPC resources. By configuring an IAM agency, you can authorize CSS to access its VPC resources through an associated account.</p> <ul style="list-style-type: none"><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-upgrade-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>VPC Administrator</b> role or the <b>VPC FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/><pre>"vpc:subnets:get",<br/>"vpc:ports:*"</pre></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</pre></li><li>To use an IAM agency, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</pre></li></ul> |

6. Click **Submit**.
7. In the displayed dialog box, configure upgrade check and acceleration options.
  - a. Select whether to enable **Check Full Index Snapshot**.

When enabled, CSS checks for valid snapshots by matching indexes by name. Snapshots help prevent potential data loss caused by upgrade failures.

CSS cannot check the content or backup times of snapshots. You should manually check existing snapshots. If any of them is over one month old, create the latest snapshot.
  - b. Select whether to enable **Check Cluster Loads**.

During an upgrade, data migration and node restarts will consume cluster resources and increase loads. When this option is enabled, CSS evaluates overload risks for the cluster and reduce the likelihood of a cluster upgrade failure caused by an overload.

The check items are as follows:

- **nodes.thread\_pool.search.queue < 1000**: Check that the maximum number of requests in the search queue is less than 1000.
- **nodes.thread\_pool.write.queue < 200**: Check that the maximum number of requests in the write queue is less than 200.
- **nodes.process.cpu.percent < 90**: Check that the maximum CPU usage of cluster nodes is less than 90%.
- **nodes.os.cpu.load\_average/Number of vCPUs < 80%**: Check that the number of running processes plus the number of processes waiting for CPUs is less than 80% of the total number of vCPUs.

If any of the results is abnormal, wait until the load drops or actively optimize it before performing the upgrade.


c. Set **Data migration concurrency control**.

Increasing the data migration concurrency can accelerate the upgrade process, but faster data migration leads to higher I/O usage. A higher concurrency is likely to lead to higher cluster load, which may impact cluster performance. You are advised to retain the default value **1**. The value should not exceed half of the number of data nodes.

8. Click **OK** to start the pre-upgrade check. The system automatically performs the check based on **Pre-Upgrade Check**, and if the check is successful, proceeds to the upgrade.

If the upgrade check fails, error information is displayed in the upper right corner of the console. Adjust the cluster configuration accordingly and try again.

**Figure 3-43** Upgrade check errors



When **Task Status** in the task list below changes to **Successful**, the upgrade is completed.

9. Confirm the upgrade result.
- a. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  - b. On the **Overview** tab, check the cluster version in the **Cluster Information** area to determine whether the upgrade is successful.

**Figure 3-44** Checking the cluster version

## Checking the Update Task

On the **Version Upgrade** page, find the current upgrade task in the **Jobs** list.

Expand the task, and click **View Progress** to check the upgrade progress.

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.

After an upgrade task is terminated, nodes that have failed to be upgraded will remain in the current state, while those that have been successfully upgraded will not be rolled back to the old version. As a result, the nodes in the same cluster may have different versions. You are advised to restart the upgrade as soon as relevant issues are handled, ensuring that all nodes have the same version.

---

**WARNING**

During a cross-version upgrade, if **Task Status** is **Failed**, you can terminate the upgrade task only if all nodes have not been upgraded.

---

## 3.13 Managing OpenSearch Clusters

### 3.13.1 Viewing OpenSearch Cluster Information

On the CSS management console, you can check an OpenSearch cluster's status, version, nodes, configurations and settings, and resource allocation on the cluster list and the cluster information page.

#### Introduction to the Cluster List

A cluster list displays all clusters of the same type under the current account. If there are a large number of clusters, the cluster list will consist of multiple pages, showing clusters of all states.

Clusters are listed in chronological order by default in the cluster list, with the most recent clusters displayed at the top.

You can click **Export** to download the full cluster list.

**Table 3-142** Cluster list parameters

| Parameter               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name/ID                 | Name and ID of a cluster. You can click a cluster name to switch to the <b>Cluster Information</b> page. The cluster ID is automatically generated by the system and uniquely identifies a cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Cluster Status          | <p>Current status of a cluster.</p> <ul style="list-style-type: none"><li>• Available: The cluster is running properly.</li><li>• Creating: The cluster is being created.</li><li>• Processing: The cluster is being restarted, scaled out, backed up, or restored.</li><li>• Abnormal: The cluster failed to be created or is unavailable.</li></ul> <p><b>WARNING</b></p> <p>If a cluster is unavailable, you can delete it or use a snapshot created when it was available to restore its data to another cluster. However, you cannot scale out this cluster, access it through OpenSearch Dashboards, create snapshots for it, or use an existing snapshot to restore its data to itself. Do not try to import data into a cluster that is unavailable, as the data may get lost. For such a cluster, you can check its monitoring data or restart it to try to fix the problem. If the fault persists, contact technical support.</p> |
| Task Status             | Status of a task, such as cluster restart, scale-out, or scale-in.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Version                 | <p>Cluster version and image version.</p> <p>For example, if the version is <b>1.3.6(1.3.6_x.x.x_x.x.x)</b>, <b>1.3.6</b> is the cluster version and <b>1.3.6_x.x.x_x.x.x</b> enclosed in parentheses is the image version.</p> <p><b>CAUTION</b></p> <ul style="list-style-type: none"><li>• For a cluster created a long time ago, its image version may not be available because it is likely not recorded in the database.</li><li>• If <b>EOS</b> is displayed below the version number, the current version has reached EOS and hence poses certain service availability risks. You are advised to upgrade the cluster version as soon as possible. For details, see <a href="#">Upgrading the Version of an OpenSearch Cluster</a>.</li></ul>                                                                                                                                                                                        |
| Private Network Address | Private network address and port number of a cluster. You can use it to access the cluster. If the cluster has multiple nodes, the private network addresses and port numbers of all nodes are displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| VPC Endpoint Address    | IP address or private domain name of the VPC endpoint associated with the cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Security Mode           | Whether the security mode is enabled for the cluster. Only authorized users can access a security-mode cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |



| Parameter           | Description                                                                                                                                                    |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enterprise Project  | Enterprise project that a cluster belongs to.                                                                                                                  |
| Cluster Description | Provides supplementary information about the cluster.                                                                                                          |
| Created             | Cluster creation time.                                                                                                                                         |
| Billing Mode        | Billing mode of a cluster. It can be <b>Pay-per-use</b> or <b>Yearly/Monthly</b> .                                                                             |
| Operation           | Operations that can be performed on a cluster, such as checking metrics, restart, and deletion. If an operation is unavailable, the menu button is grayed out. |

## Viewing Basic Cluster Information

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, check the cluster's basic information, configuration, and node information.

**Table 3-143** Cluster information

| Type                | Parameter           | Description                                                                                                                                                                                       |
|---------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Information | Name                | Cluster name. The name can be customized.<br><br>You can click  on the right to change the cluster name.     |
|                     | Initial Name        | Name of the cluster when it was created for the first time, which is unique and cannot be changed.                                                                                                |
|                     | ID                  | Unique ID of a cluster, which is automatically generated by the system.<br><br>Each cluster in the same region has a unique ID.                                                                   |
|                     | Cluster Description | Provides supplementary information about the cluster. You can click  on the right to modify the description. |

| Type          | Parameter                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               | Version                       | <p>Cluster version.</p> <p>For example, if the version is <b>1.3.6(1.3.6_x.x.x_x.x.x)</b>, <b>1.3.6</b> is the cluster version and <b>1.3.6_x.x.x_x.x.x</b> enclosed in parentheses is the image version.</p> <p><b>CAUTION</b></p> <ul style="list-style-type: none"><li>For a cluster created a long time ago, its image version may not be available because it is likely not recorded in the database.</li><li>If <b>EOS</b> is displayed below the version number, the current version has reached EOS and hence poses certain service availability risks. You are advised to upgrade the cluster version as soon as possible. For details, see <a href="#">Upgrading the Version of an OpenSearch Cluster</a>.</li></ul> |
|               | Cluster Status                | Current status of a cluster                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|               | Task Status                   | Current task status of a cluster. If there is no ongoing task, -- 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 a cluster belongs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

| Type | Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|      | Current Subnet | <p>Subnet to which a cluster belongs.</p> <p>If the current subnet cannot meet your service requirements, click <b>Switch Subnet</b>. In the displayed dialog box, select another subnet. If no existing subnets meet your requirements, click <b>Create Subnet</b> to create a new subnet.</p> <p><b>CAUTION</b></p> <ul style="list-style-type: none"><li>• After switching the subnet, the new subnet is used for adding nodes, including scaling, adding dedicated master/client nodes, and enabling VPC Endpoint. IP addresses will be allocated to newly added nodes from the new subnet.</li><li>• A cluster can be associated with a maximum of two subnets.</li><li>• Subnet switching does not affect the network settings of existing nodes.</li></ul> |
|      | Used Subnet    | Subnet once used by the cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|      | Security Group | <p>Security group to which a cluster belongs.</p> <p>To modify the security group of a cluster, click <b>Change Security Group</b> on the right.</p> <p><b>WARNING</b></p> <p>The selected security group must allow port 9200 in the inbound direction. Otherwise, the cluster may be inaccessible to external services. Furthermore, to enable OpenSearch Dashboards private network access, port 5601 must also be allowed.</p>                                                                                                                                                                                                                                                                                                                                |
|      | Security Mode  | <p>Security mode of a cluster.</p> <ul style="list-style-type: none"><li>• Enabled: The security mode is enabled for the current cluster.</li><li>• Disabled: The security mode is disabled for the current cluster.</li></ul> <p>For details about how to change the security mode of a cluster, see <a href="#">Changing the Security Mode of an OpenSearch Cluster</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                    |

| Type | Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|      | Reset Password     | <p>This parameter is available only for security-mode clusters.</p> <p>Click <b>Reset</b> to change the <b>admin</b> password for a security-mode cluster.</p> <p>Requirements for admin passwords:</p> <ul style="list-style-type: none"><li>• The password can contain 8 to 32 characters.</li><li>• 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 allowed: ~!@#\$%^&amp;*()-_+=\ []{};,:&lt;.&gt;/?</li><li>• Do not use the administrator name, or the administrator name spelled backwards, as the password.</li><li>• Change the password periodically.</li></ul> |
|      | Cluster Routing    | <p>Cluster route information. You can view, add, or modify routes for a cluster. For details, see <a href="#">Configuring Routes for an OpenSearch Cluster</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|      | Enterprise Project | <p>Enterprise project that a cluster belongs to.</p> <p>You can click a project name to check the basic project information.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|      | Public IP Address  | <p>Public IP address that enables public network access. It is displayed only for security-mode clusters.</p> <ul style="list-style-type: none"><li>• For a security-mode cluster with public network access enabled, the configured public IP address is displayed.</li><li>• For a security-mode cluster with public network access disabled, -- is displayed.</li></ul> <p>When public network access is enabled for a cluster, you are advised to use a whitelist for access control. For details about how to configure public network access, see <a href="#">Configuring Public Network Access for an OpenSearch Cluster</a>.</p>                                                                        |

| Type | Parameter                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|      | Public Network Access Control | <p>You can configure public network access control for a security-mode cluster by IP addresses or IP address ranges. This parameter is displayed only for clusters with public network access enabled.</p> <p>Click <b>Modify</b> to modify the access control whitelist.</p> <ul style="list-style-type: none"><li>• If a whitelist is configured, only IP addresses that are on this whitelist can access the cluster over the public network. Click <b>+Add</b>. In the displayed text box, enter IP addresses or CIDR blocks that are allowed to access the cluster from the public network. Separate them using commas (.). Each value must be unique. An example of valid values: <b>192.168.1.1,10.0.0.0/24</b>. Examples of invalid values: <b>0.0.0.0, xx.xx.xx.x0, 172.16.0.0-172.16.255.255</b>, non-standard formats (e.g., <b>192.168.1</b>), and duplicate values.</li><li>• If no whitelist is configured, all public IP addresses can access the cluster. However, this can be a security risk and should be avoided.</li></ul> |
|      | Bandwidth                     | <p>The bandwidth for public network access. This parameter is displayed only for clusters with public network access enabled.</p> <p>Click <b>Edit</b> to change the bandwidth size.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

| Type             | Parameter            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                  | HTTPS Access         | <p>Whether to enable HTTPS for a cluster.</p> <ul style="list-style-type: none"><li>• If HTTPS is disabled, the cluster uses HTTP instead. For a cluster with the security mode disabled, HTTPS access is disabled by default.</li><li>• If it is enabled, HTTPS is used for cluster access. Only security-mode clusters support HTTPS. If <b>HTTPS Access</b> is enabled, you can click <b>Download Certificate</b> to obtain the certificate used for access control. Currently, the security certificate cannot be used in a public network environment.</li></ul> <p>For details about how to switch the network protocol for a security-mode cluster, see <a href="#">Switching between HTTP and HTTPS in Security Mode</a>.</p> |
|                  | Private IPv4 Address | <p>Private IPv4 address and port number of a cluster. If the cluster has only one node, the IP address and port number of the only node are displayed, for example, <b>10.62.179.32:9200</b>. If the cluster has multiple nodes, the IP addresses and port numbers of all nodes are displayed, for example, <b>10.62.179.32:9200,10.62.179.33:9200</b>.</p>                                                                                                                                                                                                                                                                                                                                                                           |
|                  | Private IPv6 Address | <p>Private IPv6 address and port number of the cluster. They are displayed only when the VPC selected for the cluster supports IPv6 addresses and that an IPv6 address has been assigned automatically.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Node Information | Node Specifications  | <p>Specifications of nodes in a cluster</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|                  | Node Storage Type    | <p>Storage capacity and storage type of nodes in a cluster</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|                  | Nodes                | <p>Number of nodes in a cluster</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## 3.13.2 Creating Users for an OpenSearch Cluster and Granting Cluster Access

CSS limits access to security-mode clusters to authorized users only. When creating a security-mode cluster, an administrator account must be created. This administrator account can use OpenSearch Dashboards to add new users for the cluster and grant them the required permissions.

### Context

CSS uses the `opendistro_security` plug-in to provide security cluster capabilities. The `opendistro_security` plug-in is built based on the RBAC model. RBAC involves three core concepts: user, action, and role. RBAC simplifies the relationship between users and actions, simplifies permission management, and facilitates permission expansion and maintenance. The following figure shows the relationship between the three.

**Figure 3-45** User, action, and role




**Table 3-144** Parameters for creating and authorizing a user on OpenSearch Dashboards

| Parameter    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Permission   | A single permission, for example, creating an index (for example, <b>indices:admin/create</b> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Action group | An action group is a group of permissions. For example, the predefined <b>SEARCH</b> action group grants roles permissions to use <b>_search</b> and <b>_msearch</b> APIs.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Role         | A role is a combination of permissions or action groups, including operation permissions on clusters, indexes, documents, or fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| User         | A user can send operation requests to an OpenSearch cluster. The user has credentials such as username and password, and zero or multiple backend roles and custom attributes.                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Role mapping | A user will be assigned a role after successful authentication. Role mapping means to map a role to a user (or a backend role). For example, the mapping from <b>Dashboards_user</b> (role) to <b>Bob</b> (user) means that Bob obtains all permissions of <b>Dashboards_user</b> after authentication. Similarly, the mapping from <b>all_access</b> (role) to <b>admin</b> (backend role) means that any user with the backend role <b>admin</b> (from the LDAP/Active Directory server) has all the permissions of role <b>all_access</b> after being authenticated. You can map each role to multiple users or backend roles. |

On OpenSearch Dashboards, you can configure user permissions on the OpenSearch cluster under **Security** to implement fine-grained access control at four levels: cluster, index, document, and field.

Users can be added or deleted for a cluster, and mapped to roles. This way, you assign roles to users.

With role mapping, you can configure the members of each role and assign roles to users based on usernames, backend roles, and host names. For each role, you can configure cluster, index, and document access permissions, as well as the permission to use OpenSearch Dashboards.

For more about security configuration for a security-mode cluster and the detailed guide, see the official OpenSearch document [About Security in OpenSearch](#).

## Constraints

You can customize the username, role name, and tenant name in the **OpenSearch Dashboards**. Chinese characters are not allowed.

## Creating a User and Granting Permissions


**Step 1** Log in to the OpenSearch Dashboards.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, click **Dashboards** in the **Operation** column, and use an administrator account to log in to OpenSearch Dashboards.
  - Username: **admin** (default administrator account name)
  - Password: Enter the administrator password you set when creating the cluster in security mode.

**Step 2** Creating a user.

1. On the **OpenSearch Dashboards** page, choose **Security**. The **Security** page is displayed.


**Figure 3-46** Going to the Security page






2.
- Choose **Internal users** on the left. The user creation page is displayed.

Figure 3-47 Creating a user




3.
- Click **Create internal user**. The user information configuration page is displayed.
4.
- In the **Credentials** area, enter the username and password.

Figure 3-48 Entering the username and password



5.
- Click **Create**. After the user is created, it is displayed in the user list.

Figure 3-49 User information



Step 3 Create a role and grant permissions to the role.

1.
- Select **Roles** from the **Security** drop-down list box.
2.
- On the **Roles** page, click **Create role**. The role creation page is displayed.
3.
- In the **Name** area, set the role name.

**Figure 3-50** Setting a role name

**Name**

Specify a descriptive and unique role name. You cannot edit the name once the role is created.

Role1

The role name must contain from 2 to 50 characters. Valid characters are A-Z, a-z, 0-9, (\_) underscore, (-) hyphen and unicode characters.

4. On the **Cluster Permissions** page, set the cluster permission. Set cluster permissions based on service requirements. If this parameter is not specified for a role, the role has no cluster-level permissions.

**Figure 3-51** Assigning cluster-level permissions

**Cluster permissions**

Specify how users in this role can access the cluster. By default, no cluster permission is granted. [Learn more](#)

**Cluster Permissions**

Specify permissions using either action groups or single permissions. An action group is a list of single permissions. You can often

kibana\_all\_write ×

Create new permission group

5. In the **Index Permissions** area, set the index permission.

**Figure 3-52** Setting index permissions

**Index permissions**

Index permissions allow you to specify how users in this role can access the specific indices. By default, no index permission is granted. [Learn more](#)

my\_store

**Index**

my\_store ×

Specify index pattern using \*

**Index permissions**

You can specify permissions using both action groups or single permissions. A permission group is a list of single permissions. You can often achieve your t

search ×

Create new permission group

6. On the **Tenant Permissions** page, set role permissions.

**Figure 3-53** Role permissions

**Tenant permissions**

Tenants are useful for safely sharing your work with other OpenSearch Dashboards users. You can control which roles have access to a tenant and v

**Tenant**

global\_tenant ×

Read and Write


Remove

Add another tenant permission


After the setting is complete, you can view the created role on the **Roles** page.

**Step 4** Map a user with a role to bind them.

1. Select **Roles** from the **Security** drop-down list box.
2. On the **Roles** page, select the role to be mapped. The role mapping page is displayed.



3. On the **Mapped users** tab page, click **Map users** and select the user to be mapped from the **users** drop-down list box.



4. Click **Map**.
5. After the configuration is complete, you can check whether the configuration has taken effect in OpenSearch Dashboards.

----End

### 3.13.3 Setting Tags for an OpenSearch Cluster

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

Tags can be added during or after cluster creation. If you want to use the same tag to identify multiple cloud resources for better resource grouping, we recommend that you predefine tags in Tag Management Service (TMS). For details, see [Predefined Tags Overview](#).

## Constraints

- Each cluster can have a maximum of 20 tags.
- If your organization has configured tag policies for CSS, add tags to clusters based on these policies. If a tag does not comply with existing tag policies, cluster creation may fail. Contact the administrator to learn more about these policies.

## Managing Tags

In CSS, you can add, delete, modify, and query cluster tags.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Select the **Tags** tab, where 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 **Edit Tag** in the upper left corner. In the displayed dialog box, click **Add Tag**, enter the tag key and value, and click **OK**.
  - Modify  
Click **Edit Tag** in the upper left corner. In the displayed dialog box, modify the key and value of a tag, and click **OK**.
  - Delete  
Click **Edit Tag** in the upper left corner. In the displayed dialog box, click **Delete** in the row that contains a tag, and click **OK**.

## Searching for Clusters by Tag

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the search box, and select a tag key and value to search for the target cluster.

You can select a tag key and value from the drop-down list. The system returns a list of clusters that exactly match the tag key and value. If you select multiple tags, only clusters including all of the selected tags will be returned.

The filter can include a maximum of 10 tags.

### 3.13.4 Configuring Default Parameters in the .yaml Configuration File of an OpenSearch Cluster

You can modify an OpenSearch cluster's configuration file for purposes like accelerating queries, modifying cross-domain access configuration, adjusting the internal cache size, and managing the task queue size. The core configuration information of an OpenSearch cluster is stored in a named **opensearch.yaml**. You can modify specific parameters in this file on the CSS management console.

## Configuring the .yaml File

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Cluster Settings > Parameter Settings**.
5. Click **Edit** to modify module parameters as required.

**Table 3-145** Module parameters

| Module Name         | Parameter                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cross-domain Access | http.cors.allow-credentials | Whether to carry authentication information in cross-domain requests, that is, whether to contain the Access-Control-Allow-Credentials field in the response header.<br><br>The value can be: <ul style="list-style-type: none"><li>• <b>true</b>: The response header contains the Access-Control-Allow-Credentials field.</li><li>• <b>false</b> (default value): The response header does not contain the Access-Control-Allow-Credentials field.</li></ul> |
|                     | http.cors.allow-origin      | Source IP addresses or domain names allowed for cross-domain access. When there are multiple values, separate them with commas (,).<br><br>Example: <b>192.168.122.122:9200</b> or <b>192.168.1.1:9200, 192.168.1.2:9200</b>                                                                                                                                                                                                                                   |
|                     | http.cors.allow-headers     | Request header fields for cross-domain access. When there are multiple values, separate them using a comma (,) and a space.<br><br>Example: <b>X-Requested-With, Content-Type, Content-Length</b><br><br>The value can contain a maximum of 1024 characters. It can only contain letters, digits, hyphens (-), underscores (_), colons (:), and slashes (/).                                                                                                   |
|                     | http.cors.max-age           | Maximum retention duration of pre-check responses in the browser cache.<br><br>Unit: s<br>Value range: 0 to 31,536,000 (0 to 1 year)<br>Default value: 1,728,000 (20 days)                                                                                                                                                                                                                                                                                     |

| Module Name                 | Parameter                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                             | http.cors.enabled            | Whether to enable cross-domain access.<br>The value can be: <ul style="list-style-type: none"><li>• <b>true</b>: Enable cross-domain access.</li><li>• <b>false</b> (default): Disable cross-domain access.</li></ul>                                                                                                                                                                                                                                                                                                                                  |
|                             | http.cors.allow-methods      | HTTP methods allowed for cross-domain access. When there are multiple values, separate them using a comma (,) and a space.<br>Example: <b>OPTIONS, GET, POST</b><br>Value range: OPTIONS, HEAD, GET, POST, PUT, and DELETE                                                                                                                                                                                                                                                                                                                             |
| Reindexing                  | reindex.remote.whitelist     | Which remote hosts are allowed for remote reindexing operations. Use commas (,) to separate multiple values.<br>Example: <b>192.168.122.122:9200</b> or <b>192.168.1.1:9200, 192.168.1.2:9200</b>                                                                                                                                                                                                                                                                                                                                                      |
| Custom Cache                | indices.queries.cache.size   | Maximum heap space allocated to the query cache.<br>Unit: percentage (%)<br>Value range: 1 to 100<br>Default value: 10%                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Queue Size in a Thread Pool | thread_pool.force_merge.size | Size of the thread pool used for force merge operations.<br>Value range: a positive integer<br>Default value: 1                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Custom                      | Custom parameters            | You can add any custom parameters supported by Elasticsearch.<br>The parameter values must meet the following requirements: <ul style="list-style-type: none"><li>• If a custom parameter has multiple values, use a comma (,) and a space to separate them. The input format is [value1, value1, value1...].</li><li>• Colons (:) are not allowed.</li></ul> <b>WARNING</b><br>Set the values of any custom parameters to those supported by OpenSearch. Otherwise, the cluster may fail to restart. Exercise caution when performing this operation. |

6. 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.
7. 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 not updated** will be displayed in the **Task Status** column in the cluster list.
  - If you restart the cluster after the modification, and **Task Status** displays **Configuration error** in the cluster list, the parameter configuration file has failed to be modified.

### 3.13.5 Associating an OpenSearch Cluster with an Enterprise Project

In CSS, you can associate your OpenSearch cluster with 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 these users and user groups.

You can associate your cluster with an enterprise project either when creating the cluster or afterward.

#### Prerequisites

To use the enterprise project function, you need to assign permissions to the corresponding account. You can [submit a service ticket](#) to apply for the permissions.

You have [created an enterprise project](#).

#### Managing Enterprise Projects


1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, find **Enterprise Project** in the **Configuration** area, and click the enterprise project name next to it. The **Project Management** page is displayed.

**Figure 3-54** Enterprise Project




5. On the **Resources** tab, select the region of the current cluster, and select **CloudSearchService** as **Service**. Then click **Search**. The CSS clusters are displayed in the resource list below.

Figure 3-55 Filtering CSS clusters




6. Select the cluster whose enterprise project you want to modify and click **Remove**.
7. In the **Remove Resource** dialog box, specify **Mode** and select **Destination Enterprise Project**. Then click **OK**.

Figure 3-56 Remove Resource



8. After the change is complete, return to the OpenSearch cluster list on the CSS console and check the cluster's new enterprise project.

Figure 3-57 Enterprise Project



### 3.13.6 Changing the Node Specifications of an OpenSearch Cluster

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



**Table 3-146** Change scenarios


| Change Type                  | Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Change Process                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Changing node specifications | <p>Typically, you increase node specifications instead of decreasing them. Common scenarios include:</p> <ul style="list-style-type: none"><li>• If the allocation of new indexes or shards takes too long or the coordination and scheduling of nodes are inefficient, increase the master node specifications.</li><li>• If too many requests need to be handled or too many results need to be aggregated, increase the client node specifications.</li><li>• If data nodes are becoming slower in responding to data writing and query requests, increase the specifications of the data nodes.</li><li>• If cold data query becomes slow, increase the cold data node specifications.</li><li>• If the CPU or JVM metrics of cluster nodes show performance bottlenecks, increase node specifications to improve cluster performance.</li></ul> <p>Alternatively, you may also decrease node specifications, but doing so will decrease the cluster's data processing and storage capacities. Exercise caution.</p> | <ol style="list-style-type: none"><li>1. Bring a node offline.</li><li>2. Change the node specifications.</li><li>3. Restart the node and restore its data.</li><li>4. After the node recovers, proceed to another node and repeat the steps above. This goes on until all nodes are changed.</li></ol> <p>The node specifications are changed one node at a time. This is to ensure that there are sufficient resources to keep services running.</p> |

| Change Type                                | Scenario                                                                                                                 | Change Process                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Changing the node storage type (disk type) | Change the node storage type if disk I/O has become a performance bottleneck, which impacts query and write performance. | <ol style="list-style-type: none"><li>1. Select a random node and migrate its data to other nodes.</li><li>2. Rebuild the node using the new storage type.</li><li>3. Add the node back into the cluster. The system automatically triggers a shard reallocation, migrating some of the shards to the new node.</li><li>4. After the node recovers, proceed to another node and repeat the steps above.</li></ol> <p>The nodes are changed one at a time to prevent service interruptions.</p> |

## Prerequisites

- The cluster status is **Available**, and there are no ongoing tasks.
- Your CSS resource quotas are sufficient for the capacity expansion you are about to perform. You can check available resources on the **Modify Configuration** page.

**Figure 3-58** Checking available resources



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

## Impact on Billing

For a pay-per-use cluster, you can see its new price when confirming the node specifications or storage type change on the console. After the change is complete, the new price will apply. For pricing details, see [Cloud Search Service Price Calculator](#).

For a yearly/monthly cluster, you can also see the new price when you confirm the change. In the case of an upgrade, a new order is triggered when you confirm the

change; and in the case of a downgrade, a refund request is triggered when you do so.

## Constraints

- The node specifications and storage type cannot be changed for nodes that use local disks.
- The node specifications and storage type cannot be changed at the same time.
- The node storage type can be changed only for data nodes and cold data nodes.
- When you change the node storage type, data needs to be migrated between different nodes. The timeout for data migration per node is 48 hours. Upgrade will fail if this timeout expires. When the cluster has large quantities of data, you are advised to manually adjust the data migration rate and avoid performing the migration during peak hours.
- For a cluster without master nodes, the node storage type can be changed only if the number of data nodes plus cold data nodes is at least three.
- For a cluster with master nodes, this operation is allowed only if the cluster has at least two data nodes.
- During a node storage type change, there is always one node that is unavailable. To ensure service continuity, make sure the **total number of data nodes and cold data nodes is greater than the maximum number of index replicas plus 1**. For a single-AZ or dual-AZ cluster, also make sure there are at least two nodes of each type for the cluster in each of its AZs.
- During a node specifications change, nodes are brought offline in order to make the changes. To ensure service continuity, make sure all shards have replicas.
- Make sure the disk usage is always less than 80% during the change.

## Change Impact

Before changing a cluster's node specifications or storage type, it is essential to assess the potential impacts and review operational recommendations. This enables proper scheduling of the change, minimizing service interruptions.

- Performance impact (only when the node storage type is changed)

Changing the node storage type does not interrupt services. However, data migration that occurs during this process consumes I/O performance, and taking individual nodes offline still has some impact on the overall cluster performance.

To minimize this impact, it is advisable to adjust the data migration rate based on the cluster's traffic cycle: increase the data migration rate during off-peak hours to shorten the task duration, and decrease it **before** peak hours arrive to ensure optimal cluster performance. The data migration rate is determined by the **indices.recovery.max\_bytes\_per\_sec** parameter. The default value of this parameter is the number of vCPUs multiplied by 32 MB. For example, for four vCPUs, the data migration rate is 128 MB. Set this parameter to a value between 40 MB and 1000 MB based on your service requirements.

```
PUT /_cluster/settings
{
 "transient": {
 "indices.recovery.max_bytes_per_sec": "1000MB"
 }
}
```

- Impact on request handling  
Taking nodes offline one at a time usually does not interrupt services. However, requests sent to offline nodes may fail. To mitigate this impact, the following measures may be taken:
  - Use a VPC endpoint or a dedicated load balancer to handle access requests to your cluster, which makes sure that requests are automatically routed to available nodes.
  - Enable an exponential backoff & retry mechanism on the client (configure three retries).
  - Perform this operation during off-peak hours.
- Impact on index replicas  
Shards that have no replicas will become unavailable when the nodes that store them are taken offline, causing service interruptions. You are advised to add replicas for all important indexes before making the change described in this topic.
- Impact on OpenSearch Dashboards and Cerebro  
Changing the node storage type for a cluster will cause OpenSearch Dashboards and Cerebro to be rebuilt. During this period, they are temporarily unavailable. During a node specifications change, if OpenSearch Dashboards and Cerebro become unavailable because the node that runs them is taken offline, refresh the web page or try to log in again, and the system will reschedule them to an available node.
- Characteristics of this process  
Once started, a change task cannot be stopped until it succeeds or fails. A change task failure only impacts a single node, and does not interrupt services if there are data replicas, but the failed node still needs to be restored promptly.

If a change of the node specifications or storage type is urgently needed, submit a [service ticket](#) to contact technical support and perform a pre-change evaluation and the necessary checks.

## Change Durations

- The following formulas can be used to estimate how long the node specifications change will take:  
**Change duration (min) = 10 (min) x Total number of nodes to change + Data recovery duration (min)**  
where,
  - 10 minutes indicates how long non-data recovery operations (e.g., initialization) typically take per node. It is an empirical value.
  - The total number of nodes is the sum of the number of data nodes, master nodes, client nodes, and cold data nodes in the cluster.**Data recovery duration (min) = Total data size (MB)/[Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s)]**

where,

- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.
- The formulas above use estimates under ideal conditions. The actual data recovery speed depends on cluster load.
- The following formulas can be used to estimate how long the node storage type change will take:

**Change duration (min) = 15 (min) x Total number of nodes to change + Data migration duration (min)**

where,

- 15 minutes indicates how long non-data migration operations (e.g., initialization) typically take per node. It is an empirical value.
- The total number of nodes is the sum of the number of data nodes, master nodes, client nodes, and cold data nodes in the cluster.

**Data migration duration (min) = Total data size (MB)/[Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s)]**

where,


- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.
- The formulas above use estimates under ideal conditions. The actual migration speed depends on cluster load.

## Changing the Node Specifications and Storage Type

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. Check that all service data has replicas so that services will not be interrupted during the specifications change.
  - a. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
  - b. In the left navigation pane, choose **Dev Tools**.
  - c. Run the **GET \_cat/indices?v** command in OpenSearch Dashboards.
    - If the returned **rep** value is greater than **0**, data replicas exist. Go to the next step.
    - If the returned **rep** value is **0**, there are no data replicas. You are advised to manually create a snapshot for the cluster before moving on to the next step. For details, see [Creating Snapshots to Back Up the Data of an OpenSearch Cluster](#).
4. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
5. Click the **Scale Cluster** tab.
6. Click **Change specifications** to set parameters.

**Table 3-147** Specifications change

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                   |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Action    | Select <b>Change specifications</b> .                                                                                                                                                                                                                                                                                                                                         |
| Resources | Shows the change of resources for this operation.                                                                                                                                                                                                                                                                                                                             |
| Nodes     | Configure the changes you want to make.<br><ol style="list-style-type: none"><li>1. Select a node type in the <b>Node Type</b> column.</li><li>2. Select a new flavor in the <b>Node Specifications</b> column, or select a new storage type in the <b>Node Storage</b> column.</li></ol> <p>The node specifications and storage type cannot be changed at the same time.</p> |

**Figure 3-59** Changing specifications

7. Click **Next**.
8. Confirm the information and click **Submit**.
9. In the displayed dialog box, confirm check items, and click **OK** to start the specifications change.
  - Check items for node specifications change: **Verify index copies** and **Cluster status check**.
  - Check items for node storage type change: **Check cluster load**.

**Table 3-148** Check item description

| Item                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Verify index copies | <p>By default, CSS checks for indexes that do not have any replicas created for them. You can skip this step, but the lack of index replicas may impact service availability during a node specifications change.</p> <ul style="list-style-type: none"><li>• If you select <b>Verify index copies</b> and the cluster has no master nodes, each index is required to have at least one replica and the cluster must have at least three nodes.</li><li>• If you select <b>Verify index copies</b> and the cluster has a master node, each index is required to have at least one replica, but there is no requirement on how many nodes the cluster must have.</li></ul> |

| Item                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster status check | By default, the cluster status is checked before specifications change. The specifications of nodes are changed one by one to ensure service availability and data security. If a cluster is overloaded and services are faulty, the request for a specifications change may fail to be delivered. In this case, you can skip cluster status check, but doing so may cause the cluster to become faulty and services become interrupted. Exercise caution when performing this operation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Check cluster load   | <p>During a node storage type change, data migration between nodes and the stopping and restarting of nodes will consume cluster resources, causing the cluster load to increase. A cluster load check can identify possible overload risks for a cluster and reduce the likelihood of an overload condition causing the node storage type change to fail.</p> <p>The cluster load check items are as follows:</p> <ul style="list-style-type: none"><li>• <b>nodes.thread_pool.search.queue &lt; 1000</b>: Check that the maximum number of requests in the search queue is less than 1000.</li><li>• <b>nodes.thread_pool.write.queue &lt; 200</b>: Check that the maximum number of requests in the write queue is less than 200.</li><li>• <b>nodes.process.cpu.percent &lt; 90</b>: Check that the maximum CPU usage is less than 90%.</li><li>• <b>nodes.os.cpu.load_average/Number of vCPUs &lt; 80%</b>: Check that the number of running processes plus the number of processes waiting for CPUs is less than 80% of the total number of vCPUs.</li></ul> |

 **NOTE**

If the change request fails to be submitted and a message is displayed indicating that the cluster needs to be upgraded, it means the current cluster version does not support a node storage type change. Upgrade the cluster to the latest image version and then try again. For a detailed upgrade guide, see [Upgrading the Version of an OpenSearch Cluster](#).

10. Click **Back to Cluster List** to go back to the **Clusters** page. The **Cluster Status** is **Configuration modified**. When **Cluster Status** changes to **Available**, the cluster node specifications have been successfully changed.

### 3.13.7 Replacing Specified Nodes for an OpenSearch Cluster

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

The node replacement process is as follows:

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

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

## Prerequisites

- The cluster status is **Available**, and there are no ongoing tasks.
- All mission-critical data has been backed up.

## Constraints

- Only one node can be replaced at a time. Each new node is rebuilt using the ID, IP address, specifications, and AZ of the node it is replacing.
- The configurations you modified manually will not be retained after node replacement. For example, if you have manually added a return route for the original node, you need to add it again for the new node after the node replacement is complete.
- If the node you want to replace is a data node or cold data node, pay attention to the following precautions:
  - When a data node or cold data node is replaced, its data is first migrated to other data nodes. This means the **total number of data nodes and cold data nodes must be greater than the maximum number of index replicas plus 1**.
  - In the AZ that contains the data node or cold data node to be replaced, there has to be at least another data node or cold data node.
  - If the cluster has not master nodes, the total number of data nodes plus cold data nodes must be at least three.
  - The precautions above do not apply if you are replacing a faulty node, regardless of its type. This is because faulty nodes are not included in `_cat/nodes`.

## Change Impact

Before replacing a node, it is essential to assess the potential impacts and review operational recommendations. This enables proper scheduling of the node replacement, minimizing service interruptions.

- Impact on performance  
Replacing a node does not interrupt services. However, data migration that occurs during this process consumes I/O performance, and taking individual nodes offline still has some impact on the overall cluster performance.  
To minimize this impact, it is advisable to adjust the data migration rate based on the cluster's traffic cycle: increase the data migration rate during off-peak hours to shorten the task duration, and decrease it **before** peak hours arrive to ensure optimal cluster performance. The data migration rate is



determined by the **indices.recovery.max\_bytes\_per\_sec** parameter. The default value of this parameter is the number of vCPUs multiplied by 32 MB. For example, for four vCPUs, the data migration rate is 128 MB. Set this parameter to a value between 40 MB and 1000 MB based on your service requirements.

PUT /\_cluster/settings

```
{
 "transient": {
 "indices.recovery.max_bytes_per_sec": "1000MB"
 }
}
```

- Impact on request handling

While a node is being replaced, requests sent to it may fail. To mitigate this impact, the following measures may be taken:

- Use a VPC endpoint or a dedicated load balancer to handle access requests to your cluster, which makes sure that requests are automatically routed to available nodes.
- Enable an exponential backoff & retry mechanism on the client (configure three retries).
- Perform this operation during off-peak hours.

- Characteristics of this process

Once started, a node replacement task cannot be stopped until it succeeds or fails. A task failure only impacts a single node, and does not interrupt services if there are data replicas, but the failed node still needs to be restored promptly.

## Node Replacement Duration

The following formula can be used to estimate how long it will take to replace a specified node of a cluster:

**Change duration (min) = 15 (min) + Data migration duration (min)**

where, 15 minutes indicates how long non-data migration operations (e.g., initialization) typically take per node. It is an empirical value.

**Data migration duration (min) = Total data size (MB)/[Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s)]**

where,

- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.
- The formulas above use estimates under ideal conditions. The actual migration speed depends on cluster load.

## Replacing a Specified Node

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.

4. On the **Modify Configuration** page, click the **Replace Node** tab.
5. On the **Replace Node** tab, set the parameters as needed.

**Table 3-149** Replacing a specified node

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>After a node is replaced, NICs need to be reattached to the new node. This means you need to have VPC permissions. Select an IAM agency to grant the current account the permission to access and use VPC.</p> <ul style="list-style-type: none"><li>• This parameter is available only when the new IAM plane is connected.</li><li>• If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-upgrade-agency</b>.</li><li>• If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>VPC Administrator</b> role or the <b>VPC FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/>"vpc:subnets:get",<br/>"vpc:ports:*"</li><li>• To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</li><li>• To use an IAM agency, the following minimum permissions are required:<br/>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</li></ul> |
| Node Type | Select the node you want to replace. You can expand a node type to check all the nodes under it.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

6. Click **Submit**. In the data migration confirmation dialog box, choose to migrate data, and click **OK**.

During data migration, the system migrates all data from the to-be-replaced node to the remaining nodes, and replace the node upon completion of the data migration. If the data on the to-be-replaced nodes has replicas on other

nodes, data migration can be skipped and the cluster change can be completed faster.

- Click **Back to Cluster List** to go back to the **Clusters** page. The **Task Status** is **Upgrading**. When **Cluster Status** changes to **Available**, the node has been successfully replaced.

### 3.13.8 Changing the Security Mode of an OpenSearch Cluster

With CSS, when an OpenSearch cluster's security needs change, you can change its security mode settings.

Configure the security mode based on the security needs of your cluster.

#### CAUTION

HTTPS access can be disabled only for OpenSearch 1.3.6 and 2.19.0 security-mode clusters. For other versions, HTTPS access is forcibly enabled and cannot be disabled.

**Table 3-150** Cluster security modes

| Cluster Type              | Key Settings                   | Description                                                                                                     | Applicable Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|--------------------------------|-----------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Non-security mode cluster | <b>Security Mode:</b> disabled | Access to such a cluster requires no user authentication, and data will be transmitted in plaintext using HTTP. | <p>Use when creating a cluster for internal testing or workloads that have a low security standard.</p> <ul style="list-style-type: none"><li>• Advantage: easy to access the cluster.</li><li>• Disadvantage: poor security, as anyone can access the cluster. When the security mode is disabled, public network access and Kibana public network access cannot be enabled.</li></ul> <p>Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.</p> |

| Cluster Type          | Key Settings                                                                                                                                          | Description                                                                                                                                                        | Applicable Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Security-mode cluster | Security-mode cluster + HTTP: <ul style="list-style-type: none"><li>• <b>Security Mode:</b> enabled</li><li>• <b>HTTPS Access:</b> disabled</li></ul> | Such a cluster requires user authentication. It supports access control and data encryption, and it uses HTTP to transmit data in plaintext.                       | Use to balance security and performance. <ul style="list-style-type: none"><li>• Advantage: User authentication improves cluster security. HTTP-based access ensures high performance of the cluster. User permissions can be configured to ensure proper isolation.</li><li>• Disadvantage: Public network access is not supported.</li></ul> Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.                                                  |
|                       | Security-mode cluster + HTTPS: <ul style="list-style-type: none"><li>• <b>Security Mode:</b> enabled</li><li>• <b>HTTPS Access:</b> enabled</li></ul> | Such a cluster requires user authentication. It supports access control and data encryption, and it uses HTTPS to encrypt communication and enhance data security. | Use when security takes precedence over performance and public network access is required. <ul style="list-style-type: none"><li>• Advantage: User authentication improves cluster security. HTTPS enhances cluster security by encrypting all communication over the public network. User permissions can be configured to ensure proper isolation.</li><li>• Disadvantage: When HTTPS is used, data encryption and decryption introduce computational overhead and impact the cluster's read and write performance.</li></ul> |

**Table 3-151** lists the various types of security mode changes supported for CSS clusters.

**Table 3-151** Security mode change scenarios

| Action                                                | Scenario                                                                                                                                                                                                                          | Change Process                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Switching from the non-security mode to security mode | Non-security mode → Security mode + HTTP: Change a cluster from non-security mode to security mode + HTTP.                                                                                                                        | <ol style="list-style-type: none"><li>1. Select a node and modify the OpenSearch, OpenSearch Dashboards, and Cerebro configuration files.</li><li>2. Restart the OpenSearch, OpenSearch Dashboards, and Cerebro processes and restore data.</li><li>3. After the node recovers, proceed to another node and repeat the steps above. This goes on until all nodes are changed.</li></ol> |
|                                                       | Non-security mode → Security mode + HTTPS: Change a cluster from non-security mode to security mode + HTTPS.                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                         |
| Switching from security mode to non-security mode     | Security mode → Non-security mode: <ul style="list-style-type: none"><li>• Change a cluster from security mode + HTTP to non-security mode.</li><li>• Change a cluster from security mode + HTTPS to non-security mode.</li></ul> |                                                                                                                                                                                                                                                                                                                                                                                         |
| Switching between HTTP and HTTPS in security mode     | HTTP → HTTPS: Change a cluster from security mode + HTTP to security mode + HTTPS.                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                                         |
|                                                       | HTTPS → HTTP: Change a cluster from security mode + HTTPS to security mode + HTTP.                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                                         |

## Prerequisites

- Before changing the security mode for a cluster, you are advised to back up its data.
- The cluster status is **Available**, and there are no ongoing tasks.
- Check whether load balancing is enabled for the cluster. If yes, disable load balancing first. Enable load balancing again after the security mode is changed. This prevents errors in accessing the cluster through the load balancer during the change.

## Change Impact

Before changing the security mode for a cluster, it is essential to assess the potential impacts and review operational recommendations. This enables proper scheduling of the change, minimizing service interruptions.

**Table 3-152** Impact overview

| Action                                    | Service Interruption | Authentication Mode         | Performance | Public Network Access  | Security Account    |
|-------------------------------------------|----------------------|-----------------------------|-------------|------------------------|---------------------|
| Non-security mode → Security mode + HTTP  | Yes                  | Authentication required     | N/A         | Disallowed             | N/A                 |
| Non-security mode → Security mode + HTTPS | Yes                  | Authentication required     | Degraded    | Allowed                | N/A                 |
| Security mode → Non-security mode         | Yes                  | Authentication not required | Enhanced    | Automatically disabled | Permanently deleted |
| HTTP → HTTPS                              | Yes                  | No change                   | Degraded    | Allowed                | N/A                 |
| HTTPS → HTTP                              | Yes                  | No change                   | Enhanced    | Automatically disabled | N/A                 |

Impact description:

- The security mode can be enabled or disabled for a cluster, and a security-mode cluster can use either HTTPS and HTTP.
- Service interruption: When a cluster's security mode is changed, the cluster automatically restarts, leading to temporary service unavailability.
- Authentication mode: Changing a cluster's security mode also changes its user authentication mechanism. The client authentication logic must be updated accordingly. If client adaptation cannot be completed in a timely manner, services may be interrupted.
- Performance impact: Given the same hardware configuration, switching from HTTP to HTTPS leads to a performance loss (such as throughput) of around 20% under high concurrency. Conversely, when a cluster switches from HTTPS to HTTP, cluster performance improves.

- Public network access: Only security-mode clusters that use HTTPS support public network access.
- Security account: When the security mode is disabled for a cluster, the system permanently deletes the cluster's security account.
- Tool impact: If the security mode is changed for a cluster that has an ongoing OpenSearch Dashboards session, an error message will be displayed on OpenSearch Dashboards. To rectify the error, clear the cache and launch OpenSearch Dashboards again.

Changing the security mode for a cluster changes its accessibility mode, possibly causing service interruptions. You should perform this operation before services are brought online or when service interruptions can be tolerated.

## Change Duration

The following formula can be used to estimate how long it will take to change the security mode for a cluster:

**Change duration (min) = 5 (min) x Total number of nodes to change + Data recovery duration (min)**

where,

- 5 minutes indicates how long non-data recovery operations (e.g., initialization) typically take per node. It is an empirical value.
- The total number of nodes is the sum of the number of data nodes, master nodes, client nodes, and cold data nodes in the cluster.

**Data recovery duration (min) = Total data size (MB)/[Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s)]**

where,

- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.
- The formulas above use estimates under ideal conditions. The actual data recovery speed depends on cluster load.

## Switching from the Non-Security Mode to Security Mode

You can change a non-security mode cluster to a security-mode cluster that uses HTTP or HTTPS. After a cluster's security mode is enabled, user 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 > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. Choose the **Change Security Mode** tab.
5. Enable the security mode. Enter and confirm the administrator password of the cluster.

**Figure 3-60** Switching from the non-security mode to security mode

The screenshot shows a configuration interface with the following elements:

- Security Mode:** A toggle switch is turned on (blue). To its right, text reads: "If you enable this option, you will need to verify your identity to access the cluster".
- Administrator Username:** A text input field containing the value "admin".
- Administrator Password:** A password input field with a masked character (•) and an eye icon to toggle visibility.
- Confirm Password:** A second password input field, also masked and with an eye icon.
- HTTPS Access:** A toggle switch is turned on (blue). To its right, text reads: "If you enable this option, communication will be encrypted when you access the cluster."

6. Enable or disable **HTTPS Access**.
  - If you enable **HTTPS Access**, HTTPS is used to encrypt cluster communication, and public network access can be enabled for the cluster.
  - If you disable **HTTPS Access**, HTTP is used, and public network access is disallowed for the cluster.
7. Click **Submit**, and confirm the information. The cluster list is displayed.

The **Task Status** of the cluster is **Changing security mode**. When the cluster status changes to **Available**, the security mode has been successfully changed.

## Switching from Security Mode to Non-Security Mode

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

---

### WARNING


- Clusters with the security mode disabled do not require user authentication, and HTTP is used to transmit data in plaintext. Make sure such a cluster are deployed in a secure environment, and do not expose the cluster's network interface to the public network.
- When a cluster is switched from the security mode to non-security mode, indexes used by the security-mode cluster will be deleted. Back up data before changing the security mode.
- If a public IP address has been assigned to a security-mode cluster, unassign it before changing the security mode.
- If OpenSearch Dashboards public network access is enabled for a security-mode cluster, 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 > OpenSearch**.
  3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.



4. Choose the **Change Security Mode** tab.
5. Disable the security mode.

**Figure 3-61** Switching from security mode to non-security mode



6. Click **Submit**. In the displayed dialog box, confirm the information. The cluster list page is displayed.

The **Task Status** of the cluster is **Changing security mode**. When the cluster status changes to **Available**, the security mode has been successfully changed.

## Switching between HTTP and HTTPS in Security Mode


You can change the protocol of a security cluster.

### WARNING

- If a public IP address has been assigned to a security-mode cluster, unassign it before changing from HTTPS to HTTP.
- If the OpenSearch cluster version is 2.x or later, HTTPS is enabled by default and cannot be disabled when the security mode is enabled. This means you cannot change from "Security Mode + HTTPS" to "Security Mode + HTTP".

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. Choose the **Change Security Mode** tab.
5. Enable or disable **HTTPS Access**.

**Figure 3-62** Configuring the protocol



- If you enable **HTTPS Access**, HTTPS is used to encrypt cluster communication, and you can enable public network access for the cluster.
- If you disable **HTTPS Access**, an alarm message is displayed. Click **OK** to disable the protocol.  
Cluster communication is no longer encrypted and public network access cannot be enabled.

6. Click **Submit**, and confirm the information. The cluster list is displayed.  
The **Task Status** of the cluster is **Changing security mode**. When the cluster status changes to **Available**, the security mode has been successfully changed.

### 3.13.9 Switching AZs for an OpenSearch Cluster

CSS supports cross-AZ deployment. By switching AZs for a cluster, you can make more resources available to the cluster or improve the cluster's availability.

You can **Add AZ** or **Migrate AZ**.

- **Add AZ**: Add one or two AZs for a single-AZ cluster, or add one AZ for a dual-AZ cluster to enhance cluster availability.
- **Migrate AZ**: Migrate data from the current AZ to another AZ that has sufficient resources for the cluster.

The AZ change process is as follows:

1. Migrate data from one node to other available nodes.
2. Rebuild the node for the target AZ. Keep the node ID, IP address, and specifications unchanged.
3. Add the new node into the cluster. The system automatically triggers a shard reallocation, migrating some of the shards to the new node. Then repeat these steps on all the other nodes.

The nodes are changed one at a time to prevent service interruptions.

#### Prerequisites

- The cluster status is **Available**, and there are no ongoing tasks.
- Make sure that the cluster has not had any non-standard modifications, such as modifying return routes and system parameters. Such modifications, if ever made, will be lost after the AZ change, which may affect your services.
- All mission-critical data has been backed up before an AZ switchover. This is to prevent data loss.

#### Constraints

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

- When you add nodes for a cluster (**Add AZ**), the current AZ must be retained. When adding one or two AZs for a single-AZ cluster, you must change AZs for all nodes. When adding one AZ for a dual-AZ cluster, you can choose to change AZs for specific types of nodes or all nodes in the cluster. For example, for a dual-AZ cluster, you can use three AZs for its master nodes, while still using two AZs for the rest of its nodes. To complete this AZ change, the system will try to move as few nodes as possible to rebuild the cluster. During this process, the YML configuration file of the nodes that are not modified is also updated. You need to restart the cluster to make the change take effect.
- When migrating an AZ (**Migrate AZ**), you can select only one target AZ. For this operation, you can choose to migrate only specific types of nodes or all nodes in the cluster. For example, for a dual-AZ cluster, you may move the master nodes from one AZ to the other, while keeping the rest of the nodes unchanged. If you migrate from one AZ to another, there is no need to restart the cluster. If multiple AZs are migrated, you must restart the cluster to make the change take effect.

## Change Impact

Before changing a cluster's AZs, it is essential to assess the potential impacts and review operational recommendations. This enables proper scheduling of the change, minimizing service interruptions.

- Performance

Changing AZs for a cluster does not interrupt services. However, data migration that occurs during this process consumes I/O performance, and taking individual nodes offline still has some impact on the overall cluster performance.

To minimize this impact, it is advisable to adjust the data migration rate based on the cluster's traffic cycle: increase the data migration rate during off-peak hours to shorten the task duration, and decrease it **before** peak hours arrive to ensure optimal cluster performance. The data migration rate is determined by the **indices.recovery.max\_bytes\_per\_sec** parameter. The default value of this parameter is the number of vCPUs multiplied by 32 MB. For example, for four vCPUs, the data migration rate is 128 MB. Set this parameter to a value between 40 MB and 1000 MB based on your service requirements.

```
PUT /_cluster/settings
```

```
{
 "transient": {
 "indices.recovery.max_bytes_per_sec": "1000MB"
 }
}
```

- Impact on request handling

While a node is offline, requests sent to it may fail. To mitigate this impact, the following measures may be taken:

- Use a VPC endpoint or a dedicated load balancer to handle access requests to your cluster, which makes sure that requests are automatically routed to available nodes.
- Enable an exponential backoff & retry mechanism on the client (configure three retries).
- Perform this operation during off-peak hours.

- Characteristics of this process  
Once started, an AZ change task cannot be stopped until it succeeds or fails.

## Change Duration

The following formula can be used to estimate how long it will take to change AZs for a cluster:

**Change duration (min) = 15 (min) x Total number of nodes involved + Data migration duration (min)**

where, 15 minutes indicates how long non-data migration operations (e.g., initialization) typically take per node. It is an empirical value.

**Data migration duration (min) = Total data size of the nodes to be migrated (MB) ÷ [Total number of vCPUs of the data nodes x 32 (MB/s) x 60 (s)]**

where,

- 32 MB/s indicates that each vCPU can process 32 MB of data per second. It is an empirical value.
- The formulas above use estimates under ideal conditions. The actual migration speed depends on cluster load.

## Changing AZs

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. Click the **Change AZ** tab.
5. On the **Change AZ** page, set parameters.

---

### CAUTION

If the number of data nodes plus cold data nodes in a cluster is not divisible by the number of AZs, data in the cluster may be unevenly distributed after the AZ change is completed. For example, if the cluster has two AZs and three data nodes, one node may need to store twice as much data as the other two nodes. This uneven distribution of data may affect the cluster's query or write performance.

---

**Table 3-153** Parameters for changing AZs

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Operation Type | <ul style="list-style-type: none"><li>● <b>Add AZ:</b> Add one or two AZs for a single-AZ cluster, or add one AZ for a dual-AZ cluster.<br/>To complete this AZ change, the system will try to move as few nodes as possible to rebuild the cluster. During this process, the YML configuration file of the nodes that are not modified is also updated. You need to restart the cluster to make the change take effect.</li><li>● <b>Migrate AZ:</b> Migrate data from one AZ to another.<br/>After adding AZs, you need to restart the cluster to make the modification take effect.</li></ul> |
| Node Type      | Select a node type for the AZ change.<br>Select one node type or <b>All nodes</b> to change their AZ. When adding one or two AZs for a single-AZ cluster, you can only select <b>All nodes</b> .                                                                                                                                                                                                                                                                                                                                                                                                 |
| Current AZ     | Current AZ of a cluster                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Target AZ      | AZ to add or migrate to. <ul style="list-style-type: none"><li>● <b>Add AZ:</b> Select up to three AZs, which must include all your current AZs.</li><li>● <b>Migrate AZ:</b> Select only one target AZ, which cannot be your current AZ.</li></ul>                                                                                                                                                                                                                                                                                                                                              |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>After a node is rebuilt, NICs need to be reattached to the new node. This means you need to have VPC permissions. Select an IAM agency to grant the current account the permission to access and use VPC.</p> <ul style="list-style-type: none"><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-upgrade-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>VPC Administrator</b> role or the <b>VPC FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<div>"vpc:subnets:get",<br/>"vpc:ports:*"</div></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<div>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</div></li><li>To use an IAM agency, the following minimum permissions are required:<div>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</div></li></ul> |

6. Click **Submit**. In the displayed dialog box, choose whether to check for the presence of a full index snapshot, and click **OK** to start the change.

 **NOTE**

You are advised to select **Check full index snapshot**. This ensures that all data has been backed up, so that in case the task fails, the data can be restored using this snapshot.

**Figure 3-63** Checking full index snapshots

**Check Full Index Snapshot**

You are advised to check the full index snapshot before the change.  
If you have backed up data, check whether the backup period exceeds one month. If yes, back up the latest data.  
Currently, only the index name can be matched, and the specific content and backup time cannot be verified.

☒ Check full index snapshot

Cancel

OK

When **Task Status** in the task list below changes to **Successful**, the AZ change is completed.

7. Confirm the result.
  - a. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  - b. On the **Overview** tab, check the cluster's AZs in the **Configuration** area to determine whether the switchover is successful.

**Figure 3-64** Checking the cluster's AZs

### Configuration

Region

AZ

AZ2

VPC

vpc-

## Checking the Progress of an AZ Change Task

In the **Jobs** area, find the AZ change task.

Expand the task, and click **View Progress** to check its progress.

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

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

If by the time an AZ change task is terminated, the AZ has not been changed for some nodes, you can try to restore these nodes by performing [Replacing Specified Nodes for an OpenSearch Cluster](#).

### CAUTION

If the AZ change is only partially successful, the cluster will have nodes spanning different AZs, and there is no convenient way to continue to change the AZ for the remaining nodes. To avoid this situation, do not terminate the task in this state.

## 3.13.10 Configuring Custom Word Dictionaries for an OpenSearch Cluster

### 3.13.10.1 Introduction to Word Dictionaries for OpenSearch Clusters

#### Context

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

CSS uses the IK and synonym analyzers. The IK analyzer uses a main word dictionary and a stop word dictionary. The synonym analyzer uses a synonym word dictionary. The IK analyzer uses the **ik\_max\_word** and **ik\_smart** word segmentation policies. The synonym analyzer uses the **ik\_synonym** word segmentation policy.

- **ik\_max\_word**: splits the text at a fine granularity.
- **ik\_smart**: splits the text at a coarse granularity.

CSS provides four preset word dictionaries: static main word dictionary, static stop word dictionary, extra main word dictionary, and extra stop word dictionary.

- If these preset word dictionaries can already support your application needs, there is no need to configure custom word dictionaries.
- If these preset word dictionaries are no longer sufficient, you can add a custom main word dictionary, stop word dictionary, or synonym dictionary, or modify the four preset word dictionaries to enable search by keywords or synonyms.

**Table 3-154** describes the different word dictionaries used by CSS.

**Table 3-154** Introduction to word dictionaries

| Word Dictionary Type |                      | Description                                                                                                                                         | File Requirements                                                                                                                                                                             |
|----------------------|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Custom               | Main word dictionary | Main words are words that users want to use as basis for performing word segmentation. A main word dictionary is a custom collection of such words. | The word dictionary file must be a text file encoded using UTF-8 without BOM, with one word per row. The maximum size of a main word dictionary file is 100 MB. Letters must be in lowercase. |
|                      | Stop word dictionary | Stop words are words that users wish to be ignored during word segmentation. The stop word dictionary is a custom collection of stop words.         | The word dictionary file must be a text file encoded using UTF-8 without BOM, with one word per row. The maximum size of a stop word dictionary file is 100 MB.                               |



| Word Dictionary Type |                             | Description                                                                                                                                                                                                                                                                                                                 | File Requirements                                                                                                                                                                                     |
|----------------------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                      | Synonym dictionary          | Synonyms are words with the same or similar meanings. A synonym dictionary is a custom collection of synonyms.                                                                                                                                                                                                              | The word dictionary file must be a text file encoded using UTF-8 without BOM. Each row contains a group of synonyms separated by commas (,). The maximum size of a synonym dictionary file is 100 MB. |
| Built-in             | Static main word dictionary | The static main word dictionary is a collection of common main words preconfigured in CSS. To view the static main word dictionary, visit <a href="https://github.com/infinilabs/analysis-ik/blob/master/config/main.dic">https://github.com/infinilabs/analysis-ik/blob/master/config/main.dic</a> .                       | The word dictionary file must be a text file encoded using UTF-8 without BOM, with one word per row. The maximum size of a static main word dictionary file is 100 MB.                                |
|                      | Static stop word dictionary | The static stop word dictionary is a collection of common stop words preconfigured in CSS. To view the static stop word dictionary, visit <a href="https://github.com/infinilabs/analysis-ik/blob/master/config/stopword.dic">https://github.com/infinilabs/analysis-ik/blob/master/config/stopword.dic</a> .               | The word dictionary file must be a text file encoded using UTF-8 without BOM, with one word per row. The maximum size of a static stop word dictionary file is 100 MB.                                |
|                      | Extra main word dictionary  | The Extra main word dictionary is a collection of uncommon main words preconfigured in CSS. To access the Extra main word dictionary, visit <a href="https://github.com/infinilabs/analysis-ik/blob/master/config/extra_main.dic">https://github.com/infinilabs/analysis-ik/blob/master/config/extra_main.dic</a> .         | The word dictionary file must be a text file encoded using UTF-8 without BOM, with one word per row. The maximum size of an Extra main word dictionary file is 100 MB.                                |
|                      | Extra stop word dictionary  | The Extra stop word dictionary is a collection of uncommon stop words preconfigured in CSS. To access the Extra stop word dictionary, visit <a href="https://github.com/infinilabs/analysis-ik/blob/master/config/extra_stopword.dic">https://github.com/infinilabs/analysis-ik/blob/master/config/extra_stopword.dic</a> . | The word dictionary file must be a text file encoded using UTF-8 without BOM, with one word per row. The maximum size of an Extra stop word dictionary file is 100 MB.                                |

## Constraints

- Custom word dictionaries are used for Chinese word segmentation. They can also be used to segment English words based on special characters other than `#&+-.@_`.
- The word dictionary files must meet the requirements described in [Table 3-154](#).
- Make sure the account you are using has the following permissions (choose **Permissions > Policies/Roles** on the IAM console to check your permissions):
  - **css:IKThesaurus:\***: the permission to configure custom word dictionaries.
  - Permission to view the OBS bucket list and read permissions on OBS buckets and the objects in them. For details, see [Permissions Management](#).

### 3.13.10.2 Configuring and Using Custom Word Dictionaries for an OpenSearch Cluster

#### Prerequisites

You have uploaded the custom word dictionary files to an OBS bucket. The OBS bucket must be in the same region as the target CSS cluster, and the storage class must be **Standard**.

#### Configuring Custom Word Dictionaries

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Cluster Settings > Custom Word Dictionaries**.
5. On the **Word Dictionaries** page, configure custom word dictionaries for the cluster or modify preset ones.
  - a. To configure custom word dictionaries, see [Table 3-155](#).

**Table 3-155** Configuring custom word dictionaries

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBS Bucket | Select the OBS location for storing the word dictionary file.<br><br>If no OBS buckets meet your requirements, click <b>Create Bucket</b> to go to the OBS console and create one. For details, see <a href="#">Creating a Bucket</a> .<br><br>To grant an IAM user access to an OBS bucket, you need to grant the <b>GetBucketStoragePolicy</b> , <b>GetBucketLocation</b> , <b>ListBucket</b> , and <b>ListAllMyBuckets</b> permissions to that user. |

| Parameter            | Description                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Main Word Dictionary | Main Word Dictionary is a custom word dictionary. Its initial state is empty. By default, <b>No Update</b> is selected, meaning not to configure this word dictionary. <ul style="list-style-type: none"><li>• To add a custom main word dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>• If you do not want to use a certain word dictionary, click <b>Do Not Use</b>.</li></ul> |
| Stop Word Dictionary | Stop Word Dictionary is a custom word dictionary. Its initial state is empty. By default, <b>No Update</b> is selected, meaning not to configure this word dictionary. <ul style="list-style-type: none"><li>• To add a custom stop word dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>• If you do not want to use a certain word dictionary, click <b>Do Not Use</b>.</li></ul> |
| Synonym Dictionary   | Synonym Dictionary is a custom word dictionary. Its initial state is empty. By default, <b>No Update</b> is selected, meaning not to configure this word dictionary. <ul style="list-style-type: none"><li>• To add a custom synonym dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>• If you do not want to use a certain word dictionary, click <b>Do Not Use</b>.</li></ul>     |

- b. To modify a preset word dictionary, toggle on **Modify Preset Word Dictionary**, and then modify that word dictionary as needed.

 **NOTE**

If the four preset word dictionaries (static main word, static stop word, extra main word, and extra stop word) are not displayed, the current cluster version does not support the deletion or modification of them. To use this function, you are advised to upgrade the cluster version, or create a new cluster and migrate data to it.

**Table 3-156** Configuring preset word dictionaries

| Parameter                   | Description                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Static Main Word Dictionary | Static Main Word Dictionary is a preset collection of common main words. By default, <b>No Update</b> is selected, meaning this word dictionary will be used. <ul style="list-style-type: none"><li>• To modify Static Main Word Dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>• If you do not want to use Static Main Word Dictionary, click <b>Do Not Use</b>.</li></ul> |

| Parameter                   | Description                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Static Stop Word Dictionary | Static Stop Word Dictionary is a preset collection of common stop words. By default, <b>No Update</b> is selected, meaning this word dictionary will be used. <ul style="list-style-type: none"><li>To modify Static Stop Word Dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>If you do not want to use Static Stop Word Dictionary, click <b>Do Not Use</b>.</li></ul> |
| Extra Main Word Dictionary  | Extra Main Word Dictionary is a preset collection of uncommon main words. By default, <b>No Update</b> is selected, meaning this word dictionary will be used. <ul style="list-style-type: none"><li>To modify Extra Main Word Dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>If you do not want to use Extra Main Word Dictionary, click <b>Do Not Use</b>.</li></ul>  |
| Extra Stop Word Dictionary  | Extra Stop Word Dictionary is a preset collection of uncommon stop words. By default, <b>No Update</b> is selected, meaning this word dictionary will be used. <ul style="list-style-type: none"><li>To modify Extra Stop Word Dictionary, click <b>Update</b> and select a .txt word dictionary file.</li><li>If you do not want to use Extra Stop Word Dictionary, click <b>Do Not Use</b>.</li></ul>  |

- Click **Save**. In the dialog box that is displayed, click **OK**. The word dictionary information is displayed in the lower part of the page. The word dictionary status is **Updating**. In approximately 1 minute, the word dictionaries are configured, and **Word Dictionary Status** changes to **Successful**.
- The deletion or update of the two static word dictionaries requires a cluster restart to take effect. The update of other word dictionaries happens dynamically, and there is no need for cluster restart. For details about how to restart a cluster, see [Restarting an OpenSearch Cluster](#).

## Example

Configure a custom word dictionary for the cluster, set main words, stop words, and synonyms. Search for the target text by keyword and synonym and check the search results.

**Step 1** Configure a custom word dictionary and check the word segmentation result. If there is no need for a custom word dictionary because the preset word dictionaries are already sufficient, skip this step.

- Prepare a word dictionary file (a text file encoded using UTF-8 without BOM) and upload it to the target OBS path.  
Set the main word dictionary file, stop word dictionary file, and synonym word dictionary file.

 **NOTE**

The built-in static stop word dictionary contains common stop words such as **are** and **the**. If the built-in stop word dictionary was never deleted or updated, you do not need to upload such stop words.

2. Configure the word dictionary by referring to [Configuring Custom Word Dictionaries](#).
3. After the word dictionary takes effect, return to the cluster list. Locate the target cluster and click **Dashboards** in the **Operation** column to access the cluster.
4. On the OpenSearch Dashboards page, click **Dev Tools** in the navigation tree on the left. The operation page is displayed.
5. Run the following commands to check the performance of the ik\_smart word segmentation policy and ik\_max\_word word segmentation policy.
  - Use the ik\_smart word segmentation policy to split a piece of target text.

Example code:

```
POST /_analyze
{
 "analyzer":"ik_smart",
 "text":"Text used for word segmentation"
}
```

After the operation is completed, view the word segmentation result.

```
{
 "tokens": [
 {
 "token": "word-1",
 "start_offset": 0,
 "end_offset": 4,
 "type": "CN_WORD",
 "position": 0
 },
 {
 "token": "word-2",
 "start_offset": 5,
 "end_offset": 8,
 "type": "CN_WORD",
 "position": 1
 }
]
}
```

- Use the ik\_max\_word word segmentation policy to split a piece of target text.

Example code:

```
POST /_analyze
{
 "analyzer":"ik_max_word",
 "text":"Text used for word segmentation"
}
```

After the operation is completed, view the word segmentation result.

```
{
 "tokens" : [
 {
 "token": "word-1",
 "start_offset" : 0,
 "end_offset" : 4,
 "type" : "CN_WORD",
 "position" : 0
 },
 {
```

```
{
 "token" : "word-3",
 "start_offset" : 0,
 "end_offset" : 2,
 "type" : "CN_WORD",
 "position" : 1
},
{
 "token" : "word-4",
 "start_offset" : 0,
 "end_offset" : 1,
 "type" : "CN_WORD",
 "position" : 2
},
{
 "token" : "word-5",
 "start_offset" : 1,
 "end_offset" : 3,
 "type" : "CN_WORD",
 "position" : 3
},
{
 "token" : "word-6",
 "start_offset" : 2,
 "end_offset" : 4,
 "type" : "CN_WORD",
 "position" : 4
},
{
 "token" : "word-7",
 "start_offset" : 3,
 "end_offset" : 4,
 "type" : "CN_WORD",
 "position" : 5
},
{
 "token" : "word-2",
 "start_offset" : 5,
 "end_offset" : 8,
 "type" : "CN_WORD",
 "position" : 6
},
{
 "token" : "word-8",
 "start_offset" : 5,
 "end_offset" : 7,
 "type" : "CN_WORD",
 "position" : 7
},
{
 "token" : "word-9",
 "start_offset" : 6,
 "end_offset" : 8,
 "type" : "CN_WORD",
 "position" : 8
},
{
 "token" : "word-10",
 "start_offset" : 7,
 "end_offset" : 8,
 "type" : "CN_WORD",
 "position" : 9
}
]
```

**Step 2** Create an index and configure a word segmentation policy. After data is imported, search for data by keyword.

The example code is as follows:

1. Create the **book** index and configure a word segmentation policy.  
In this example, both **analyzer** and **search\_analyzer** are set to **ik\_max\_word**. You can also use **ik\_smart**.

```
PUT /book
{
 "settings": {
 "number_of_shards": 2,
 "number_of_replicas": 1
 },
 "mappings": {
 "properties": {
 "content": {
 "type": "text",
 "analyzer": "ik_max_word",
 "search_analyzer": "ik_max_word"
 }
 }
 }
}
```

2. Import the text information to the **book** index.

```
PUT /book/_doc/1
{
 "content": "Imported text"
}
```

3. Use a keyword to search for the text and view the search results.

```
GET /book/_doc/_search
{
 "query": {
 "match": {
 "content": "Keyword"
 }
 }
}
```

#### Search result

```
{
 "took" : 16,
 "timed_out" : false,
 "_shards" : {
 "total" : 2,
 "successful" : 2,
 "skipped" : 0,
 "failed" : 0
 },
 "hits" : {
 "total" : {
 "value" : 1,
 "relation" : "eq"
 },
 "max_score" : 1.7260926,
 "hits" : [
 {
 "_index" : "book",
 "_type" : "_doc",
 "_id" : "1",
 "_score" : 1.7260926,
 "_source" : {
 "content" : "Imported text"
 }
 }
]
 }
}
```

**Step 3** Create an index and configure a synonym policy. After data is imported, search for data by synonyms.

The example code is as follows:

1. Create the **myindex** index and configure the word segmentation policy.

```
PUT myindex
{
 "settings": {
 "analysis": {
 "filter": {
 "my_synonym": {
 "type": "dynamic_synonym"
 }
 },
 "analyzer": {
 "ik_synonym": {
 "filter": [
 "my_synonym"
],
 "type": "custom",
 "tokenizer": "ik_smart"
 }
 }
 },
 "mappings": {
 "properties": {
 "desc": {
 "type": "text",
 "analyzer": "ik_synonym"
 }
 }
 }
 }
}
```

2. Import the text information to the **myindex** index.

```
PUT /myindex/_doc/1
{
 "desc": "Imported text"
}
```

3. Conduct search based on the synonym and view the search results.

```
GET /myindex/_search
{
 "query": {
 "match": {
 "desc": "Keyword"
 }
 }
}
```

#### Search result

```
{
 "took" : 1,
 "timed_out" : false,
 "_shards" : {
 "total" : 1,
 "successful" : 1,
 "skipped" : 0,
 "failed" : 0
 },
 "hits" : {
 "total" : {
 "value" : 1,
 "relation" : "eq"
 },
 "max_score" : 0.1519955,
 "hits" : [
 {
 "_index" : "myindex",
 "_type" : "_doc",
 "_id" : "1",
```



```
 "_score" : 0.1519955,
 "_source" : {
 "desc" : "Imported text"
 }
]
}
```

----End

### 3.13.11 Switching Between Simplified and Traditional Chinese for Data Search in an OpenSearch Cluster

This topic describes how to switch between simplified and traditional Chinese for data search in an OpenSearch cluster.

#### Scenario

The simplified-traditional Chinese conversion plugin converts between simplified and traditional Chinese. With this plugin, you can search index data containing the corresponding simplified Chinese based on the traditional Chinese keyword, and vice versa.

This plugin is installed by default. You do not need to install it by yourself.

The simplified-traditional Chinese conversion plugin can be used as the analyzer, tokenizer, token-filter, or char-filter.

The simplified-traditional Chinese conversion plugin provides the following two conversion types:

- s2t: converts simplified Chinese to traditional Chinese.
- t2s: converts traditional Chinese to simplified Chinese.

#### Switching Between Simplified and Traditional Chinese for Data Search

1. Log in to the OpenSearch Dashboards.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > OpenSearch**.
  - c. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
  - d. In the left navigation pane, choose **Dev Tools**.
2. Run the following command to create the **stconvert** index and specify the user-defined mapping to define the data type:

```
PUT /stconvert
{
 "settings": {
 "number_of_shards": 1,
 "number_of_replicas": 0,
 "analysis": {
 "analyzer": {
 "ts_ik": {
 "tokenizer": "ik_smart",
 "char_filter": [
 "tsconvert",
 "stconvert"
]
 }
 }
 }
 }
}
```

```
]
 }
},
"char_filter": {
 "tsconvert": {
 "type": "stconvert",
 "convert_type": "t2s"
 },
 "stconvert": {
 "type": "stconvert",
 "convert_type": "s2t"
 }
}
},
"mappings": {
 "properties": {
 "desc": {
 "type": "text",
 "analyzer": "ts_ik"
 }
 }
}
}
```

The command output is similar to the following:

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

3. Run the following command to import data to index **stconvert**:

```
POST /stconvert/_doc/1
{
 "desc": "Text in traditional Chinese"
}
```

If the value of **failed** in the command output is **0**, the data is imported successfully.

4. Run the following command to search for the keyword and check the search result:

```
GET /stconvert/_search
{
 "query": {
 "match": {
 "desc": "Keyword"
 }
 }
}
```

The search result is similar to the following:

```
{
 "took" : 15,
 "timed_out" : false,
 "_shards" : {
 "total" : 1,
 "successful" : 1,
 "skipped" : 0,
 "failed" : 0
 },
 "hits" : {
 "total" : 1,
 "max_score" : 0.5753642,
 "hits" : [
 {
 "_index" : "stconvert",
 "_type" : "type",
 "_id" : "1",
```

```
{
 "_score": 0.5753642,
 "_source": {
 "desc": "Text in traditional Chinese"
 }
}
```

### 3.13.12 Restarting an OpenSearch Cluster

If an OpenSearch cluster has stopped working, you can try to recover it by restarting it. Two restart options are available: quick restart and rolling restart. You are advised to perform the restart during off-peak hours.

#### Constraints

- Avoid restarting a cluster during peak hours, as doing so may impact data integrity and service availability.
- During a quick restart, the cluster is unavailable for the entire duration of the restart.
- Rolling restart is supported only for clusters with three or more nodes, including the master node, client node, and cold data nodes. During a rolling restart, only the nodes that are being restarted are unavailable. When the cluster stores a large amount of data, the restart may take a long time.

#### Prerequisites

- The cluster you want to restart is not frozen and has no task in progress.
- The cluster is available and is not processing any service requests, such as importing and searching for data. Otherwise, data loss may occur during a cluster restart.

#### Quick Restart

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Restart** in the **Operation** column.
4. On the **Restart Cluster** page, set parameters.

**Table 3-157** Configuring a quick restart

| Parameter    | Description                                                                                                                    |
|--------------|--------------------------------------------------------------------------------------------------------------------------------|
| Restart Mode | Select <b>Quick Restart</b> .<br><b>WARNING</b><br>The cluster will be unavailable for the entire duration of a quick restart. |

| Parameter | Description                                                                                                                                                                                                                                                                                                                       |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Select By | Select <b>Node type</b> or <b>Node name</b> . <ul style="list-style-type: none"><li>• <b>Node type</b>: You can select multiple node types to restart them at the same time.</li><li>• <b>Node name</b>: You can restart only one node at a time.</li></ul>                                                                       |
| Value     | Select one or more node types or a node name, depending on the setting of <b>Select By</b> .<br>Values for <b>Node type</b> : <ul style="list-style-type: none"><li>• <b>ess</b>: data node</li><li>• <b>ess-master</b>: master node</li><li>• <b>ess-client</b>: client node</li><li>• <b>ess-cold</b>: cold data node</li></ul> |

5. Click **OK** to start the restart task.

Refresh the page and check the cluster status. During the restart, the cluster status is **Processing**, and the task status is **Restarting**. When the cluster status changes to **Available**, the cluster has restarted successfully.

## Rolling Restart

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Restart** in the **Operation** column.
4. On the **Restart Cluster** page, set parameters.

**Table 3-158** Configuring a rolling restart

| Parameter    | Description                                                                                                                                                                                                                                                                                                                                                                             |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Restart Mode | Select <b>Rolling Restart</b> . <ul style="list-style-type: none"><li>• Rolling restart is supported only for clusters with three or more nodes, including the master node, client node, and cold data nodes.</li><li>• During a rolling restart, the nodes of a cluster are restarted one at a time. This may take a long time if the cluster has a large number of indexes.</li></ul> |
| Select By    | Only <b>Node type</b> is supported. You can select multiple node types to restart them at the same time.                                                                                                                                                                                                                                                                                |

| Parameter | Description                                                                                                                                                                                                                                                             |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Value     | Select node types for the restart.<br>Values for <b>Node type</b> : <ul style="list-style-type: none"><li>• <b>ess</b>: data node</li><li>• <b>ess-master</b>: master node</li><li>• <b>ess-client</b>: client node</li><li>• <b>ess-cold</b>: cold data node</li></ul> |

5. Click **OK** to start the restart task.

Refresh the page and check the cluster status. During the restart, the cluster status is **Processing**, and the task status is **Restarting**. When the cluster status changes to **Available**, the cluster has restarted successfully.

### 3.13.13 Deleting an OpenSearch Cluster

You can delete clusters that you no longer need.

#### Constraints

- Deleting a cluster will delete all its data. Any data that has not been backed up will be permanently lost. Exercise caution.
- A cluster that has no snapshots created for it cannot be restored once it is deleted. The snapshots of the cluster stored in OBS, if any, will not be deleted with the cluster. You can restore a deleted cluster using its snapshots stored in the OBS bucket. For details, see [Can I Restore a Deleted Cluster?](#)
- For a cluster for which the VPCEP service has been enabled, deleting the cluster may not delete the VPC endpoint associated with the cluster if the current account does not have the necessary VPCEP permissions. For details about the VPCEP permissions, see [Permissions](#).

#### Deleting a Pay-per-Use Cluster

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Delete** in the **Operation** column.
4. In the displayed dialog box, manually type in **DELETE**, and click **OK**.

#### Deleting a Yearly/Monthly Cluster

To delete a yearly/monthly cluster, unsubscribe from it. After unsubscription, the cluster's resources and data will be permanently deleted. If you want to retain data, manually back up data before submitting the unsubscription request.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Unsubscribe/Release** in the **Operation** column.

4. In the displayed dialog box, manually type in **YES**, and click **OK**.  
On the **Unsubscribe from Resource** page, confirm the resource information and refund amount.
5. Select a reason for unsubscription, select the acknowledgement check boxes, and click **Confirm**.

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

 **NOTE**

If the cluster is in **Available** state, an unsubscription order will be generated for refund, and the cluster will be deleted. If the cluster has expired or has been frozen, it will be deleted directly. For more information, see [Billing Center](#).

## 3.14 Managing Index Policies for OpenSearch Clusters

### 3.14.1 Creating and Managing OpenSearch Index Policies

This topic describes how to create and manage index lifecycle policies for OpenSearch clusters.

#### Context

Index State Management (ISM) of OpenSearch is a plugin that allows you to automate periodic, administrative operations on indexes by triggering them based on changes in the index age, index size, or number of documents. With ISM, you can define custom policies to automate index rollovers and deletion, thus optimizing cluster search performance or cutting storage costs. The procedure for using ISM is as follows:

1. [Creating an Index Lifecycle Policy](#): Create an index lifecycle policy on OpenSearch Dashboards.
2. [Associating Indexes with a Lifecycle Policy](#): Associate indexes with a lifecycle policy.
3. [Managing Index Policies](#): Modify, retry, and change index lifecycle policies.

For more information about ISM, see [Index State Management](#).

The OpenSearch Dashboards GUI may vary depending on the software version. OpenSearch Dashboards 1.3.6 is used as an example here.

#### Creating an Index Lifecycle Policy

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
4. Expand the OpenSearch Dashboards menu in the upper-left corner, and choose **Index Management**.
5. Click **Create policy** to create an index policy. In the **Configuration method** dialog box, select **JSON editor** and click **Continue**. The page for creating an index policy is displayed.

- **Policy ID:** use a custom policy name.
- **Define policy:** define a custom policy based on the reference example provided.

**Figure 3-65** Configuring a policy

Create policy

Name policy

Policy ID: hot\_cold\_workflow

Define policy

```
{
 "policy": {
 "description": "A simple default policy that changes the replica count between hot and cold states.",
 "default_state": "hot",
 "states": [
 {
 "name": "hot",
 "actions": [
 {
 "replica_count": {
 "number_of_replicas": 5
 }
 }
],
 "transitions": [
 {
 "state_name": "cold",
 "conditions": [
 {
 "min_index_age": "384d"
 }
]
 }
]
 },
 {
 "name": "cold",
 "actions": [
 {
 "replica_count": {
 "number_of_replicas": 2
 }
 }
]
 }
]
 }
}
```

In the following example, a policy named **policy1** is created. The **ism\_template** part specifies the indexes to apply the policy to. After the policy is created, indexes whose names start with **test** are automatically associated with this policy.


```
PUT _plugins/_ism/policies/policy1
{
 "policy": {
 "description": "an example policy",
 "default_state": "hot",
 "states": [
 // Ignore the states part.
],
 "ism_template": {
 "index_patterns": [
 "test*"
],
 "priority": 100
 }
 }
}
```

6. Click **Create**.

## Associating Indexes with a Lifecycle Policy


You can bind a policy to one or more indexes, or associate a policy with an index template, so that the policy will be applied to all indexes created using this index template. Use this method when associating existing indexes with a lifecycle policy.

1. On the **Index Management** page of the OpenSearch Dashboards, choose **Indexes**.
2. In the **Indexes** list, select one or more indexes that you want to bind with a policy.
3. Click **Apply policy** in the upper-right corner. In the displayed dialog box, select **Policy ID**.

**Figure 3-66** Apply policy

The screenshot shows the 'Indices' page in the OpenSearch Dashboards. At the top, there is a search bar and a toggle for 'Show data stream indices' (set to Off). Below this is a table with columns: Index, Health, Managed by Policy, Status, Total size, Primary size, Total documents, Deleted documents, Primaries, and Replicas. Three indices are listed: 'opensearch\_dashboards\_sample\_data\_ecommerce', '.opensearch-observability', and 'kibana'. All have a 'green' health status and are not managed by a policy. An 'Apply policy' button is in the top right corner.

| Index                                       | Health | Managed by Policy | Status | Total size | Primary size | Total documents | Deleted documents | Primaries | Replicas |
|---------------------------------------------|--------|-------------------|--------|------------|--------------|-----------------|-------------------|-----------|----------|
| opensearch_dashboards_sample_data_ecommerce | green  | No                | open   | 4.4mb      | 4.4mb        | 4675            | 0                 | 1         | 0        |
| .opensearch-observability                   | green  | No                | open   | 208b       | 208b         | 0               | 0                 | 1         | 0        |
| kibana                                      | green  | No                | open   | 75.1kb     | 75.1kb       | 40              | 0                 | 1         | 0        |

**Figure 3-67** Selecting a policy

The 'Apply policy' dialog box is shown. It contains a text input field labeled 'Policy ID' with a search icon and a dropdown arrow. The field contains the text 'new01'. Below the field are 'Cancel' and 'Apply' buttons.

4. Click **Apply**.

After you bind a policy to an index, ISM creates a job that runs every 5 minutes by default to execute the policy, check criteria, and change index states.

## Managing Index Policies

1. On the **Index Management** page of the OpenSearch Dashboards, choose **Managed Indices**.  
The displayed page shows index policies configured for the current cluster.
2. Manage policies in the index policy list.
  - For a policy whose status is abnormal, click **Retry policy**.
  - To update a policy for an index, select the associated index, and click **Change policy**. After the policy is updated, click **Change** to apply the new policy.
  - To remove a policy from an index, select the index, and click **Remove policy**. In the displayed dialog box, click **Remove**.

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

### 3.14.2 Automating Index Rollover in an OpenSearch Cluster Through Index Lifecycle Management

As time series data is continuously written into OpenSearch and the index size keeps growing, you can configure Index Lifecycle Management (ISM) to periodically roll over data to new indexes and delete old indexes.


In this example, a lifecycle policy is created for an OpenSearch cluster. According to this policy, when the size of an index reaches 1 TB or the index is one day old, a new index will be automatically generated; seven days after the index is created,



data duplicates will be disabled; 30 days after the index is created, the index will be deleted.

Assume that an index generates approximately 2.4 TB of data each day. The index alias is **log-alias**. The following figure illustrates the data format in the cluster. During reads, it points to all indexes whose name starts with **test**. During writes, it points to the latest index.

**Figure 3-68** log-alias format



#### NOTE

The one day in the rollover time refers to 24 hours following the index creation time, not a calendar day.

## Prerequisites

The target OpenSearch cluster is available.

## Automating Index Rollover Through Index Lifecycle Management

1. Log in to the OpenSearch Dashboards.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > OpenSearch**.
  - c. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
  - d. In the left navigation pane, choose **Dev Tools**.
2. Create a rollover lifecycle policy named **rollover\_workflow**.

Policy description: When the size of an index reaches 1 TB or its age exceeds one day, index rollover is performed automatically. When the index age reaches seven days, data replicas are disabled. When it reaches 30 days, the index is deleted.

```
PUT _plugins/_ism/policies/rollover_workflow
{
 "policy": {
 "description": "rollover test",
 "default_state": "hot",
 "states": [
 {
 "name": "hot",
 "actions": [
 {
 "rollover": {
 "min_size": "1tb",
 "min_index_age": "1d"
 }
 }
]
 }
]
 }
}
```

```
 },
 "transitions": [
 {
 "state_name": "warm",
 "conditions": {
 "min_index_age": "7d"
 }
 }
]
 },
 {
 "name": "warm",
 "actions": [
 {
 "replica_count": {
 "number_of_replicas": 0
 }
 }
],
 "transitions": [
 {
 "state_name": "delete",
 "conditions": {
 "min_index_age": "30d"
 }
 }
]
 },
 {
 "name": "delete",
 "actions": [
 {
 "delete": {}
 }
]
 }
],
"ism_template": {
 "index_patterns": [
 "test*"
],
 "priority": 100
}
}
```

After a lifecycle policy is created, run the following command to query the policy details:

```
GET _plugins/_ism/policies/rollover_workflow
```

### 3. Create the index template **template\_test**.

Template description: All the new indexes starting with **test** are automatically associated with the rollover lifecycle policy **rollover\_workflow**. The alias **log\_alias** is used during rollover.

```
PUT _template/template_test
```

```
{
 "index_patterns": "test*",
 "settings": {
 "number_of_replicas": 1,
 "number_of_shards": 1,
 "index.opendistro.index_state_management.rollover_alias": "log_alias"
 },
 "mappings": {
 "properties": {
 "name": {
 "type": "text"
 }
 }
 }
}
```

```
}
}
}
```

**Table 3-159** Parameter description

| Parameter                                              | Description                    |
|--------------------------------------------------------|--------------------------------|
| number_of_shards                                       | Number of index shards         |
| number_of_replicas                                     | Number of index shard replicas |
| index.opendistro.index_state_management.rollover_alias | Index alias for rollover       |

After an index template is created, you can run the following command to query the template details:

```
GET _template/template_test
```

4. Create an index, specify **aliases**, and set **is\_write\_index** to **true**. The index template **template\_test** is automatically used for the index and is associated with the lifecycle policy **rollover\_workflow** based on the index template configuration. In this way, when the index size reaches 1 TB or the index is created for more than one day, the rollover automatically starts. After an index is created for seven days, the data copy is disabled. After an index is created for 30 days, the index is deleted.

The following index is the URL code of **<test-{now/d}-000001>**. By default, an index name contains the creation date. For example, if an index is created on 2022-06-02, the index name is **test-2022.06.02-000001**.

```
PUT %3Ctest-%7Bnow%2Fd%7D-000001%3E
```

```
{
 "aliases": {
 "log_alias": {
 "is_write_index": true
 }
 }
}
```

5. The alias **log\_alias** is used to during data write, and **log\_alias** always points to the last index.

```
POST log_alias/_bulk
```

```
{"index":{}}
{"name":"name1"}
{"index":{}}
{"name":"name2"}
{"index":{}}
{"name":"name3"}
{"index":{}}
{"name":"name4"}
{"index":{}}
{"name":"name5"}
{"index":{}}
{"name":"name6"}
```

6. Query data and check whether the rollover has taken effect.
  - One day after the indexes are created, check the indexes starting with **test**.

```
GET _cat/indices/test*?s=i
```

There are supposed to be at least two indexes, for example:

```
green open test-<Date>-000001 r8ab5NX6T3Ox_hoGUanogQ 1 1 6 0 416b 208b
green open test-<Date>-000002 sfwkVgy8RSSEw7W-xYjM2Q 1 1 0 0 209b 209b
```

In the preceding information, **test-<Date>-000001** is the index created in 4, and **test-<Date>-000002** is the index generated through rollover.

- To query the index associated with the alias **log\_alias**, run the following command:

```
GET _cat/aliases/log_alias?v
```

The alias is supposed to point to multiple indexes, for example:

| alias     | index              | filter | routing.index | routing.search | is_write_index |
|-----------|--------------------|--------|---------------|----------------|----------------|
| log_alias | test-<Date>-000001 | -      | -             | -              | false          |
| log_alias | test-<Date>-000002 | -      | -             | -              | true           |

## FAQ: How Do I Skip Rollover for an Index?

### Scenarios

- A manual rollover is performed: A time-based ISM policy was configured to automate index rollover. Before the rollover criteria were met, however, a manual rollover was performed. When the auto rollover time configured by ISM was reached, an automatic rollover was attempted but failed because the index had already been rolled over manually. As a result, the ISM task stopped. To solve this problem, you need to skip the index rollover and then retry the ISM task to resume automatic index rollover.
- For an index that has already been rolled over automatically, a remove or add policy operation is performed: After the rollover policy is modified (through the remove policy or add policy operation), the ISM task restarts from the beginning. When rollover starts, the task fails. To solve this, you also need to skip index rollover and then retry the ISM task to resume automatic index rollover.



### WARNING

After rollover is skipped for an index, ISM will no longer attempt rollover or generate rolled-over indexes. This means skipping index rollover may lead to data loss. Please exercise caution.

### Procedure

1. If an ISM task stops due to an index rollover failure, run the following command to skip index rollover:

```
PUT index_name/_settings
{
 "index.plugins.index_state_management.rollover_skip": true
}
```

If **true** is returned, the configuration is successful.

2. After index rollover is skipped, run the following command to retry the ISM task:

```
POST _plugins/_ism/retry/index_name
```

If the following information is returned, the retry is successful:

```
{
 "updated_indices": 1,
 "failures": false,
```

```
"failed_indices": []
}
```


### 3.14.3 Decoupling Index Storage and Compute in an OpenSearch Cluster Through Index Lifecycle Management

#### Overview

CSS supports decoupled storage and compute. That is, indexes can be frozen in OBS to reduce the storage cost of cold data. This document describes how to use index lifecycle management to automatically freeze indexes at a specific time to decouple storage and compute.

In this section, a lifecycle policy is configured to automatically freeze an index three days after it is created and dump data to OBS. The index will be deleted seven days after it is created.

**Figure 3-69** Storage-compute decoupling



#### Prerequisites

- There are available CSS clusters.
- OpenSearch 1.3.6 or 2.19.0 is used.

#### Decoupling Index Storage and Compute Through Index Lifecycle Management

1. Log in to the OpenSearch Dashboards.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > OpenSearch**.
  - c. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
  - d. In the left navigation pane, choose **Dev Tools**.
2. Create a lifecycle policy named **hot\_warm\_policy**.

Policy description: Three days after an index is created, the API for freezing indexes is automatically called to dump data to OBS. Seven days after index creation, the index is deleted.

```
PUT _plugins/_ism/policies/hot_warm_policy
{
 "policy": {
 "description": "hot warm delete workflow",
 "error_notification": null,
 "default_state": "hot",
 "states": [

```

```
{
 "name": "hot",
 "actions": [],
 "transitions": [
 {
 "state_name": "warm",
 "conditions": {
 "min_index_age": "3d"
 }
 }
]
},
{
 "name": "warm",
 "actions": [
 {
 "freeze_low_cost": {}
 }
],
 "transitions": [
 {
 "state_name": "delete",
 "conditions": {
 "min_index_age": "7d"
 }
 }
]
},
{
 "name": "delete",
 "actions": [
 {
 "delete": {}
 }
],
 "transitions": []
}
],
"ism_template": {
 "index_patterns": [
 "data*"
],
 "priority": 100
}
}
```

3. Create an index named data-2022-06-06. The index is automatically associated with the lifecycle policy **hot\_warm\_policy**, which freezes the index three days after it is created and deletes it seven days after creation.

POST data-2022-06-06/\_bulk

```
{"index":{}}
{"name":"name1"}
{"index":{}}
{"name":"name2"}
{"index":{}}
{"name":"name3"}
{"index":{}}
{"name":"name4"}
{"index":{}}
{"name":"name5"}
{"index":{}}
{"name":"name6"}
```

4. Query data and check whether storage and compute is automatically decoupled.

- Three days after the index is created, check the frozen index.

GET \_cat/freeze\_indices?s=i&v

The index generated three days ago is expected to be frozen.

| health | status | index           | uuid                   | pri | rep | docs.count | docs.deleted | store.size |
|--------|--------|-----------------|------------------------|-----|-----|------------|--------------|------------|
| green  | open   | data-2022-06-06 | x8ab5NX6T3Ox_xoGUanogQ | 1   | 1   | 6          | 0            |            |
| 7.6kb  | 3.8kb  |                 |                        |     |     |            |              |            |

- Seven days after the index is created, check the frozen index. The index is expected to be deleted.

## 3.15 OpenSearch Cluster Monitoring and Log Management

### 3.15.1 Monitoring Metrics for OpenSearch Clusters in Cloud Eye

You can use Cloud Eye to monitor CSS cluster metrics in real time and quickly handle exceptions.

#### Overview

You can use Cloud Eye to monitor CSS cluster metrics in real time and quickly handle exceptions. This topic describes CSS cluster and node metrics that can be monitored through Cloud Eye, as well as their namespaces and dimensions. You can query metrics and alerts generated by CSS on the [Cloud Eye console](#) or an [API](#)

#### Namespaces

SYS.ES

#### Cluster Monitoring Metrics

##### NOTE

Accumulated value: The value is accumulated from when a node is started. After the node is restarted, the value is reset to zero and starts accumulating again.

Table 3-160 CSS cluster metrics

| Metric ID | Metric                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                    | Value Range | Unit | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------|---------------|------------------------------|---------------------------|
| status    | Cluster Health Status | Health status of the monitored object <ul style="list-style-type: none"><li>0: The cluster is 100% available.</li><li>1: All primary shards are allocated, but some replicas are 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.</li></ul> | 0, 1, 2, 3  | N/A  | N/A           | CSS cluster                  | 1 minute                  |



| Metric ID              | Metric                     | Description                                                                                                                                                                | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|------------------------|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
|                        |                            | <ul style="list-style-type: none"><li>• <b>2:</b> Data is missing and the cluster cannot work properly.</li><li>• <b>3:</b> Failed to obtain the cluster status.</li></ul> |             |       |               |                              |                           |
| disk_util              | Max. Disk Usage            | Maximum disk usage of nodes in a CSS cluster                                                                                                                               | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| max_jvm_heap_usage     | Max. JVM Heap Usage        | Maximum JVM heap usage of nodes in a CSS cluster.                                                                                                                          | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| max_jvm_young_gc_time  | Max. JVM Young GC Duration | Maximum accumulated JVM Young GC duration of nodes in a CSS cluster.                                                                                                       | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |
| max_jvm_young_gc_count | Max. JVM Young GC Count    | Maximum accumulated JVM Young GC count of nodes in a CSS cluster.                                                                                                          | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_jvm_old_gc_time    | Max. JVM Old GC Duration   | Maximum accumulated JVM Old GC duration of nodes in a CSS cluster.                                                                                                         | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                              | Metric                                  | Description                                                                            | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|----------------------------------------|-----------------------------------------|----------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| max_jvm_old_gc_count                   | Max. JVM Old GC Count                   | Maximum accumulated JVM Old GC count of nodes in a CSS cluster.                        | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| total_fs_size                          | Total Size of File Systems              | Total size of file systems in a CSS cluster.                                           | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| free_fs_size                           | Available Size of File Systems          | Available size of file systems in a CSS cluster.                                       | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| max_cpu_usage                          | Max. CPU Usage                          | Maximum node CPU usage in a CSS cluster.                                               | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| 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.         | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |
| 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.                    | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| max_current_opened_http_count          | Current Max. Opened HTTP Connections    | Maximum number of HTTP connections that are currently open for nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                   | Metric                             | Description                                                                   | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------|------------------------------------|-------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| max_total_opened_http_count | Total Max. Opened HTTP Connections | Maximum number of HTTP connections that were open for nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| indices_count               | Indexes                            | Number of indexes in a CSS cluster                                            | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| total_shards_count          | Shards                             | Number of shards in a CSS cluster                                             | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| primary_shards_count        | Primary Shards                     | Number of primary shards in a CSS cluster                                     | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| docs_count                  | Documents                          | Number of documents in a CSS cluster                                          | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| docs_deleted_count          | Deleted Documents                  | Number of documents deleted in a CSS cluster                                  | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| nodes_count                 | Nodes                              | Number of nodes in a CSS cluster                                              | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| data_nodes_count            | Data Nodes                         | Number of data nodes in a CSS cluster                                         | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| coordinating_nodes_count    | Coordinator Nodes                  | Number of coordinator nodes in a CSS cluster                                  | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                 | Metric                     | Description                                                                       | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|---------------------------|----------------------------|-----------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| master_nodes_count        | Master Nodes               | Number of master nodes in a CSS cluster                                           | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| ingest_nodes_count        | Client Nodes               | Number of client nodes in a CSS cluster                                           | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_load_average          | Max. Node Load             | Maximum number of average queued tasks per minute on nodes in a cluster           | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_cpu_usage             | Avg. CPU Usage             | Average node CPU usage in a CSS cluster.                                          | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| avg_load_average          | Avg. Node Load             | Average number of queued tasks per minute on nodes in a CSS cluster               | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_jvm_heap_usage        | Avg. JVM Heap Usage        | Average node JVM heap usage in a CSS cluster                                      | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| max_open_file_descriptors | Max. Open File Descriptors | Maximum number of node file descriptors that are currently open in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                         | Metric                        | Description                                                                               | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------------|-------------------------------|-------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_open_file_descriptors         | Avg. Open File Descriptors    | Average number of node file descriptors that are currently open in a CSS cluster.         | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_max_file_descriptors          | Max. Allowed File Descriptors | Maximum number of allowed node file descriptors in a CSS cluster.                         | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_open_file_descriptors         | Open File Descriptors         | Number of node file descriptors that are currently open in a cluster.                     | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_write_queue       | Tasks in Write Queue          | Number of tasks in queue for the write thread pool.                                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_search_queue      | Tasks in Search Queue         | Total number of tasks in queue for the search thread pool on nodes in a CSS cluster.      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_force_merge_queue | Tasks in ForceMerge Queue     | Total number of tasks in queue for the force merge thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                            | Metric                             | Description                                                                                 | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|--------------------------------------|------------------------------------|---------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_write_rejected       | Rejected Tasks in Write Queue      | Total number of tasks rejected by the write thread pool on nodes in a CSS cluster.          | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_search_rejected      | Rejected Tasks in Search Queue     | Total number of tasks rejected by the search thread pool on nodes in a CSS cluster.         | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_force_merge_rejected | Rejected Tasks in ForceMerge Queue | Total number of tasks rejected by the force merge thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_search_queue         | Max. Tasks in Search Queue         | Maximum number of tasks in queue for the search thread pool on nodes in a CSS cluster.      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_force_merge_queue    | Max. Tasks in ForceMerge Queue     | Maximum number of tasks in queue for the force merge thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                           | Metric                         | Description                                                                            | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------------|--------------------------------|----------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_write_threads       | Size of Write Thread Pool      | Total size of the write thread pools on nodes in the CSS cluster.                      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_search_threads      | Size of Search Thread Pool     | Total size of the search thread pools on nodes in the CSS cluster.                     | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_force_merge_threads | Size of ForceMerge Thread Pool | Total size of the force merge thread pools on nodes in the CSS cluster.                | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_write_queue         | Avg. Tasks in Write Queue      | Average number of tasks in queue for the write thread pool on nodes in a CSS cluster.  | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_search_queue        | Avg. Tasks in Search Queue     | Average number of tasks in queue for the search thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                           | Metric                              | Description                                                                                 | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------------|-------------------------------------|---------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_thread_pool_force_merge_queue   | Avg. Tasks in ForceMerge Queue      | Average number of tasks in queue for the force merge thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_search_threads      | Avg. Size of Search Thread Pool     | Average size of the search thread pool on nodes in a CSS cluster.                           | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_write_threads       | Avg. Size of Write Thread Pool      | Average size of the write thread pool on nodes in a CSS cluster.                            | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_force_merge_threads | Avg. Size of ForceMerge Thread Pool | Average size of the force merge thread pool on nodes in a CSS cluster.                      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_write_rejected      | Avg. Rejected Tasks in Write Queue  | Average number of tasks rejected by the write thread pool on nodes in a CSS cluster.        | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_flush_queue         | Tasks in Flush Queue                | Total number of tasks in queue for the flush thread pool on nodes in a CSS cluster.         | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |



| Metric ID                      | Metric                         | Description                                                                           | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|--------------------------------|--------------------------------|---------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_flush_rejected | Rejected Tasks in Flush Queue  | Total number of tasks rejected by the flush thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_flush_queue    | Max. Tasks in Flush Queue      | Maximum number of tasks in queue for the flush thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_flush_threads  | Size of Flush Thread Pool      | Total size of the flush thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_flush_queue    | Avg. Tasks in Flush Queue      | Average number of tasks in queue for the flush thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_flush_threads  | Avg. Size of Flush Thread Pool | Average number of tasks in queue for the flush thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                        | Metric                          | Description                                                                             | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|----------------------------------|---------------------------------|-----------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_generic_queue    | Tasks in Generic Queue          | Total number of tasks in queue for the generic thread pool on nodes in a CSS cluster.   | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_generic_rejected | Rejected Tasks in Generic Queue | Total number of tasks rejected by the generic thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_generic_queue    | Max. Tasks in Generic Queue     | Maximum number of tasks in queue for the generic thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_generic_threads  | Size of Generic Thread Pool     | Total size of the generic thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_generic_queue    | Avg. Tasks in Generic Queue     | Average number of tasks in queue for the generic thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                           | Metric                             | Description                                                                                | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------------|------------------------------------|--------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_thread_pool_generic_threads     | Avg. Size of Generic Thread Pool   | Average number of tasks in queue for the generic thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_management_queue    | Tasks in Management Queue          | Total number of tasks in queue for the management thread pool on nodes in a CSS cluster.   | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_management_rejected | Rejected Tasks in Management Queue | Total number of tasks rejected by the management thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_management_queue    | Max. Tasks in Management Queue     | Maximum number of tasks in queue for the management thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_management_threads  | Size of Management Thread Pool     | Total size of the management thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                          | Metric                              | Description                                                                                | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|------------------------------------|-------------------------------------|--------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_thread_pool_management_queue   | Avg. Tasks in Management Queue      | Average number of tasks in queue for the management thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_management_threads | Avg. Size of Management Thread Pool | Average number of tasks in queue for the management thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_refresh_queue      | Tasks in Refresh Queue              | Total number of tasks in queue for the refresh thread pool on nodes in a CSS cluster.      | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_refresh_rejected   | Rejected Tasks in Refresh Queue     | Total number of tasks rejected by the refresh thread pool on nodes in a CSS cluster.       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_refresh_queue      | Max. Tasks in Refresh Queue         | Maximum number of tasks in queue for the refresh thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                             | Metric                               | Description                                                                                | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|---------------------------------------|--------------------------------------|--------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_refresh_threads       | Size of Refresh Thread Pool          | Total size of the refresh thread pools on nodes in a CSS cluster.                          | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_refresh_queue         | Avg. Tasks in Refresh Queue          | Average number of tasks in queue for the refresh thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_refresh_threads       | Avg. Size of Refresh Thread Pool     | Average number of tasks in queue for the refresh thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_searcher_queue    | Tasks in OBS Searcher Queue          | Total number of tasks in queue for the OBS Searcher thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_searcher_rejected | Rejected Tasks in OBS Searcher Queue | Total number of tasks rejected by the OBS Searcher thread pool on nodes in a CSS cluster.  | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                            | Metric                                | Description                                                                                  | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|--------------------------------------|---------------------------------------|----------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| max_thread_pool_obs_searcher_queue   | Max. Tasks in OBS Searcher Queue      | Maximum number of tasks in queue for the OBS Searcher thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_searcher_threads | Size of OBS Searcher Thread Pool      | Total size of the OBS Searcher thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_searcher_queue   | Avg. Tasks in OBS Searcher Queue      | Average number of tasks in queue for the OBS Searcher thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_searcher_threads | Avg. Size of OBS Searcher Thread Pool | Average number of tasks in queue for the OBS Searcher thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_queue            | Tasks in OBS Queue                    | Total number of tasks in queue for the OBS thread pool on nodes in a CSS cluster.            | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                    | Metric                       | Description                                                                         | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|------------------------------|------------------------------|-------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_obs_rejected | Rejected Tasks in OBS Queue  | Total number of tasks rejected by the OBS thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_obs_queue    | Max. Tasks in OBS Queue      | Maximum number of tasks in queue for the OBS thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_threads  | Size of OBS Thread Pool      | Total size of the OBS thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_queue    | Avg. Tasks in OBS Queue      | Average number of tasks in queue for the OBS thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_threads  | Avg. Size of OBS Thread Pool | Average number of tasks in queue for the OBS thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                           | Metric                             | Description                                                                                | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------------|------------------------------------|--------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_thread_pool_obs_upload_queue    | Tasks in OBS Upload Queue          | Total number of tasks in queue for the OBS Upload thread pool on nodes in a CSS cluster.   | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_upload_rejected | Rejected Tasks in OBS Upload Queue | Total number of tasks rejected by the OBS Upload thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_obs_upload_queue    | Max. Tasks in OBS Upload Queue     | Maximum number of tasks in queue for the OBS Upload thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_upload_threads  | Size of OBS Upload Thread Pool     | Total size of the OBS Upload thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_upload_queue    | Avg. Tasks in OBS Upload Queue     | Average number of tasks in queue for the OBS Upload thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |



| Metric ID                             | Metric                               | Description                                                                                  | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|---------------------------------------|--------------------------------------|----------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_thread_pool_obs_upload_threads    | Avg. Size of OBS Upload Thread Pool  | Average number of tasks in queue for the OBS Upload thread pool on nodes in a CSS cluster.   | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_download_queue    | Tasks in OBS Download Queue          | Total number of tasks in queue for the OBS Download thread pool on nodes in a CSS cluster.   | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_download_rejected | Rejected Tasks in OBS Download Queue | Total number of tasks rejected by the OBS Download thread pool on nodes in a CSS cluster.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_thread_pool_obs_download_queue    | Max. Tasks in OBS Download Queue     | Maximum number of tasks in queue for the OBS Download thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_thread_pool_obs_download_threads  | Size of OBS Download Thread Pool     | Total size of the OBS Download thread pools on nodes in a CSS cluster.                       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                            | Metric                                 | Description                                                                                  | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|--------------------------------------|----------------------------------------|----------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_thread_pool_obs_download_queue   | Avg. Tasks in OBS Download Queue       | Average number of tasks in queue for the OBS Download thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_thread_pool_obs_download_threads | Avg. Size of OBS Download Thread Pool  | Average number of tasks in queue for the OBS Download thread pool on nodes in a CSS cluster. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| min_free_fs_size                     | Min. Available Storage Space           | Minimum available storage space of nodes in a CSS cluster.                                   | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| avg_jvm_old_gc_count                 | Avg. GCs of Old-Generation JVM         | Average number of old-generation garbage collections of nodes in a CSS cluster               | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_jvm_old_gc_time                  | Avg. GC Duration of Old-Generation JVM | Average old-generation garbage collection duration of nodes in a CSS cluster                 | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                | Metric                                   | Description                                                                                    | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|--------------------------|------------------------------------------|------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_jvm_young_gc_count   | Avg. GCs of Young-Generation JVM         | Average number of young-generation garbage collections of nodes in a CSS cluster               | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_jvm_young_gc_time    | Avg. GC Duration of Young-Generation JVM | Average young-generation garbage collection duration of nodes in a CSS cluster                 | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |
| 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. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_mem_free_in_bytes    | Avg. Available Memory                    | Average unused memory capacity of nodes in a CSS cluster.                                      | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| avg_mem_free_percent     | Avg. Available Memory Percentage         | Average percentage of unused memory of nodes in a CSS cluster.                                 | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| avg_mem_used_in_bytes    | Avg. Used Memory                         | Average used memory of nodes in a CSS cluster.                                                 | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |

| Metric ID             | Metric                                  | Description                                                                   | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------|-----------------------------------------|-------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_mem_used_percent  | Avg. Used Memory Percentage             | Average percentage of used memory of nodes in a CSS cluster.                  | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| max_mem_free_in_bytes | Max. Available Memory                   | Maximum unused memory of nodes in a CSS cluster.                              | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| max_mem_free_percent  | Max. Available Memory Percentage        | Maximum percentage of unused memory of nodes in a CSS cluster.                | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| max_mem_used_in_bytes | Max. Used Memory                        | Maximum used memory of nodes in a CSS cluster.                                | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| max_mem_used_percent  | Max. Used Memory Percentage             | Maximum percentage of used memory of nodes in a CSS cluster.                  | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| sum_jvm_old_gc_count  | Total GCs of Old-Generation JVM         | Number of old-generation garbage collections of nodes in a CSS cluster.       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_jvm_old_gc_time   | Total GC Duration of Old-Generation JVM | Total old-generation garbage collection duration of nodes in the CSS cluster. | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                     | Metric                                    | Description                                                                     | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------|-------------------------------------------|---------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| sum_jvm_young_gc_count        | Total GCs of Young-Generation JVM         | Number of young-generation garbage collections of nodes in a CSS cluster.       | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_jvm_young_gc_time         | Total GC Duration of Young-Generation JVM | Total young-generation garbage collection duration of nodes in the CSS cluster. | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |
| sum_current_opened_http_count | Currently Open HTTP Connections           | Number of HTTP connections that are open on nodes in a CSS cluster.             | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_total_opened_http_count   | Historical Open HTTP Connections          | Number of HTTP connections that were open on nodes in a CSS cluster.            | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| IndexingLatency               | Average Index Latency                     | Average time required for a shard to complete an indexing operation.            | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |
| IndexingRate                  | Average Index Rate                        | Average number of index operations per second in a cluster.                     | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                    | Metric                           | Description                                                                    | Value Range | Unit      | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|------------------------------|----------------------------------|--------------------------------------------------------------------------------|-------------|-----------|---------------|------------------------------|---------------------------|
| SearchLatency                | Average Search Latency           | Average time required for a shard to complete a search operation.              | $\geq 0$    | ms        | N/A           | CSS cluster                  | 1 minute                  |
| SearchRate                   | Average QPS                      | Average queries per second (QPS) in a cluster.                                 | $\geq 0$    | Count     | N/A           | CSS cluster                  | 1 minute                  |
| task_max_running_time        | Max. Task Runtime                | Duration of the most time-consuming search or write task in the cluster.       | $\geq 0$    | ms        | N/A           | CSS cluster                  | 1 minute                  |
| number_of_pending_tasks      | Number of Pending Tasks in Queue | Number of pending tasks to be processed by the master node in the CSS cluster. | $\geq 0$    | Count     | N/A           | CSS cluster                  | 1 minute                  |
| sum_disk_read_requests_rate  | Total Disk Read IOPS             | Number of disk read requests to all nodes in the CSS cluster per second.       | $\geq 0$    | request/s | N/A           | CSS cluster                  | 1 minute                  |
| sum_disk_write_requests_rate | Total Disk Write IOPS            | Number of disk write requests to all nodes in the CSS cluster per second.      | $\geq 0$    | request/s | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                               | Metric                                                    | Description                                                                                                                                                           | Value Range | Unit   | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------------------|-----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|--------|---------------|------------------------------|---------------------------|
| sum_disk_read_bytes_rate                | Total Disk Read Rate                                      | Number of bytes of data read from the disks of all nodes in the CSS cluster per second.                                                                               | $\geq 0$    | Byte/s | 1024          | CSS cluster                  | 1 minute                  |
| sum_disk_write_bytes_rate               | Total Disk Write Rate                                     | Number of bytes of data written into the disks of all nodes in the CSS cluster per second.                                                                            | $\geq 0$    | Byte/s | 1024          | CSS cluster                  | 1 minute                  |
| avg_vector_index_off_heap_used_in_bytes | Average Off-Heap Memory Usage of Vector Indexes           | Average off-heap memory (bytes) used for querying vector indexes on each node in a CSS cluster.                                                                       | $\geq 0$    | byte   | 1024          | CSS cluster                  | 1 minute                  |
| avg_vector_index_off_heap_usage         | Average Off-Heap Memory Utilization (%) of Vector Indexes | Average off-heap memory (in %) used for querying vector indexes on each node in a CSS cluster.<br><br>This metric is calculated only when vector index queries occur. | 0~100       | %      | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                               | Metric                                                    | Description                                                                                                                                                                                                 | Value Range | Unit | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------------------|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------|---------------|------------------------------|---------------------------|
| max_vector_index_off_heap_used_in_bytes | Maximum Off-Heap Memory Usage of Vector Indexes           | Maximum off-heap memory (bytes) used for querying vector indexes on each node in a CSS cluster.                                                                                                             | ≥ 0         | byte | 1024          | CSS cluster                  | 1 minute                  |
| max_vector_index_off_heap_usage         | Maximum Off-Heap Memory Utilization (%) of Vector Indexes | Maximum off-heap memory (in %) used for querying vector indexes on each node in a CSS cluster.<br><br>This metric is calculated only when vector index queries occur.                                       | 0~100       | %    | N/A           | CSS cluster                  | 1 minute                  |
| vector_index_circuit_breaker_status     | Vector Index Circuit Breaker Status                       | Circuit breaker status for vector indexes. <ul style="list-style-type: none"><li>• 0: Not enabled.</li><li>• 1: Enabled. Check whether the cluster's off-heap memory usage exceeds the threshold.</li></ul> | 0, 1        | N/A  | N/A           | CSS cluster                  | 1 minute                  |



| Metric ID                         | Metric                                                               | Description                                                                                                                                                                                                                | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------------|----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| cold_data_storage                 | The size of cold data stored                                         | The total size of cold data stored in a CSS cluster, which is the total size of the data of frozen indexes stored in an OBS bucket. This metric is measured hourly. Users cannot check this OBS bucket on the OBS console. | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| number_of_index_creation_failures | Number of index creation failures                                    | Number of index creation failures in a CSS cluster.                                                                                                                                                                        | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| shard_doc_exceed_threshold_count  | Number of shards where the number of documents exceeds the threshold | Number of shards where the number of documents exceeds the threshold in a CSS cluster.                                                                                                                                     | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| write_fail_alias_error            | Number of Write Failures Due to Alias Exceptions                     | Number of write failures caused by abnormal index aliases in a CSS cluster.<br>Only OpenSearch 1.3.6 and 2.19.0 clusters support this metric.                                                                              | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                           | Metric                                              | Description                                                                                                                                     | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------------|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| write_fail_documents_count_exceeded | Number of Write Failures Due to Excessive Documents | Number of write failures caused by too many index documents in a CSS cluster.<br>Only OpenSearch 1.3.6 and 2.19.0 clusters support this metric. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| write_fail_fields_count_exceed      | Number of Write Failures Due to Excessive Fields    | Number of write failures caused by too many index fields in a CSS cluster.<br>Only OpenSearch 1.3.6 and 2.19.0 clusters support this metric.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| write_fail_shards_count_exceed      | Number of Write Failures Due to Excessive Shards    | Number of write failures caused by too many index shards in a CSS cluster.<br>Only OpenSearch 1.3.6 and 2.19.0 clusters support this metric.    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                     | Metric                                                  | Description                                                                                                                                                                        | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-------------------------------|---------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| write_fail_index_blocks_write | Number of Write Failures Due to Index Write Prohibition | Number of write failures caused by the setting of write prohibition in index configuration in a CSS cluster.<br><br>Only OpenSearch 1.3.6 and 2.19.0 clusters support this metric. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| write_fail_mapping_malformed  | Number of Write Failures Due to Incorrect Field Types   | Number of write failures caused by incorrect index field types in a CSS cluster.<br><br>Only OpenSearch 1.3.6 and 2.19.0 clusters support this metric.                             | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| index_storage                 | Index Data Volume                                       | Total data size of all indexes in a CSS cluster.                                                                                                                                   | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| backup_failed_task_count      | Number of Failed Backup Tasks                           | Number of failed snapshot backup tasks in a CSS cluster.                                                                                                                           | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID                   | Metric                            | Description                                                                          | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|-----------------------------|-----------------------------------|--------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| disk_total                  | Total Disk Space                  | Total data disk capacity of all nodes in a CSS cluster.                              | $\geq 0$    | byte  | 1024          | CSS cluster                  | 1 minute                  |
| number_of_unassigned_shards | Number of Unallocated Shards      | Total number of unallocated shards in a CSS cluster.                                 | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| disk_iops_qos_num           | Disk IOPS Overload Times          | Total number of times the disk IOPS exceeds the predefined threshold.                | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| disk_jobw_qos_num           | Disk I/O Bandwidth Overload Times | Total number of times the disk I/O bandwidth usage exceeds the predefined threshold. | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_disk_io_util            | Maximum Disk I/O Usage            | Maximum disk I/O usage among all disks.                                              | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| avg_disk_io_util            | Average Disk I/O Usage            | Average disk I/O usage among all disks.                                              | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| shard_quota_usage           | Shard Quota Usage                 | Ratio of allocated shards to the upper limit of the shard quantity.                  | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |

Node Monitoring Metrics

Table 3-161 CSS node monitoring metrics

| Metric ID             | Metric Name           | Description                                         | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|-----------------------|-----------------------|-----------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| jvm_heap_usage        | JVM Heap Usage        | JVM heap memory usage of a node.                    | 0~100       | %     | N/A           | CSS cluster - cloud service node | 1 minute                  |
| cpu_usage             | CPU Usage             | CPU usage.                                          | 0~100       | %     | N/A           | CSS cluster - cloud service node | 1 minute                  |
| disk_util             | Max. Disk Usage       | Disk usage of a node                                | 0~100       | %     | N/A           | CSS cluster - cloud service node | 1 minute                  |
| load_average          | Average Load          | Average number of queued tasks per minute on a node | ≥ 0         | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| open_file_descriptors | Open File Descriptors | Number of opened file descriptors on a node         | ≥ 0         | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID                     | Metric Name                   | Description                                               | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|-------------------------------|-------------------------------|-----------------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| max_file_descriptors          | Max. Allowed File Descriptors | Maximum number of allowed file descriptors                | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_write_queue       | Tasks in Write Queue          | Number of tasks in queue for the write thread pool.       | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_search_queue      | Tasks in Search Queue         | Number of tasks in queue for the search thread pool.      | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_force_merge_queue | Tasks in ForceMerge Queue     | Number of tasks in queue for the force merge thread pool. | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_write_rejected    | Rejected Tasks in Write Queue | Number of tasks rejected by the write thread pool.        | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID                        | Metric Name                        | Description                                             | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|----------------------------------|------------------------------------|---------------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| thread_pool_search_rejected      | Rejected Tasks in Search Queue     | Number of tasks rejected by the search thread pool.     | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_force_merge_rejected | Rejected Tasks in ForceMerge Queue | Number of tasks rejected by the force merge thread pool | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_write_threads        | Size of Write Thread Pool          | Size of a write thread pool                             | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_search_threads       | Size of Search Thread Pool         | Size of a search thread pool                            | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_force_merge_threads  | Size of ForceMerge Thread Pool     | Size of a force merge thread pool                       | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID                    | Metric Name                     | Description                                           | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|------------------------------|---------------------------------|-------------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| thread_pool_flush_queue      | Tasks in Flush Queue            | Number of tasks in queue for the flush thread pool.   | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_flush_rejected   | Rejected Tasks in Flush Queue   | Number of tasks rejected by the flush thread pool.    | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_flush_threads    | Size of Flush Thread Pool       | Size of the flush thread pool.                        | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_generic_queue    | Tasks in Generic Queue          | Number of tasks in queue for the generic thread pool. | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_generic_rejected | Rejected Tasks in Generic Queue | Number of tasks rejected by the generic thread pool.  | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |



| Metric ID                       | Metric Name                        | Description                                              | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|---------------------------------|------------------------------------|----------------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| thread_pool_generic_threads     | Size of Generic Thread Pool        | Size of the generic thread pool.                         | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_management_queue    | Tasks in Management Queue          | Number of tasks in queue for the management thread pool. | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_management_rejected | Rejected Tasks in Management Queue | Number of tasks rejected by the management thread pool.  | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_management_threads  | Size of Management Thread Pool     | Size of the management thread pool.                      | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_refresh_queue       | Tasks in Refresh Queue             | Number of tasks in queue for the refresh thread pool.    | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID                         | Metric Name                          | Description                                                | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|-----------------------------------|--------------------------------------|------------------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| thread_pool_refresh_rejected      | Rejected Tasks in Refresh Queue      | Number of tasks rejected by the refresh thread pool.       | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_refresh_threads       | Size of Refresh Thread Pool          | Size of the refresh thread pool.                           | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_searcher_queue    | Tasks in OBS Searcher Queue          | Number of tasks in queue for the OBS Searcher thread pool. | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_searcher_rejected | Rejected Tasks in OBS Searcher Queue | Number of tasks rejected by the OBS Searcher thread pool.  | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_searcher_threads  | Size of OBS Searcher Thread Pool     | Size of the OBS Searcher thread pool.                      | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID                       | Metric Name                        | Description                                              | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|---------------------------------|------------------------------------|----------------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| thread_pool_obs_queue           | Tasks in OBS Queue                 | Number of tasks in queue for the OBS thread pool.        | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_rejected        | Rejected Tasks in OBS Queue        | Number of tasks rejected by the OBS thread pool.         | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_threads         | Size of OBS Thread Pool            | Size of the OBS thread pool.                             | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_upload_queue    | Tasks in OBS Upload Queue          | Number of tasks in queue for the OBS Upload thread pool. | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_upload_rejected | Rejected Tasks in OBS Upload Queue | Number of tasks rejected by the OBS Upload thread pool.  | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID                         | Metric Name                          | Description                                                | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|-----------------------------------|--------------------------------------|------------------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| thread_pool_obs_upload_threads    | Size of OBS Upload Thread Pool       | Size of the OBS Upload thread pool.                        | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_download_queue    | Tasks in OBS Download Queue          | Number of tasks in queue for the OBS Download thread pool. | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_download_rejected | Rejected Tasks in OBS Download Queue | Number of tasks rejected by the OBS Upload thread pool.    | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| thread_pool_obs_download_threads  | Size of OBS Download Thread Pool     | Size of the OBS Download thread pool.                      | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| free_fs_size                      | Available Size of File Systems       | Available size of file systems in a CSS cluster.           | $\geq 0$    | byte  | 1024          | CSS cluster - cloud service node | 1 minute                  |

| Metric ID          | Metric Name                             | Description                                         | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|--------------------|-----------------------------------------|-----------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| total_fs_size      | Total Size of File Systems              | Total size of file systems in a CSS cluster.        | $\geq 0$    | byte  | 1024          | CSS cluster - cloud service node | 1 minute                  |
| jvm_old_gc_count   | Total GCs of Old-Generation JVM         | Number of old-generation garbage collection times   | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| jvm_old_gc_time    | Total GC Duration of Old-Generation JVM | Time spent on old-generation garbage collection     | $\geq 0$    | ms    | N/A           | CSS cluster - cloud service node | 1 minute                  |
| jvm_young_gc_count | Total GCs of Young-Generation JVM       | Number of young-generation garbage collection times | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| jvm_young_gc_time  | GC Duration of Young-Generation JVM     | Time spent on young-generation garbage collection   | $\geq 0$    | ms    | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID                 | Metric Name                     | Description                                         | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|---------------------------|---------------------------------|-----------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| mem_free_in_bytes         | Available Memory                | Unused memory space of a node.                      | $\geq 0$    | byte  | 1024          | CSS cluster - cloud service node | 1 minute                  |
| mem_free_percent          | Available Memory Percentage     | Percentage of unused memory space on a node.        | 0~100       | %     | N/A           | CSS cluster - cloud service node | 1 minute                  |
| mem_used_in_bytes         | Used Memory                     | Used memory space of a node.                        | $\geq 0$    | byte  | 1024          | CSS cluster - cloud service node | 1 minute                  |
| current_opened_http_count | Currently Open HTTP Connections | Number of HTTP connections currently open on a node | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| total_opened_http_count   | Total Open HTTP Connections     | Total number of HTTP connections on a node          | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID                | Metric Name     | Description                                                                             | Value Range | Unit      | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|--------------------------|-----------------|-----------------------------------------------------------------------------------------|-------------|-----------|---------------|----------------------------------|---------------------------|
| disk_read_requests_rate  | Disk Read IOPS  | Number of disk read requests to a node in the CSS cluster per second.                   | $\geq 0$    | request/s | N/A           | CSS cluster - cloud service node | 1 minute                  |
| disk_write_requests_rate | Disk Write IOPS | Number of disk write requests to a node in the CSS cluster per second.                  | $\geq 0$    | request/s | N/A           | CSS cluster - cloud service node | 1 minute                  |
| disk_read_bytes_rate     | Disk Read Rate  | Number of bytes of data read from the disks of a node in the CSS cluster per second.    | $\geq 0$    | Byte/s    | 1024          | CSS cluster - cloud service node | 1 minute                  |
| disk_write_bytes_rate    | Disk Write Rate | Number of bytes of data written into the disks of a node in the CSS cluster per second. | $\geq 0$    | Byte/s    | 1024          | CSS cluster - cloud service node | 1 minute                  |
| shards_count             | Shards          | Number of shards per node in the CSS cluster.                                           | $\geq 0$    | Count     | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID                           | Metric Name                                       | Description                                                                          | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|-------------------------------------|---------------------------------------------------|--------------------------------------------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| vector_index_off_heap_used_in_bytes | Off-Heap Memory Usage of Vector Indexes           | Off-heap memory (bytes) used for querying vector indexes on nodes in a CSS cluster.  | $\geq 0$    | byte  | 1024          | CSS cluster - cloud service node | 1 minute                  |
| vector_index_off_heap_usage         | Off-Heap Memory Utilization (%) of Vector Indexes | Off-heap memory (in %) used for querying vector indexes on nodes in a CSS cluster.   | 0~100       | %     | N/A           | CSS cluster - cloud service node | 1 minute                  |
| disk_iops_qos_num                   | Disk IOPS Overload Times                          | Total number of times the disk IOPS exceeds the predefined threshold.                | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| disk_iobw_qos_num                   | Disk I/O Bandwidth Overload Times                 | Total number of times the disk I/O bandwidth usage exceeds the predefined threshold. | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| max_disk_io_util                    | Maximum Disk I/O Usage                            | Maximum disk I/O usage among all disks.                                              | 0~100       | %     | N/A           | CSS cluster - cloud service node | 1 minute                  |



| Metric ID         | Metric Name            | Description                                                         | Value Range | Unit | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|-------------------|------------------------|---------------------------------------------------------------------|-------------|------|---------------|----------------------------------|---------------------------|
| avg_disk_io_util  | Average Disk I/O Usage | Average disk I/O usage among all disks.                             | 0~100       | %    | N/A           | CSS cluster - cloud service node | 1 minute                  |
| shard_quota_usage | Shard Quota Usage      | Ratio of allocated shards to the upper limit of the shard quantity. | 0~100       | %    | N/A           | CSS cluster - cloud service node | 1 minute                  |

## Dimensions

**Table 3-162** Dimension description

| Key                    | Value                 |
|------------------------|-----------------------|
| cluster_id             | CSS cluster ID        |
| cluster_id,instance_id | CSS cluster - node ID |

### 3.15.2 Using Cloud Eye to Monitor OpenSearch Clusters

With CSS, you can use Cloud Eye to monitor created clusters. After configuring 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:** Configure custom alarm rules for the monitoring metrics. Once a metric exceeds the threshold, the system will notify you via email or HTTP/HTTPS.
2. **Configuring Monitoring Metrics:** Configure monitoring metrics for a cluster or a node in the cluster.
3. **Viewing Cluster Monitoring Information:** Monitor the trends and changes of the monitoring metrics over specific periods.

## Recommended Monitoring Metrics

- For routine monitoring of a cluster, you are advised to focus on disk usage and cluster health status.
- 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 search 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.

## Prerequisites

- The cluster is in the **Available** or **Processing** state.
- The cluster has been running properly for more than 10 minutes.

## 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 to filter alarm rules.

If no alarm rules are available, create one by referring to [Creating an Alarm Rule](#). [Table 3-163](#) describes the key parameters. For more parameters, customize them as needed.

**Table 3-163** Alarm rule configuration parameters

| Parameter      | Description                                                                                                                                                                                                                                                                   |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Alarm Type     | Select <b>Metric</b> .                                                                                                                                                                                                                                                        |
| Cloud Service  | Select <b>Cloud Search Service</b> .                                                                                                                                                                                                                                          |
| Resource Layer | Select the layer by which you are going to configure alarm rules. <ul style="list-style-type: none"><li>• <b>CSS Clusters</b>: Alarm rules are specified by cluster.</li><li>• <b>CSS Clusters - CSS Instances</b>: Alarm rules are specified by node in a cluster.</li></ul> |

## Configuring Monitoring Metrics


1. Create a monitoring dashboard by referring to [Creating a Dashboard](#). If an available monitoring dashboard has been created, skip this step.
2. Add CSS monitoring graphs by referring to [Adding a Graph](#).

## Viewing Cluster Monitoring Information

In the cluster list, check the monitoring information of clusters and cluster nodes.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and choose **More > Monitoring Metrics** in the **Operation** column.
4. On the **Monitoring Metrics** page, check the cluster and node metrics.

**Figure 3-70** Monitoring Metrics



- Click the **Instances** tab to check cluster metrics.
  - **Instances:** Select a cluster to check its metrics. The current cluster is selected by default. Multiple clusters can be selected so that you may compare between them.
  - You can also set monitoring periods so that you can compare metrics between different periods for the same cluster.
  - **Add Graph Group:** You can add a custom group in addition to the default group. When you move the pointer over a group name on the left, the icons for changing the group name and deleting the group are displayed on the right. You can modify the group.
  - **Select Metric:** In the displayed **Select Metric** dialog box, set **All Metrics** and **Top N Metrics**.
- Click the **CSS Instances** tab to check the metrics of cluster nodes. Only the node monitoring information of the first cluster selected on the **Instances** tab is displayed.
  - **CSS Instances:** Select the cluster nodes whose metrics and information you want to check. By default, the first node of the current cluster is selected. Multiple nodes can be selected so that you may compare between them.
  - You can also set monitoring periods so that you can compare metrics between different periods for the same cluster.
  - **Add Graph Group:** You can add a custom group in addition to the default group. When you move the pointer over a group name on

the left, the icons for changing the group name and deleting the group are displayed on the right. You can modify the group.

- **Select Metric:** In the displayed **Select Metric** dialog box, set **All Metrics**.

### 3.15.3 Configuring Index Monitoring for an OpenSearch Cluster

Index monitoring helps you monitor the index status, statistics, and trends of a cluster, allowing you to promptly handle potential risks to ensure cluster reliability. By collecting index statistics, you can gain insights into various aspects of index performance and resource utilization, and make informed O&M decisions.

#### How the Feature Works

Index monitoring consists of two steps:

- **Data ingestion**
  - Index stats (such as the number of documents, data storage, and shard status) are ingested periodically.
  - Data is written into the **monitoring-eye-css-\*** index, where a data aging mechanism is applied.
  - The default data retention duration is seven days. Data is automatically deleted when this duration expires.
- **Monitoring chart generation**
  - Preset dashboards and custom visualizations on OpenSearch Dashboards enable intuitive index monitoring.
  - Time series analysis, trend analysis and comparison, and anomaly detection are supported.

#### Constraints

- Only OpenSearch 2.19.0 supports index monitoring.
- Indexes starting with **monitoring-eye-css-\*** are identified as monitoring indexes and will not be monitored. Do not use this prefix for regular indexes.
- The index pattern of **monitoring-eye-css-\*** must not be deleted while index monitoring is enabled. Otherwise, monitoring charts will become abnormal.

#### Logging In to OpenSearch Dashboards

Log in to OpenSearch Dashboards and go to the command execution page. OpenSearch clusters support multiple access methods. This topic uses OpenSearch Dashboards as an example to describe the operation procedures.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
4. In the left navigation pane, choose **Dev Tools**.

## Enabling Index Monitoring

1. Run the following command to enable index monitoring:

```
PUT _cluster/settings
{
 "persistent": {
 "css.monitoring.index.enabled": "true"
 }
}
```

2. To monitor a single index, run the following command:

```
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 3-164** Parameters for index monitoring

| Parameter                     | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| css.monitoring.index.enabled  | Boolean | Whether to enable index monitoring. Setting this parameter to <b>true</b> enables index monitoring.<br>Default value: <b>false</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| css.monitoring.index.interval | Time    | Index monitoring cycle.<br>Minimum value: <b>1s</b><br>Default value: <b>10s</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 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 specific type of indexes to monitor.<br>Example: <ul style="list-style-type: none"><li>• <b>"css.monitoring.index.indices": ["index_name"]</b> indicates only <i>index_name</i> is monitored.</li><li>• <b>"css.monitoring.index.indices": ["log_*"]</b> indicates that only indexes starting with <b>log_</b> are monitored.</li><li>• <b>"css.monitoring.index.indices": ["index1", "index2"]</b> indicates that <b>index1</b> and <b>index2</b> are monitored.</li></ul> Default value: <b>*</b> (indicating that all indexes are monitored) |

| Parameter                       | Type | Description                                                                                                                      |
|---------------------------------|------|----------------------------------------------------------------------------------------------------------------------------------|
| css.monitoring.history.duration | Time | Retention period of index monitoring data. The default period is a week.<br>Minimum value: <b>1d</b><br>Default value: <b>7d</b> |

## Checking the Read and Write Traffic of Indexes

After index monitoring is enabled for a cluster, you can check the read/write traffic of this cluster's indexes during specific periods of time.

- Run the following command to check the read and write traffic of all indexes:  
`GET _cat/monitoring`
- Run the following command to check the read and write traffic of a specified index:  
`GET _cat/monitoring/{indexname}`  
**{indexName}** indicates the name of the index whose read and write traffic you want to check.
- Run the following command to check the read and write traffic of an index during specified periods of time:  
`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 3-165** Parameters for checking read and write traffic

| Parameter | Type   | Description                                                                                                                                                                           |
|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| indexname | String | Index name.<br>You cannot check the traffic of system indexes, whose names start with a period (.).                                                                                   |
| begin     | Date   | Start time (UTC time) of the monitoring data you want to view.<br>Time format: strict_date_optional_time epoch_millis<br>The default start time is 5 minutes before the current time. |
| end       | Date   | End time (UTC time) of the monitoring data you want to view.<br>Time format: strict_date_optional_time epoch_millis<br>The default end time is the current time.                      |

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 3-166** Parameters in the returned information

| Parameter      | Description                                                                  |
|----------------|------------------------------------------------------------------------------|
| index          | Index name                                                                   |
| 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. |


## Checking Index Monitoring Data in OpenSearch Dashboards Charts

You can check preconfigured index monitoring charts on the **Dashboard** and **Visualizations** pages of OpenSearch Dashboards. You can also customize tables and charts.

- Check index monitoring results in preconfigured dashboards.
  - a. On the OpenSearch Dashboards console, expand the menu in the upper-left corner, and choose **Dashboards**.

- b. Click **[Monitoring] Index monitoring Dashboard** to check preconfigured dashboards.

**Figure 3-71** Preconfigured dashboards



The preconfigured dashboard displays the number of read and write operations per second in the cluster and the top 10 indexes that have the most read and write operations per second.

**Table 3-167** 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 space 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 that store 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.        |

- Check index monitoring results in custom visualizations.  
The index monitoring module periodically stores the index/stats information in the **monitoring-eyes-css** index. You can create custom charts on OpenSearch Dashboards to visualize the information.  
The following procedure describes how to check the trend of document quantities in a chart as an example.
  - a. Click **Visualize** in the navigation tree on the left of the OpenSearch Dashboards console.
  - b. Click **Create visualization** and select **TSVB**.
  - c. Set chart parameters and view the visualizations.  
On the **Data** tab page, set the parameters as needed.
    - Select **Max** for **Aggregation**, and select **index\_stats primaries.docs.count** in **Field**, indicating the number of documents in a primary shard.
    - Select **Derivative** from **Aggregation** to indicate differences between aggregation buckets. Set **Units** to **1s** to visualize network rates as "per second".
    - Set **Aggregation** to **Positive Only** to prevent negative numbers after resetting.
    - To show statistics by index, set **Group by** to **Terms** and **By** to **index\_stats.index**. Statistics will be grouped by index name.
    - 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 3-72 TSVB page





Figure 3-73 Setting the interval




If the index monitoring charts are not displayed, re-load them on OpenSearch Dashboards.

 **NOTE**

If the preset dashboards and visualizations cannot be found for a security-mode OpenSearch cluster, try switching to the Private or Global space. If the issue persists, import the charts and graphs again.

1. Create the **monitoring-dashboards.ndjson** file by referring to [dashboards-monitor Configuration File](#).
2. On OpenSearch Dashboards, choose **Management > Dashboards Management > Saved objects**.
3. Click **Import** and upload the **monitoring-dashboards.ndjson** file created in 1.
4. After the file is uploaded, click **done**. The index monitoring charts are imported successfully.

Figure 3-74 Index monitoring charts imported successfully



**dashboards-monitor Configuration File**

The content of **dashboards-monitor** configuration file is as follows. You are advised to save the file as **monitoring-dashboards.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\": \"metrics\", \"aggs\": [], \"params\": {
 \"id\": \"61ca57f0-469d-11e7-af02-69e470af7417\", \"type\": \"timeseries\", \"series\": [{
 \"id\": \"61ca57f1-469d-11e7-af02-69e470af7417\", \"color\": \"#68BC00\", \"split_model\": \"terms\",
 \"split_color_model\": \"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\",
 \"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\": {
 \"monitoring\"
 }
 }, \"uiStateJSON\": {
 \"version\": 1, \"visState\": {
 \"aggs\": [], \"params\": {
 \"axis_formatter\": \"number\", \"axis_position\": \"left\", \"axis_scale\": \"normal\", \"default_index_pattern\": \"monitoring-eye-css-*\",
 \"default_timefield\": \"timestamp\", \"filter\": {
 \"language\": \"kuery\", \"query\": \"\"
 }, \"id\": \"61ca57f0-469d-11e7-af02-69e470af7417\", \"index_pattern\": \"monitoring-eye-css-*\",
 \"interval\": \"\", \"isModelInvalid\": false, \"series\": [{
 \"axis_position\": \"right\", \"chart_type\": \"line\", \"color\": \"rgba(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\": \"index_stats.total.segments.count\", \"id\": \"61ca57f2-469d-11e7-af02-69e470af7417\", \"type\": \"max\",
 \"point_size\": 1, \"separate_axis\": 0, \"split_color_model\": \"kibana\", \"split_model\": \"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\": {
 \"monitoring\"
 }
 }, \"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\": \"rgba(0,32,188,1)\", \"split_model\": \"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_model\": \"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\": \"rgba(0,33,224,1)\", \"split_model\": \"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_model\": \"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": "timeseries", "aggs": [{"id": "61ca57f1-469d-11e7-af02-69e470af7417", "type": "timeseries", "series": [{"id": "61ca57f1-469d-11e7-af02-69e470af7417", "color": "#68BC00", "split_model": "everything", "split_color_model": "kibana", "metrics": [{"unit": "bytes", "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": "1", "axis_position": "left", "axis_formatter": "number", "axis_scale": "normal", "show_legend": 1, "show_grid": 1, "tooltip_model": "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": "timeseries", "aggs": [{"id": "61ca57f1-469d-11e7-af02-69e470af7417", "type": "timeseries", "series": [{"id": "61ca57f1-469d-11e7-af02-69e470af7417", "color": "#68BC00", "split_model": "everything", "split_color_model": "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": "1", "axis_position": "left", "axis_formatter": "number", "axis_scale": "normal", "show_legend": 1, "show_grid": 1, "tooltip_model": "show_all", "default_index_pattern": "monitoring-eye-css-*", "default_timefield": "timestamp", "isModelInvalid": false, "drop_last_bucket": 1, "legend_position": "bottom"}], "id": "cfbb4e20-664c-11ed-8cd7-973626cf6f70", "references": [{"type": "visualization", "updated_at": "2022-12-01T12:41:01.165Z", "version": "WzlxMiwXQ=="}, {"attributes": {"description": "", "kibanaSavedObjectMeta": {"searchSourceJSON": {}}, "title": "[monitoring] total store size in bytes", "uiStateJSON": {}}, {"version": 1, "visState": {"title": "[monitoring] total store size in bytes", "type": "timeseries", "aggs": [{"id": "61ca57f0-469d-11e7-af02-69e470af7417", "type": "timeseries", "series": [{"id": "61ca57f1-469d-11e7-af02-69e470af7417", "color": "#68BC00", "split_model": "everything", "split_color_model": "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": "1", "axis_position": "left", "axis_formatter": "number", "axis_scale": "normal", "show_legend": 1, "show_grid": 1, "tooltip_model": "show_all", "default_index_pattern": "monitoring-eye-css-*", "default_timefield": "timestamp", "isModelInvalid": false, "legend_position": "bottom", "background_color_rules": [{"id": "7712e550-664f-11ed-8b5d-8db37e5b4cc4"}, {"bar_color_rules": [{"id": "77680a30-664f-11ed-8b5d-8db37e5b4cc4"}]}, {"id": "c7f72ae0-664e-11ed-8cd7-973626cf6f70", "references": [{"type": "visualization", "updated_at": "2022-12-01T12:41:01.165Z", "version": "WzlxMywyXQ=="}, {"attributes": {"description": "", "kibanaSavedObjectMeta": {"searchSourceJSON": {}}, "title": "[monitoring] indexing rate of index for top10(/s)", "uiStateJSON": {}}, {"version": 1, "visState": {"title": "[monitoring] indexing rate of index for top10(/s)", "type": "timeseries", "aggs": [{"id": "61ca57f0-469d-11e7-af02-69e470af7417", "type": "timeseries", "series": [{"id": "61ca57f1-469d-11e7-af02-69e470af7417", "color": "#68BC00", "split_model": "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": "1", "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_model": "rainbow"}], "time_field": "timestamp", "index_pattern": "monitoring-eye-css-*", "interval": "1", "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_model": "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": "timeseries", "aggs": [{"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\\\":\\\"fd9dfad0-a5ef-11ec-aa10-992297d21a2e\\\",\\\"type\\\":\\\"derivative\\\",\\\"field\\\":\\\"61ca57f2-469d-11e7-af02-69e470af7417\\\",{\\\"unit\\\":\\\"\\\",\\\"id\\\":\\\"0aaa26a0-a5f0-11ec-aa10-992297d21a2e\\\",\\\"type\\\":\\\"positive_only\\\",\\\"field\\\":\\\"fd9dfad0-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_model\\\":\\\"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_model\\\":\\\"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_model\\\":\\\"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\\\":0\\\",\\\"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-bf36-6cba3b09586c\"}, \"panelIndex\": \"dd14182d-d8b9-47f2-bf36-6cba3b09586c\", \"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\": []}
```

### 3.15.4 Setting Alarm Alerting via SMN for an OpenSearch Cluster

This topic describes how to use an open-source alerting or notification plugin to configure alarm alerting via the Simple Message Notification (SMN) service for OpenSearch clusters in OpenSearch Dashboards.

#### Scenarios

In CSS, the open-source OpenSearch alerting plugin (opensearch-alerting) and OpenSearch notification plugin (opensearch-notifications and opensearch-notifications-core) are installed by default. They trigger alerts when data meets predefined conditions.

- By default, the open-source alert plugin **opensearch-alerting** is installed for OpenSearch clusters of version 1.3.6. This plugin consists of three components: **Alerts**, **Monitors**, and **Destinations**. CSS integrates the SMN service in the **Destinations** component, allowing it to send alarm notifications to the SMN service as a destination.
- For OpenSearch 2.19.0 clusters, the open-source OpenSearch alerting plugin (opensearch-alerting) and OpenSearch notification plugin (opensearch-notifications and opensearch-notifications-core) are installed by default. The Destinations component of the OpenSearch alerting plugin is now an independent notification plugin that manages notification channels. CSS integrates the SMN service in its **Notifications** component, so it can use SMN as a notification channel.

For more information about the OpenSearch alerting plugin, see [Alerting - OpenSearch Documentation](#). For more information about the OpenSearch notifications plugin, see [Notifications - OpenSearch Documentation](#).

## Prerequisites

- You have created a topic on the SMN console. For details, see [Creating a Topic](#).
- You have obtained the CSS administrator account. This account has the permission to access CSS, check the agency list, create agencies, and grant permissions to agencies.

## Authorizing Users to Use SMN

1. Log in to the [CSS management console](#).  
You must log in using a CSS administrator account.
2. In the navigation pane, choose **Service Authorization**.
3. On the **Service Authorization** page, click **Create Agency**. In the dialog box displayed, confirm that the agency is successfully created.
  - If an agency has been created, "css\_smn\_agency exist, no need to created." is displayed in the upper right corner.
  - If you do not have the permission to create an agency, an error message will be displayed in the upper right corner indicating "no permission", in which case, check that the administrator account has been assigned the IAM permission.

## Setting Alarm Notifications via SMN (OpenSearch 2.19.0)

1. Log in to the [CSS management console](#).  
Log in using an account with CSS permissions.
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
4. On the OpenSearch Dashboards console, expand the menu in the upper-left corner, and choose **Notifications**.
5. Create an SMN channel to send alert messages.
  - a. On the **Channels** page, click **Create channel** and configure a notification channel.

**Table 3-168** Channel parameters

| Parameter   | Description                           |
|-------------|---------------------------------------|
| Name        | Custom channel name.                  |
| Description | Custom channel description.           |
| Type        | Retain the default value <b>SMN</b> . |

| Parameter | Description                                                                                        |
|-----------|----------------------------------------------------------------------------------------------------|
| Topic     | Select the SMN topic you have created in <a href="#">Prerequisites</a> for sending alert messages. |

**Figure 3-75** Create channel

**Create channel**

**Name and description**

**Name**

TestChannel

**Description - optional**

This is a test channel

**Configurations**

**Channel type**

Channel type cannot be changed after the channel is created.

SMN


**Topic**

Alerting\_test

- b. Click **Create**.
- c. Return to the Channels page. If the new channel is displayed on the Channels page, it has been created successfully.



Figure 3-76 Channels list



6. On the OpenSearch Dashboards console, expand the menu in the upper-left corner, and choose **Alerting**.
7. Create a monitor and configure alarm triggers and monitoring frequency.
  - a. Click the **Monitors** tab on the **Alerting** page and click **Create monitor** to configure monitor information.

Table 3-169 Monitor parameters

| Parameter               | Description                                                                                                                                                                                                                                                                                                                     |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Monitor details</b>  |                                                                                                                                                                                                                                                                                                                                 |
| Monitor name            | User-defined monitor name                                                                                                                                                                                                                                                                                                       |
| Monitor type            | Monitor type. The value can be <b>Per query monitor</b> (common monitoring), <b>Per bucket monitor</b> (aggregation bucket monitoring), and <b>Per cluster metrics monitor</b> (cluster metric monitoring).                                                                                                                     |
| Monitor defining method | Monitor defining method. <b>Extraction query editor</b> is recommended. <ul style="list-style-type: none"><li>• <b>Visual editor</b></li><li>• <b>Extraction query editor</b></li><li>• <b>Anomaly detector</b></li></ul> The options of <b>Monitor defining method</b> are determined by the <b>Monitor type</b> you selected. |
| Detector                | If <b>Monitor defining method</b> is set to <b>Anomaly detector</b> , select an exception detection task.                                                                                                                                                                                                                       |
| Frequency               | Select the monitoring frequency and set the monitoring interval. The options include: <ul style="list-style-type: none"><li>• <b>By interval</b></li><li>• <b>Daily</b></li><li>• <b>Weekly</b></li><li>• <b>Monthly</b></li><li>• <b>Custom cron expression</b></li></ul>                                                      |
| <b>Select data</b>      |                                                                                                                                                                                                                                                                                                                                 |


| Parameter                     | Description                                                                                                                                                                        |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Index                         | When <b>Monitor defining method</b> is set to <b>Visual editor</b> or <b>Extraction query editor</b> , you need to specify the index to be monitored.                              |
| Time field                    | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to specify the time field to define counting parameters such as <b>count</b> .                       |
| <b>Query</b>                  |                                                                                                                                                                                    |
| Metrics                       | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to set the metrics range for extracting statistics.                                                  |
| Time range for the last       | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to set the monitoring time range for plugins.                                                        |
| Data filter                   | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to set filters for data search.                                                                      |
| Group by                      | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to specify a field so that each value of the field triggers an alarm.                                |
| Define extraction query       | When <b>Monitor defining method</b> is set to <b>Extraction query editor</b> , you need to enter the query statement to define the monitoring.                                     |
| Request type                  | When <b>Monitor type</b> is set to <b>Per cluster metrics monitor</b> , you need to specify the request type to monitor cluster metrics, such as the running status and CPU usage. |
| Preview query and performance | Preview the query result and verify query performance under the current configuration.                                                                                             |

- b. Click **Add trigger** to add triggers and specify the alarm triggering conditions and actions to be triggered when an alarm is reported.
- c. On the **Triggers** page, set the alarm triggering sensitivity and message release on the destination end.



**Table 3-170** Trigger parameters

| Parameter      | Description                                                                                                                                                    |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Trigger name   | User-defined trigger name.                                                                                                                                     |
| Severity level | Sensitivity of a trigger, that is, the number of alarms that need to be triggered before an alarm message is sent. <b>1</b> indicates the highest sensitivity. |

| Parameter         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Trigger condition | Trigger condition. An alarm is triggered when the trigger condition is hit.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Action name       | Trigger action name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Channels          | Select the SMN destination created in <a href="#">5</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Message subject   | A description of the message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Message           | Alarm message body. By default, the subject and body are defined when the destination is an email address. For details, see <a href="#">Message Publishing</a> .                                                                                                                                                                                                                                                                                                                                                         |
| Perform action    | <p>When <b>Monitor type</b> is set to <b>Per bucket monitor</b>, you need to set whether to send alarms in combination. The value can be:</p> <ul style="list-style-type: none"><li>● <b>Per execution</b>: A combination alarm is sent when multiple alarm triggering conditions are hit.</li><li>● <b>Per alert</b>: Alarms are sent separately when multiple alarm triggering conditions are hit.</li></ul>                                                                                                           |
| Actionable alerts | <p>When <b>Monitor type</b> is set to <b>Per bucket monitor</b>, set this parameter to <b>Per alert</b>. You need to set the alarms that can be executed after alarm triggering conditions are hit.</p> <ul style="list-style-type: none"><li>● <b>De-duplicated</b>: Alarms that have been triggered. OpenSearch retains the existing alarms to prevent the plugin from creating duplicate alarms.</li><li>● <b>New</b>: Newly created alarms.</li><li>● <b>Completed</b>: Alarms that are no longer ongoing.</li></ul> |
| Throttling        | <p>Message sending frequency. It limits the number of notification messages can be received in a specified period.</p> <p>For example, if this parameter is set to 10 minutes, SMN sends only one alarm notification in the next 10 minutes even if the trigger condition is hit for multiple times. After 10 minutes, SMN sends another alarm notification if the alarm condition is met.</p>                                                                                                                           |

**Figure 3-77** Setting the destination of a trigger action

- d. Click **Send test message**. If a subscriber receives an email, as shown in [Figure 3-79](#), the trigger is configured successfully.

**Figure 3-78** Sending a test message**Figure 3-79** Email notification

- e. Click **Create** to return to the monitor details page. The detector is successfully created.

## Setting Alarm Notifications via SMN (OpenSearch 1.3.6)

1. Log in to the [CSS management console](#).  
Log in using an account with CSS permissions.
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.
4. On the **OpenSearch Dashboards** page, choose **OpenSearch Plugins > Alerting** in the navigation tree on the left.
5. Create an SMN destination to send alert messages.
  - a. On the **Alerting** page, click the **Destinations** tab and click **Add destination** to configure destination information.

**Table 3-171** Destinations parameters

| Parameter | Description                                                                                        |
|-----------|----------------------------------------------------------------------------------------------------|
| Name      | User-defined destination name                                                                      |
| Type      | Retain the default value <b>SMN</b> .                                                              |
| Topic     | Select the SMN topic you have created in <a href="#">Prerequisites</a> for sending alert messages. |

**Figure 3-80** Add destination

Alerting / Destinations / Create destination

## Add destination

### Destination

**Name**

Specify a name of the destination.

**Type**

SMN

### Settings

**Topic**

test

- b. Click **Create** to return to the destination list. The created SMN destination is displayed in the list.

**Figure 3-81** Destination list

Destinations

Search

All type

< 1 >

Destination name

Destination type

Last updated by

Actions

test

SMN

-

Edit

Delete

Rows per page: 20

< 1 >

6. Create a monitor and configure alarm triggers and monitoring frequency.

- a. Click the **Monitors** tab on the **Alerting** page and click **Create monitor** to configure monitor information.

**Table 3-172** Monitor parameters

| Parameter               | Description                                                                                                                                                                                                                                                                                                                     |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Monitor details</b>  |                                                                                                                                                                                                                                                                                                                                 |
| Monitor name            | User-defined monitor name                                                                                                                                                                                                                                                                                                       |
| Monitor type            | Monitor type. The value can be <b>Per query monitor</b> (common monitoring), <b>Per bucket monitor</b> (aggregation bucket monitoring), and <b>Per cluster metrics monitor</b> (cluster metric monitoring).                                                                                                                     |
| Monitor defining method | Monitor defining method. <b>Extraction query editor</b> is recommended. <ul style="list-style-type: none"><li>• <b>Visual editor</b></li><li>• <b>Extraction query editor</b></li><li>• <b>Anomaly detector</b></li></ul> The options of <b>Monitor defining method</b> are determined by the <b>Monitor type</b> you selected. |
| Detector                | If <b>Monitor defining method</b> is set to <b>Anomaly detector</b> , select an exception detection task.                                                                                                                                                                                                                       |
| Frequency               | Select the monitoring frequency and set the monitoring interval. The options include: <ul style="list-style-type: none"><li>• <b>By interval</b></li><li>• <b>Daily</b></li><li>• <b>Weekly</b></li><li>• <b>Monthly</b></li><li>• <b>Custom cron expression</b></li></ul>                                                      |
| <b>Data source</b>      |                                                                                                                                                                                                                                                                                                                                 |
| Index                   | When <b>Monitor defining method</b> is set to <b>Visual editor</b> or <b>Extraction query editor</b> , you need to specify the index to be monitored.                                                                                                                                                                           |
| Time field              | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to specify the time field to define counting parameters such as <b>count</b> .                                                                                                                                                                    |
| <b>Query</b>            |                                                                                                                                                                                                                                                                                                                                 |
| Metrics                 | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to set the metrics range for extracting statistics.                                                                                                                                                                                               |

| Parameter               | Description                                                                                                                                                                        |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time range for the last | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to set the monitoring time range for plugins.                                                        |
| Data filter             | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to set filters for data search.                                                                      |
| Group by                | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to specify a field so that each value of the field triggers an alarm.                                |
| Define extraction query | When <b>Monitor defining method</b> is set to <b>Extraction query editor</b> , you need to enter the query statement to define the monitoring.                                     |
| Request type            | When <b>Monitor type</b> is set to <b>Per cluster metrics monitor</b> , you need to specify the request type to monitor cluster metrics, such as the running status and CPU usage. |

- b. Click **Add trigger** to add triggers and specify the alarm triggering conditions and actions to be triggered when an alarm is reported.
- c. On the **Triggers** page, set the alarm triggering sensitivity and message release on the destination end.

**Table 3-173** Trigger parameters

| Parameter         | Description                                                                                                                                                     |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Trigger name      | User-defined trigger name                                                                                                                                       |
| Severity level    | Sensitivity of a trigger, that is, the number of alarms that need to be triggered before an alarm message is sent. <b>1</b> indicates the highest sensitivity.  |
| Trigger condition | Trigger condition. An alarm is triggered when the trigger condition is hit.                                                                                     |
| Action name       | Trigger action name                                                                                                                                             |
| Destination       | Select the SMN destination created in <a href="#">5</a> .                                                                                                       |
| Message           | Alarm message body By default, the subject and body are defined when the destination is an email address. For details, see <a href="#">Message Publishing</a> . |

| Parameter         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Perform action    | When <b>Monitor type</b> is set to <b>Per bucket monitor</b> , you need to set whether to send alarms in combination. The value can be: <ul style="list-style-type: none"><li>• <b>Per execution</b>: A combination alarm is sent when multiple alarm triggering conditions are hit.</li><li>• <b>Per alert</b>: Alarms are sent separately when multiple alarm triggering conditions are hit.</li></ul>                                                                                                            |
| Actionable alerts | When <b>Monitor type</b> is set to <b>Per bucket monitor</b> , set this parameter to <b>Per alert</b> . You need to set the alarms that can be executed after alarm triggering conditions are hit. <ul style="list-style-type: none"><li>• <b>De-duplicated</b>: Alarms that have been triggered. OpenSearch retains the existing alarms to prevent the plugin from creating duplicate alarms.</li><li>• <b>New</b>: Newly created alarms.</li><li>• <b>Completed</b>: Alarms that are no longer ongoing.</li></ul> |
| Throttling        | Message sending frequency. It limits the number of notification messages can be received in a specified period.<br><br>For example, if this parameter is set to 10 minutes, SMN sends only one alarm notification in the next 10 minutes even if the trigger condition is hit for multiple times. After 10 minutes, SMN sends another alarm notification if the alarm condition is met.                                                                                                                             |

**Figure 3-82** Setting the destination of a trigger action

SMN notification: Remove action

Action name  
Enter action name

Names can only contain letters, numbers, and special characters.

Destination  
test - (SMN)

Message  
Embed variables in your message using Mustache templates. [Learn more](#)

```
{
 "subject": "subject-example",
 "message": "Monitor {{ctx.monitor.name}} just entered alert status. Please investigate the issue. Trigger: {{ctx.trigger.name}} Severity: {{ctx.trigger.severity}} Period start: {{ctx.period.start}} Period end: {{ctx.period.end}}",
 "time_to_live": "3600"
}
```

☐ Preview message

Action configuration

Throttling  
☐ Enable action throttling

Send test message

- d. Click **Send test message**. If a subscriber receives an email, as shown in [Figure 3-84](#), the trigger is configured successfully.



**Figure 3-83** Sending a test message

Destination  
test - (SMN)

Message subject  
test

Message [Info](#)


```
{
 "subject": "subject",
 "message": "Monitor {{ctx.monitor.name}} just entered alert status. Please investigate the issue. - Trigger: {{ctx.trigger.name}} - Severity: {{ctx.trigger.severity}} - Period start: {{ctx.periodStart}} - Period end: {{ctx.periodEnd}}",
 "time_to_live": "3600"
}
```

Embed variables in your message using Mustache templates. [Learn more about Mustache](#)

[Send test message](#)

Message preview

```
{
 "subject": "subject",
 "message": "Monitor test-monitor just entered alert status. Please investigate the issue. - Trigger: test-trigger - Severity: 1 - Period start: 2023-04-14T02:16:36Z - Period end: 2023-04-14T02:17:36Z",
 "time_to_live": "3600"
}
```

**Figure 3-84** Email notification

- e. Click **Create** to return to the monitor details page. The detector is successfully created.

### 3.15.5 Synchronizing OpenSearch Alerts to Prometheus

Make necessary configurations in OpenSearch Dashboards to synchronize OpenSearch alerts to Prometheus so you can use Prometheus to monitor and analyze key performance metrics for OpenSearch clusters in real time.

Prometheus is an open-source monitoring system with a dimensional data model, flexible query language, efficient time series database and modern alerting approach.

#### Constraints

Only OpenSearch 2.19.0 clusters support alarm synchronization to Prometheus.

#### Prerequisites

- The Prometheus monitor server is ready, and the Pushgateway address has been obtained. Prometheus and OpenSearch must be connected. Otherwise, alerts cannot be sent.
- The target OpenSearch cluster is available.

#### Configuring Alert Synchronization

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, find the target cluster, and click **Dashboards** in the **Operation** column to log in to OpenSearch Dashboards.

4. On the **OpenSearch Dashboards** page, choose **OpenSearch Management > Notifications** in the navigation tree on the left.
5. Create a Prometheus channel to send alert messages.
  - a. On the **Channels** page, click **Create channel** to configure a channel.

**Table 3-174** Channel parameters

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                | Custom channel name.                                                                                                                                                                                                                                                                                                                                                                                      |
| Description         | Channel description.                                                                                                                                                                                                                                                                                                                                                                                      |
| Channel type        | Select <b>PROMETHEUS</b> .                                                                                                                                                                                                                                                                                                                                                                                |
| Prometheus Endpoint | <p>Enter the Pushgateway address of the Prometheus monitor server.</p> <ul style="list-style-type: none"><li>• Currently, only Prometheus Gauge dashboards can be added or deleted. Metrics are queried using specific statements and numeric values are synchronized to Pushgateway for monitoring via Prometheus.</li><li>• Two types of Pushgateway addresses are supported: HTTP and HTTPS.</li></ul> |

Figure 3-85 Create channel

# Create channel

## Name and description

**Name**

test

**Description - optional**

This is a test prometheus channel.

## Configurations

**Channel type**

Channel type cannot be changed after the channel is created.

PROMETHEUS

**Prometheus Endpoint**

http://xx.xx.xx.xx:9091

- b. Click **Create** to return to the Channels list. The created Prometheus channel is displayed in the list.

Figure 3-86 Channels list

Channels (1)

Actions

Create channel

Search

Status

Type

| <input type="checkbox"/> | Name ↑ | Notification status | Type       | Description                        |
|--------------------------|--------|---------------------|------------|------------------------------------|
| <input type="checkbox"/> | test   | Active              | PROMETHEUS | This is a test prometheus channel. |

Rows per page: 10

< 1 >

6. On the OpenSearch Dashboards console, expand the menu in the upper-left corner, and choose **Alerting**.

7. Create a monitor and configure alarm triggers and monitoring frequency.
  - a. Click the **Monitors** tab on the **Alerting** page, and click **Create monitors** to configure monitor information.

**Table 3-175** Monitor parameters

| Parameter               | Description                                                                                                                                                                                                                                                                                                                     |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Monitor details</b>  |                                                                                                                                                                                                                                                                                                                                 |
| Monitor name            | User-defined monitor name                                                                                                                                                                                                                                                                                                       |
| Monitor type            | Monitor type. The value can be <b>Per query monitor</b> (common monitoring), <b>Per bucket monitor</b> (aggregation bucket monitoring), and <b>Per cluster metrics monitor</b> (cluster metric monitoring).                                                                                                                     |
| Monitor defining method | Monitor defining method. <b>Extraction query editor</b> is recommended. <ul style="list-style-type: none"><li>• <b>Visual editor</b></li><li>• <b>Extraction query editor</b></li><li>• <b>Anomaly detector</b></li></ul> The options of <b>Monitor defining method</b> are determined by the <b>Monitor type</b> you selected. |
| Detector                | If <b>Monitor defining method</b> is set to <b>Anomaly detector</b> , select an exception detection task.                                                                                                                                                                                                                       |
| Frequency               | Select the monitoring frequency and set the monitoring interval. The options include: <ul style="list-style-type: none"><li>• <b>By interval</b></li><li>• <b>Daily</b></li><li>• <b>Weekly</b></li><li>• <b>Monthly</b></li><li>• <b>Custom cron expression</b></li></ul>                                                      |
| <b>Select data</b>      |                                                                                                                                                                                                                                                                                                                                 |
| Index                   | When <b>Monitor defining method</b> is set to <b>Visual editor</b> or <b>Extraction query editor</b> , you need to specify the index to be monitored.                                                                                                                                                                           |
| Time field              | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to specify the time field to define counting parameters such as <b>count</b> .                                                                                                                                                                    |
| <b>Query</b>            |                                                                                                                                                                                                                                                                                                                                 |
| Metrics                 | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to set the metrics range for extracting statistics.                                                                                                                                                                                               |

| Parameter                     | Description                                                                                                                                                                        |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time range for the last       | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to set the monitoring time range for plugins.                                                        |
| Data filter                   | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to set filters for data search.                                                                      |
| Group by                      | When <b>Monitor defining method</b> is set to <b>Visual editor</b> , you need to specify a field so that each value of the field triggers an alarm.                                |
| Define extraction query       | When <b>Monitor defining method</b> is set to <b>Extraction query editor</b> , you need to enter the query statement to define the monitoring.                                     |
| Request type                  | When <b>Monitor type</b> is set to <b>Per cluster metrics monitor</b> , you need to specify the request type to monitor cluster metrics, such as the running status and CPU usage. |
| Preview query and performance | Preview the query result and verify query performance under the current configuration.                                                                                             |

- b. Click **Create**. The **Create trigger** page is displayed.
- c. On the **Create trigger** page, set the alert triggering conditions and the actions to be triggered.


**Table 3-176** Trigger parameters

| Parameter         |                   | Description                                                                                                                                                                                                                                              |
|-------------------|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Define trigger    | Trigger name      | User-defined trigger name.                                                                                                                                                                                                                               |
|                   | Severity level    | Sensitivity of a trigger, that is, the number of alarms that need to be triggered before an alarm message is sent. <b>1</b> indicates the highest sensitivity.                                                                                           |
|                   | Trigger condition | Trigger condition. An alarm is triggered when the trigger condition is hit.<br><b>NOTE</b><br>You are advised to set a trigger condition that can almost always be triggered so that the queried metrics will always be synchronized to the Pushgateway. |
| Configure actions | Action name       | Trigger action name.                                                                                                                                                                                                                                     |
|                   | Destination       | Select the destination created in <a href="#">5</a> .                                                                                                                                                                                                    |


| Parameter |            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           | Message    | Defines the body of the message to be published, which must use the JSON format. See the following for an example. <pre>{  "metricsName": "hits_total_value", //Prometheus metric name  "metricsLabel":  {"label_key1": "label_value1", "label_key2": "label_value2"}, //Prometheus labels  "metricsValue": {{ctx.results.0.hits.total.value}}, // Prometheus metric values  "jobName": "job_name" //Prometheus monitor task name  "metricsHelp": "*****" //Metric explanation. Optional.}</pre> |
|           | throttling | Message sending frequency. It limits the number of notification messages can be received in a specified period.<br><br>For example, if this parameter is set to 10 minutes, Prometheus sends only one alert notification in the next 10 minutes even if the trigger condition is met multiple times. After 10 minutes, Prometheus sends another alert notification if the trigger condition is met again.                                                                                        |

- d. Click **Send test message** to send a test message to Prometheus to check whether the trigger is set successfully.

Figure 3-87 Send a test message.



As shown in [Figure 3-88](#), Prometheus can receive a triggered message, meaning the trigger is set successfully.

**Figure 3-88** Message received successfully

e. Click **Create** to return to the **Monitor** details page.

### 3.15.6 Intelligent Risk Detection for OpenSearch Clusters

CSS provides intelligent O&M to help detect potential risks for clusters, along with risk handling suggestions.

#### Scenarios

Intelligent O&M for clusters supports the following functions:

- **Creating a Scan Task:** Start a scan task to trigger an intelligent health check and diagnosis on the current cluster.
- **Checking the Risk Items of a Cluster:** After a scan task is completed, check the risks identified by the scan, confirm and handle these risks in a timely manner based on risk handling suggestions.
- **Deleting a Scan Task:** Delete scan tasks that you no longer need. After a scan task is deleted, the system deletes all diagnoses generated by it.

The check items of intelligent O&M are as follows:

- Check the current health status of a cluster. Red: Some primary shards are not allocated. Yellow: Some replica shards are not allocated. Green: All shards have been allocated.
- Check the number of nodes and the number of AZs of a distributed OpenSearch cluster to evaluate its high-availability status.
- Check whether index replicas are enabled. An index with no replicas may become unavailable in the case of a node failure. If local disks are used, this may even lead to data loss.
- Check for Kibana index conflicts in a cluster.
- Check disk usage. If the disk usage of a node is too high, new index shards may fail to be allocated to the node and the cluster performance may be affected.
- Check whether the storage usage of a cluster's data nodes or cold data nodes is balanced. Unbalanced storage distribution may result in unbalanced cluster loads and increase read/write latency.
- Monitor node disconnection or unavailability in a cluster for 5 consecutive minutes each time.
- Check for nodes with too many shards. An excessively large number of shards will consume excessive node resources, increasing read/write latency and slowing down metadata update.

- Check the sizes of all shards. A large shard may impact performance, occupy excessive node memory, and slow down shard restoration during cluster scaling or fault recovery.
- Check for new versions the current cluster may upgrade to.
- Check for snapshot creation failures or the absence of any snapshot creation records in the last seven days.

## Granting Access to SMN

To send alarm notifications via SMN after a scan task is completed, you must first be granted access to the SMN service. Additionally, you must create a topic on the SMN console in advance. For details, see [Creating a Topic](#).

1. Log in to the [CSS management console](#).  
You must log in using a CSS administrator account.
2. In the navigation pane, choose **Service Authorization**.
3. On the **Service Authorization** page, click **Create Agency**. In the dialog box displayed, confirm that the agency is successfully created.
  - If an agency has been created, "css\_smn\_agency exist, no need to created." is displayed in the upper right corner.
  - If you do not have the permission to create an agency, an error message will be displayed in the upper right corner indicating "no permission", in which case, check that the administrator account has been assigned the IAM permission.

## Creating a Scan Task

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Monitoring > Intelligent O&M**.
5. On the **Intelligent O&M** tab, click **Scan** in the upper-left corner. In the displayed dialog box, configure the scan task.

**Table 3-177** Configuring a scan task

| Parameter   | Description                                                                                                                                                              |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name        | Custom name of a scan task.<br>Enter 4 to 64 characters and start with a lowercase letter. Only lowercase letters, digits, hyphens (-), and underscores (_) are allowed. |
| Description | Brief description of the task.                                                                                                                                           |




| Parameter                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Send SMN notification upon task completion | Whether to send an SMN notification upon completion of the scan. <ul style="list-style-type: none"><li>When selected, you need to further configure <b>SMN Topic</b> and <b>Notification Level</b>. Upon completion of the scan, an SMN notification is sent if a risk higher than or equal to the notification level you set is detected.</li><li>When deselected (default), no SMN notification will be sent upon completion of the scan.</li></ul> |
| SMN Topic                                  | If you select <b>Send SMN notification upon task completion</b> , you need to set the SMN topic.                                                                                                                                                                                                                                                                                                                                                      |
| Notification Level                         | If you select <b>Send SMN notification upon task completion</b> , you need to set the risk level.<br>If the scan result contains risks at this level or higher, SMN will send notifications that list all the risk items in the result.                                                                                                                                                                                                               |

- After the configuration is complete, click **OK**.
- When the status of the scan task changes to **Completed**, you can check the cluster's risk items.

## Checking the Risk Items of a Cluster


- On the **Intelligent O&M** page, select a completed scan task, and click  on the lower-right left to check the task details, including the task creation time, summary, ID, and the risk items found by it.

**Figure 3-89** Checking scan task details



- Click  on the left of a risk item to check its details, including the check item, risk description, and risk handling suggestion.

You can handle cluster risks in a timely manner based on the suggestions provided by the system.

**Figure 3-90 Risk items**

3. Select a scan task, and click **Export Risk** in the **Operation** column to download the scan result.

## Deleting a Scan Task

You can delete scan tasks that are no longer needed. Deleting a scan task also deletes all its diagnoses and scan report.

1. On the **Intelligent O&M** page, locate the scan task you want to delete, and click **Delete** in the **Operation** column.
2. In the displayed dialog box, enter **DELETE**, and click **OK** to delete the task.

## 3.15.7 Managing OpenSearch Cluster Logs

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

- Log query: On the log management page of the CSS management console, you can query the latest log records by node, log type, and other criteria, so you can quickly locate or diagnose issues.
- Log backup: Cluster logs are periodically synchronized to OBS buckets. You can download them for in-depth analysis at any time. You can configure custom log backup policies by specifying backup schedules and storage locations. The system backs up all critical logs, including run logs, slow query logs, and deprecation logs. They provide comprehensive data for auditing and troubleshooting purposes.
- Log ingestion: Cluster logs are transmitted in real time to the current cluster or another cluster on the same network. (Certain version and network compatibility requirements must be met.) You can use OpenSearch Dashboards or other visualization tools for log search and analysis. You can also set index prefixes and retention periods to flexibly manage the log lifecycle. Logs can be transferred across clusters, facilitating centralized monitoring in a multi-cluster environment.

## Impact on Billing

When log backup is enabled, the generated log backups are stored in OBS buckets, which will result in additional costs. For details, see [Object Storage Service Billing](#).

## Prerequisites

The OBS bucket used for storing log backups has been created. The OBS bucket must meet the following requirements:

- **Storage Class: Standard.**
- **Region:** the same as that of the cluster.

## Log Search

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Logs > Log Search**. The **Log Search** page is displayed.

You can search records by log type, node, log level, or keyword. For a detailed description of each type of logs, see [Log Introduction](#).

When a log file reaches 128 MB or when the time reaches 00:00 UTC, the system automatically compresses and archives it. Only unarchived logs appear on the log search page, while archived logs remain accessible through the log backup function.

## Log Backup

Cluster logs can be backed up to OBS buckets, where you can download them for in-depth analysis at any time.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > OpenSearch**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Logs > Log Backup**. The **Log Backup** page is displayed.
5. Enable log backup.

Perform the following steps to enable log backup. If it is already enabled, skip this step.

- a. Click **Enable Backup**. In the displayed dialog box, configure necessary settings.

**Table 3-178** Log backup settings

| Parameter  | Description                                                                                                                                                                                                                                         |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBS Bucket | Select an OBS bucket for storing log backups from the drop-down list box.<br><br>If no OBS buckets meet your requirements, click <b>Create Bucket</b> to go to the OBS console and create one. For details, see <a href="#">Creating a Bucket</a> . |

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                                       |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup Path | <p>Set the log storage location in the OBS bucket.</p> <p>The backup path cannot:</p> <ul style="list-style-type: none"><li>• Contain the following characters: \:*?"&lt;&gt; '{}.</li><li>• Start with a slash (/).</li><li>• Start or end with a period (.).</li><li>• Contain more than two consecutive slashes (/) or periods (.).</li><li>• Exceed 512 characters.</li></ul> |


| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IAM Agency | <p>To back up data to an OBS bucket, you must have the write permission to it. By configuring an IAM agency, you can authorize CSS to access its OBS resources through an associated account.</p> <ul style="list-style-type: none"><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-obs-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the OBS Administrator permissions deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/><pre>"obs:bucket:GetBucketLocation",<br/>"obs:object:GetObjectVersion",<br/>"obs:object:GetObject",<br/>"obs:object:DeleteObject",<br/>"obs:bucket:HeadBucket",<br/>"obs:bucket:GetBucketStoragePolicy",<br/>"obs:object:DeleteObjectVersion",<br/>"obs:bucket:ListBucketVersions",<br/>"obs:bucket:ListBucket",<br/>"obs:object:PutObject"</pre></li><li>When OBS buckets use SSE-KMS encryption, the IAM agency must be granted KMS permissions. You can click <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b> to have the following custom policies created automatically.<br/><pre>"kms:cmk:create",<br/>"kms:dek:create",<br/>"kms:cmk:get",<br/>"kms:dek:decrypt",<br/>"kms:cmk:list"</pre></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</pre></li><li>To use an IAM agency, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</pre></li></ul> |

| Parameter | Description                                                                                                                                             |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
|           | <b>WARNING</b><br>The agency name can contain only letters (case-sensitive), digits, underscores (_), and hyphens (-). Otherwise, the backup will fail. |

Determine whether to enable **Automatic backup** based on service needs.

- To enable automatic, periodic log backup, select **Automatic backup**, and set the backup policy. For more information, see [Table 3-179](#).
  - If manual log backup is sufficient, deselect **Automatic backup**.
- b. Click **OK** to enable log backup.
- The configuration information will be displayed on the **Log Backup** tab.
6. Back up logs. Two options are available: automatic or manual.
- Automatic log backup: Logs are backed up periodically based on the preset policy.
- a. On the **Log Backup** page, click **Modify Settings** on the right.

**Figure 3-91** Modifying Settings



- b. In the **Modify Settings** dialog box, select **Automatic backup**, and configure the automatic backup policy.

**Table 3-179** Automatic backup settings

| Parameter         | Description                                                                                                                                                                     |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time Zone         | Select a time zone for the backup start time.                                                                                                                                   |
| Backup Start Time | Specify the start time of auto backup.<br>Select a value from the drop-down list. The value range is from <b>00:00</b> to <b>23:00</b> . The backup always happens on the hour. |


- c. Click **OK** to enable automatic log backup.
- Automatic log backup tasks are displayed in the task list. When **Task Status** in the task list changes to **Succeeded**, the backup is successful.
- If log backup fails, click **Failed Tasks** learn the failure causes. A maximum of 10 failed tasks can be displayed. When log backup is disabled or the cluster is deleted, the failure records are also cleared.

- d. If automatic log backup is not required, click **Modify Settings**. In the displayed dialog box, deselect **Automatic backup**, and click **OK**. In the displayed confirmation dialog box, click **OK**.
- Disabling log backup does not automatically delete log backups automatically generated earlier. Instead, you need to manually delete them on the OBS console.

Manual log backup: Back up logs right away.

- a. On the **Log Backup** page, click **Back Up Manually** under **Log Backup Tasks**.

Figure 3-92 Manually backing up logs



- b. In the displayed dialog box, confirm the log backup path and click **OK**.
- The log backup task is displayed in the task list. When **Task Status** in the task list changes to **Succeeded**, the backup is successful.
- If log backup fails, click **Failed Tasks** learn the failure causes. A maximum of 10 failed tasks can be displayed. When log backup is disabled or the cluster is deleted, the failure records are also cleared.
7. Check the backed-up log files.
- Logs are backed up incrementally. After the backup is successful, you can access the target OBS bucket to obtain the full log files by clicking **Log Path**.

Figure 3-93 OBS bucket address




Table 3-180 lists the log types, where **clustername** indicates the cluster name.

Table 3-180 Log types

| Log Name                                       | Description            |
|------------------------------------------------|------------------------|
| <i>clustername</i> _deprecation.log            | Deprecation log file   |
| <i>clustername</i> _index_indexing_slowlog.log | Slow indexing log file |

| Log Name                                    | Description         |
|---------------------------------------------|---------------------|
| <i>clustername_index_search_slowlog.log</i> | Slow query log file |
| <i>clustername.log</i>                      | Run log file        |
| <i>clustername_access.log</i>               | Access log file     |

- If the log backup function is no longer needed, you can disable it.  
On the **Log Backup** page, click **Disable Backup**. In the displayed dialog box, click **OK**. Disabling log backup does not automatically delete existing log backups. Instead, you need to manually delete them on the OBS console.

## Log Ingestion

With log ingestion, you can store the real-time logs of a cluster in itself or a different cluster on the same network, facilitating log search and analytics using OpenSearch Dashboards.

Log ingestion is supported only when the OpenSearch cluster version is 1.3.6 or 2.19.0 and the image version is no earlier than x.x.x\_24.2.0\_x.x.x.

- Log in to the [CSS management console](#).
- In the navigation pane on the left, choose **Clusters > OpenSearch**.
- In the cluster list, click the name of the target cluster. The cluster information page is displayed.
- Choose **Logs > Log Ingestion**. The **Log Ingestion** page is displayed.

---

### CAUTION

If the **Log Ingestion** tab is not displayed, the cluster does not support log ingestion.

---

- Enable log ingestion. If it is already enabled, skip this step.
  - Click **Enable Log Ingestion**. In the displayed dialog box, configure necessary settings.

**Table 3-181** Log ingestion settings

| Parameter         | Description                                                                                                                                                                                                                                                                                                               |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Index Prefix Name | If you set a prefix for the log file indexes, the log index names will use the format <b>index prefix name + log ingestion date</b> . The unit of log ingestion is days.<br><br>An index prefix name is a string of 1 to 128 characters. It can contain only digits, lowercase letters, underscores (_), and hyphens (-). |



| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Retention Period    | Set the log retention period. Ingested logs will be deleted upon expiration of this period.<br>Value range: 1 to 3650<br>Unit: days                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Log Storage Cluster | Select a cluster to store ingested logs. Options include <b>Current cluster</b> and <b>Other clusters</b> . <ul style="list-style-type: none"><li>• <b>Current cluster</b>: Store ingested logs in the current cluster.</li><li>• <b>Other clusters</b>: Store ingested logs in another cluster.<br/>Select a cluster from the <b>Cluster</b> drop-down list to use as the log storage destination. Make sure this cluster is in the same VPC and have the same version as the current cluster.<br/>When you select another cluster, click <b>Check</b> to verify network connectivity between the current cluster and the selected other cluster. Log ingestion can start only when <b>The current cluster is available</b> is displayed.</li></ul> |

- b. Click **OK** to enable log ingestion.

The configuration information will be displayed on the **Log Ingestion** tab.

6. View ingested logs.

If the task status changes to **Running**, log ingestion has started.

**Figure 3-94** Viewing ingested logs

|                                                 |                                            |                                            |                            |
|-------------------------------------------------|--------------------------------------------|--------------------------------------------|----------------------------|
| Index Prefix Name<br>test                       | Retention Period<br>30days                 | Status<br>● Running                        | Log Storage Cluster<br>css |
| Dashboards<br><a href="#">Access Dashboards</a> | Created<br>Aug 07, 2025 16:44:04 GMT+08:00 | Updated<br>Aug 07, 2025 16:44:13 GMT+08:00 |                            |

- Click **Dashboards** to log in to the target cluster and check the logs.
  - Click the cluster name at **Log Storage Cluster** to go to the cluster details page.
7. If you no longer need the log ingestion function, you can disable it to reclaim storage resources.

On the **Log Ingestion** page, click **Disable Log Ingestion**. In the displayed dialog box, click **OK**. Disabling log ingestion does not automatically remove ingested log data. Rather, it will be deleted by the system upon expiration of their retention period. You can also manually delete it from the target cluster before the retention period expires.

## Log Introduction

**Table 3-182** Log introduction

| Log Type           | Description                                                                                                                                                                                                                                                                                | Purpose                                                                                                                                                                       |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Run logs           | Run logs, or main logs, record the cluster status and key information about write and query operations. For example, write logs record operations such as index creation, index mapping update, and write queue exhaustion; and query logs record query queue status and query exceptions. | Check the status and write and query operations of each cluster node, including inter-node connectivity, full GC, index creation or deletion, and cluster-level query errors. |
| Slow indexing logs | Slow indexing logs record indexing operations (such as bulk, index, update, and delete) that took a long time to complete, helping you identify performance bottlenecks.                                                                                                                   | In the case of slow write performance, you can query slow indexing logs to locate the cause.                                                                                  |
| Slow query logs    | Slow query logs record search requests that took a long time to complete. They help you monitor and analyze time-consuming search requests, so you can identify performance bottlenecks, optimize SQL queries, and improve overall system performance.                                     | In the case of slow query performance, you can query slow query logs to locate the cause.                                                                                     |
| Deprecation logs   | Deprecation logs record deprecation warnings. Deprecation warnings are written to this log when you use APIs, configurations, or functions that are marked for removal in future versions.                                                                                                 | Check for APIs or features that are about to expire in future versions.                                                                                                       |
| Access logs        | Access logs record cluster access requests, such as the request path and source address.<br><br>You cannot check access logs on the console. To check them, you need to back them up to an OBS bucket or transfer them to a target cluster first.                                          | If there is a surge in service requests, you can analyze the request sources and paths by checking the access logs.                                                           |

- **Run log description**

Run logs record the cluster status and key information about write and query operations. For example, the log record below indicates that an index named **test** was created and afterwards the cluster status changed from **YELLOW** to **GREEN**.

**Figure 3-95** Run log sample

```
[2025-07-04T07:15:55,976][INFO][o.e.c.m.MetadataCreateIndexService] [css-ess-esn-3-1] [test] creating index, cause [auto(bulk api)], templates [], shards [1]/[1]
[2025-07-04T07:15:56,351][INFO][o.e.c.m.MetadataMappingService] [css-ess-esn-3-1] [test/KTXEJL0cSYfC1wy3AVrqA] create_mapping [_doc]
[2025-07-04T07:15:56,720][INFO][o.e.c.r.a.AllocationService] [css-ess-esn-3-1] Cluster health status changed from [YELLOW] to [GREEN] (reason: [shards started [[test]
]([0])]).
```

Log content:

- 1. Log generation time
- 2. Log level, which can be DEBUG, INFO, WARN, or ERROR
- 3. Log-generating module
- 4. Name of the log-generating node
- 5. Log content

- **Slow indexing log description**

Slow indexing logs record indexing operations that took a long time to complete. For example, the log record below shows an indexing request that lasted longer than the preset threshold. The log contains the index name, duration, and request content.

**Figure 3-96** Slow indexing log sample

```
[2025-07-04T07:24:46,898][INFO][i.i.s.index] [css-ess-esn-3-1] [test/KTXEJL0cSYfC1wy3AVrqA] took[1.5ms], took_millis[1], type[_doc], id[INFS1JcB1WhAeIP
M057R], routing[], source[{"message":"[scheduler] end","micro_service":"esim-mgmt","@timestamp":"2025-07-04T00:00:17.087Z"}]
```

Log content:

- 1. Log generation time
- 2. Log level, which can be DEBUG, INFO, WARN, or ERROR
- 3. Log-generating module
- 4. Name of the log-generating node
- 5. Index name and ID
- 6. Log content. In this example, the log recorded the request execution duration, index type, and index request body.

- **Slow query log description**

Slow query logs record search requests that took a long time to complete. For example, the log record below shows a search request that lasted longer than the preset threshold. The log contains the index name, duration, and request content.

**Figure 3-97** Slow query log sample

```
[2025-07-04T07:24:51,289][WARN][i.s.s.query] [css-ess-esn-2-1] [test][0] took[3.6ms], took_millis[3], total_hits[42 hits], types[], stats[], search_type
[QUERY_THEN_FETCH], total_shards[1], source[{"query":{"match":{"message":{"query":"[scheduler]","operator":"OR","prefix_length":0,"max_expansions":50,"fuzzy_transpositions":true,
lenient:false,"zero_terms_query":"NONE","auto_generate_synonyms_phrase_query":true,"boost":1.0}}}], id[]]
```

Log content:

- 1. Log generation time
  - 2. Log level, which can be DEBUG, INFO, WARN, or ERROR
  - 3. Log-generating module
  - 4. Name of the log-generating node
  - 5. Index name and shard ID
  - 6. Log content. In this example, the log recorded the query duration, number of hits, and query request body.
- **Deprecation log description**

Deprecation logs record deprecation warnings. For example, the log record below indicates that **GET /\_cat/master** has been deprecated and should be replaced with **GET /\_cat/cluster\_manager**.

**Figure 3-98** Deprecation log sample

```
[2025-06-27T06:47:12.185][DEPRECATION][o.o.d.r.RestController] [css-c355-ess-esn-1-1] [GET /_cat/master] is deprecated! Use [GET /_cat/cluster_manager] instead.
```

Log content:

- 1. Log generation time
  - 2. Log level, which can only be DEPRECATION.
  - 3. Log-generating module
  - 4. Name of the log-generating node
  - 5. Log content
- **Access log description**

Access logs record cluster access requests and source addresses. For example, the log record below has recorded source information for the **/\_snapshot/my\_backup/my\_snapshot/\_restore?pretty=true** operation.

**Figure 3-99** Access log sample

```
2025-07-09 08:35:00.815[elasticsearch(css-f557-ess-esn-1-1)][business_worker][T#3][INFO][357535][null][POST][_snapshot/my_backup/my_snapshot/_restore?pretty=true] received request from [Netty4#HttpChannel{localAddress=10.10.1.251:9200, remoteAddress=10.10.1.251:38257}]
```

Log content:

- 1. Log generation time
- 2. Name of the log-generating node
- 3. Name of the log-generating thread
- 4. Log level, which can be DEBUG, INFO, WARN, or ERROR.
- 5. Log request method
- 6. Request path
- 7. Source and destination addresses of the request

## 3.16 Viewing OpenSearch Cluster Audit Logs

CTS can record operations associated with CSS for subsequent query, auditing, and backtracking purposes.

## Prerequisites

CTS has been enabled.

## Key Operations Recorded by CTS

**Table 3-183** Key operations recorded by CTS

| Operation                                                                 | Resource Type | Event Name                     |
|---------------------------------------------------------------------------|---------------|--------------------------------|
| Creating a cluster                                                        | cluster       | createCluster                  |
| Changing the billing mode of a cluster from pay-per-use to yearly/monthly | cluster       | updateOndemandClusterTo-Period |
| Querying cluster details                                                  | cluster       | showClusterDetail              |
| Deleting a cluster                                                        | cluster       | deleteCluster                  |
| Renaming a cluster                                                        | cluster       | updateClusterName              |
| Querying the cluster list                                                 | cluster       | listClusters                   |
| Changing the password of a cluster                                        | cluster       | resetPassword                  |
| Restarting a cluster                                                      | cluster       | restartCluster                 |
| Expanding cluster capacity                                                | cluster       | updateExtendCluster            |
| Adding nodes and expanding storage                                        | cluster       | updateExtendInstanceStorage    |
| Changing the node specifications                                          | cluster       | updateFlavor                   |
| Obtaining the node specifications list                                    | cluster       | listFlavors                    |
| Querying all tags                                                         | cluster       | listClustersTags               |
| Querying the tags of a specified cluster                                  | cluster       | showClusterTag                 |
| Adding tags for a cluster                                                 | cluster       | createClustersTags             |
| Deleting the tags of a cluster                                            | cluster       | deleteClustersTags             |
| Adding or deleting cluster tags in batches                                | cluster       | updateBatchClustersTags        |
| Removing specified nodes                                                  | cluster       | updateShrinkNodes              |

| Operation                                             | Resource Type | Event Name            |
|-------------------------------------------------------|---------------|-----------------------|
| Changing the flavor of a specified node type          | cluster       | updateFlavorByType    |
| Removing nodes of a specific type                     | cluster       | updateShrinkCluster   |
| Downloading a security certificate                    | cluster       | downloadCert          |
| Replacing nodes                                       | cluster       | updateInstance        |
| Changing the security mode                            | cluster       | changeMode            |
| Adding dedicated Master or Client nodes               | cluster       | addIndependentNode    |
| Upgrading the cluster kernel                          | cluster       | upgradeCore           |
| Obtaining a target image ID                           | cluster       | listImages            |
| Obtaining upgrade details                             | cluster       | upgradeDetail         |
| Retrying a failed upgrade task                        | cluster       | retryUpgradeTask      |
| Changing the security group                           | cluster       | changeSecurityGroup   |
| Creating a cluster (V2)                               | cluster       | createClusterV2       |
| Restarting a cluster (V2)                             | cluster       | restartCluster        |
| Rolling restart                                       | cluster       | rollingRestart        |
| Loading a custom word dictionary                      | cluster       | createLoadIkThesaurus |
| Viewing custom word dictionary configurations         | cluster       | showIkThesaurus       |
| Deleting a custom word dictionary                     | cluster       | deleteIkThesaurus     |
| Enabling OpenSearch Dashboards public network access  | cluster       | startKibanaPublic     |
| Disabling OpenSearch Dashboards public network access | cluster       | updateCloseKibana     |

| Operation                                                          | Resource Type | Event Name                  |
|--------------------------------------------------------------------|---------------|-----------------------------|
| Modifying the OpenSearch Dashboards public network bandwidth       | cluster       | updateAlterKibana           |
| Modifying OpenSearch Dashboards public network access control      | cluster       | updatePublicKibanaWhitelist |
| Disabling OpenSearch Dashboards public network access control      | cluster       | stopPublicKibanaWhitelist   |
| Enabling logging                                                   | cluster       | startLogs                   |
| Disabling logging                                                  | cluster       | stopLogs                    |
| Querying the job list                                              | cluster       | listLogsJob                 |
| Querying basic log configurations                                  | cluster       | showGetLogSetting           |
| Modifying basic log configurations                                 | cluster       | updateLogSetting            |
| Enabling automatic log backup                                      | cluster       | startLogAutoBackupPolicy    |
| Disabling automatic log backup                                     | cluster       | stopLogAutoBackupPolicy     |
| Backing up logs                                                    | cluster       | createLogBackup             |
| Querying logs                                                      | cluster       | showLogBackup               |
| Enabling public network access                                     | cluster       | createBindPublic            |
| Disabling public network access                                    | cluster       | updateUnbindPublic          |
| Modifying public network bandwidth                                 | cluster       | updatePublicBandWidth       |
| Enabling the public network access whitelist                       | cluster       | startPublicWhitelist        |
| Disabling the public network access whitelist                      | cluster       | stopPublicWhitelist         |
| Automatically configuring the basic settings of a cluster snapshot | cluster       | startAutoSetting            |

| Operation                                             | Resource Type | Event Name               |
|-------------------------------------------------------|---------------|--------------------------|
| Modifying the basic settings of a cluster snapshot    | cluster       | updateSnapshotSetting    |
| Manually creating a snapshot                          | snapshot      | createSnapshot           |
| Restoring a snapshot                                  | snapshot      | restoreSnapshot          |
| Deleting a specified snapshot                         | snapshot      | deleteSnapshot           |
| Setting an automatic snapshot creation policy         | cluster       | createAutoCreatePolicy   |
| Querying automatic snapshot creation policies         | cluster       | showAutoCreatePolicy     |
| Querying the snapshot list                            | cluster       | listSnapshots            |
| Disabling snapshots                                   | cluster       | stopSnapshot             |
| Enabling automatic snapshot creation                  | cluster       | startAutoCreateSnapshots |
| Disabling automatic snapshot creation                 | cluster       | stopAutoCreateSnapshots  |
| Enabling VPC Endpoint                                 | cluster       | startVpecp               |
| Disabling VPC Endpoint                                | cluster       | stopVpecp                |
| Obtaining a VPCEP connection                          | cluster       | showVpcepConnection      |
| Updating a VPCEP connection                           | cluster       | updateVpcepConnection    |
| Modifying the VPCEP whitelist                         | cluster       | updateVpcepWhitelist     |
| Modifying parameter settings                          | cluster       | listYmls                 |
| Obtaining the task list of parameter settings         | cluster       | listYmlsJob              |
| Obtaining the parameter settings list                 | cluster       | updateYmls               |
| Querying ELB V3 load balancers supported by a cluster | cluster       | listElbs                 |



| Operation                                                 | Resource Type | Event Name         |
|-----------------------------------------------------------|---------------|--------------------|
| Enabling or disabling the Elasticsearch Load Balancer     | cluster       | enableOrDisableElb |
| Configuring an Elasticsearch listener                     | cluster       | createElbListener  |
| Obtaining information about an Elasticsearch ELB          | cluster       | showElbDetail      |
| Updating an Elasticsearch listener                        | cluster       | updateEsListener   |
| Querying the certificate list                             | cluster       | listElbCerts       |
| Obtaining the intelligent O&M task list and details       | cluster       | listAiOps          |
| Creating a cluster detection task                         | cluster       | createAiOps        |
| Deleting a detection task record                          | cluster       | deleteAiOps        |
| Obtaining SMN topics available for intelligent O&M alarms | cluster       | listSmnTopics      |

## Querying Real-Time Traces

After a management tracker is created on the CTS console, the system starts recording operations performed on cloud service resources. After a data tracker is created, the system starts recording operations performed on data in OBS buckets. CTS retains operation records generated in the latest seven days.

To view or export operation records of the last seven days on the CTS console, see [Querying Real-Time Traces](#).

# 4 Using Logstash for Data Migration

## 4.1 Procedure for Using Logstash

Table 4-1 Procedure for using a Logstash cluster

| Cate<br>gory | Operation                                                      |                                                                                                                                                                                                                                                    | Details                                                                                                      |
|--------------|----------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| Use          | Plann<br>ing a<br>cluste<br>r                                  | Before creating a Logstash cluster, develop a plan for it, such as whether to deploy the cluster across multiple AZs to improve availability; and the node quantity and specifications in order to ensure the desired performance and reliability. | <a href="#">Logstash Cluster Planning Suggestions</a>                                                        |
|              | Creati<br>ng a<br>cluste<br>r                                  | Create a Logstash cluster based on the plan.                                                                                                                                                                                                       | <a href="#">Creating a Logstash Cluster (New Version)</a>                                                    |
|              | Confi<br>gurin<br>g<br>route<br>s                              | If the subnets of the Logstash cluster and its data source are under different NICs, you need to configure routes to connect the two.                                                                                                              | <a href="#">Configuring Routes for a Logstash Cluster</a>                                                    |
|              | Confi<br>gurin<br>g a<br>Logst<br>ash<br>migra<br>tion<br>task | Through the configuration center of the Logstash cluster, you can create and start a Logstash configuration file to migrate data between different locations.                                                                                      | <a href="#">Creating a Logstash Configuration File</a><br><a href="#">Starting a Logstash Migration Task</a> |

| Category | Operation                     | Details                                                                                                                                                                                                                                                                                                                                                              |
|----------|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| O&M      | Scaling a cluster             | <a href="#">Scaling Out a Logstash Cluster</a><br><a href="#">Scaling In a Logstash Cluster</a>                                                                                                                                                                                                                                                                      |
|          | Managing clusters             | <a href="#">Viewing Logstash Cluster Information</a><br><a href="#">Setting Tags for a Logstash Cluster</a><br><a href="#">Associating a Logstash Cluster with an Enterprise Project</a><br><a href="#">Managing CA Certificates for a Logstash Cluster</a><br><a href="#">Forcibly Restarting a Logstash Cluster</a><br><a href="#">Deleting a Logstash Cluster</a> |
|          | Monitoring and log management | <a href="#">Monitoring Metrics for Logstash Clusters in Cloud Eye</a><br><a href="#">Using Cloud Eye to Monitor Logstash Clusters</a><br><a href="#">Managing Logstash Cluster Logs</a>                                                                                                                                                                              |
|          | Audit logs                    | <a href="#">Viewing Logstash Cluster Audit Logs</a>                                                                                                                                                                                                                                                                                                                  |

## 4.2 Logstash Cluster Planning Suggestions

## 4.2.1 Planning Cluster AZs and HA

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

Multi-AZ deployment is a high availability feature provided by CSS. Deploying a cluster across two AZs located in the same region can help prevent data loss and lower the possibility of service outages.

If you select multi-AZ deployment when creating a cluster, CSS automatically enables cross-AZ HA to ensure that cluster nodes will be evenly distributed across the selected AZs (the difference between the number of nodes in each AZ cannot exceed 1).

### Suggestions on Multi-AZ Deployment

If you select multi-AZ deployment when creating a cluster, CSS automatically enables cross-AZ HA to ensure that cluster nodes will be evenly distributed across the selected AZs (the difference between the number of nodes in each AZ cannot exceed 1).

A Logstash cluster can span a maximum of two AZs.

### Node Distribution Rules

When a multi-AZ cluster is created, nodes of all types are evenly distributed across different AZs. A maximum of two AZs can be used. [Table 4-2](#) shows the node distribution when the number of AZs varies.



#### WARNING

When creating a multi-AZ cluster, select at least two nodes. Otherwise, multi-AZ cluster deployment will fail.

**Table 4-2** Node distribution when the numbers of nodes and AZs vary

| Node Quantity | Single AZ | Two AZs       |     |
|---------------|-----------|---------------|-----|
|               | AZ1       | AZ1           | AZ2 |
| 1             | 1         | Not supported |     |
| 2             | 2         | 1             | 1   |
| 3             | 3         | 2             | 1   |
| 4             | 4         | 2             | 2   |
| ...           | ...       | ...           | ... |

## 4.2.2 Planning Node Specifications and Capacity

This topic provides suggestions on selecting node specifications and configuring the node storage type, storage capacity, and node quantity for a Logstash cluster, helping you properly plan the capacities of your cluster.

### Node Configuration Suggestions

**Table 4-3** Node configuration

| Parameter                      | Configuration Suggestions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node Specifications            | <p>In the node flavor list, <b>vCPUs   Memory</b> indicate the number of vCPUs and memory capacity available for each flavor, and <b>Recommended Storage</b> indicates the supported storage capacity range. We recommend that you select node specifications based on core metrics of your Logstash cluster, such as the CPU load, memory requirements, and I/O features.</p> <p><a href="#">Logstash Node Specifications</a> describes the application scenarios and core features of different node specifications. It can help you properly plan your cluster.</p> <p>For more information about different node specifications, see <a href="#">ECS Types</a>.</p>                                                                                                                                                                                                               |
| Node Storage Type and Capacity | <p>Select an appropriate storage type and capacity for cluster nodes.</p> <ul style="list-style-type: none"><li>• Select an EVS disk type:<ul style="list-style-type: none"><li>– If you select <b>persisted</b> (disk storage) for event buffering, you are advised to select EVS disks with better performance, such as <b>Ultra-high I/O</b> and <b>Extreme SSD</b>.</li><li>– If you use memory queues, select high I/O EVS disks.</li></ul></li></ul> <p><b>NOTE</b><br/>By default, Logstash uses in-memory bounded queues to buffer events. The corresponding configuration item is <b>queue.type</b>, which is available when you configure pipelines in the configuration file.</p> <p>For more on EVS disk performance, see <a href="#">Disk Types and Performance</a>.</p> <ul style="list-style-type: none"><li>• The node storage capacity is fixed at 40 GB.</li></ul> |

| Parameter | Configuration Suggestions                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nodes     | <p>The number of nodes in a Logstash cluster ranges from 1 to 100.</p> <p>Logstash nodes are used to ingest, parse, process, and transfer data. The number of nodes determines the data migration speed. Select the number of Logstash nodes based on service requirements.</p> <p>When the Logstash cluster has two or more nodes, all nodes use the same configuration files. This mode works when Logstash is a consumer of Kafka data.</p> |

## Logstash Node Specifications

Logstash nodes support EVS disks only. EVS disks are a virtual block storage service that is independent of ECSs. They provide high reliability and fast elasticity, making them ideal for workloads that require high data reliability and highly scalable storage capacity.

The table below describes the application scenarios and core features of different node specifications. It can help you properly plan your cluster.

**Table 4-4** Comparing different node specifications

| CPU Architecture | Node Flavor       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| x86              | Compute-intensive | <p><b>Core advantages</b></p> <ul style="list-style-type: none"><li>• High-performance CPU: designed for high computational load, ideal for CPU-intensive tasks.</li><li>• Optimized network I/O: supports high throughput in both the inbound and outbound directions (such as network plugins).</li></ul> <p><b>Application scenarios</b></p> <ul style="list-style-type: none"><li>• CPU-intensive plugins: plugins that involve heavy CPU computation, such as grok (regular expression parsing) and dissect (structured log parsing).</li><li>• Hybrid load tasks: tasks that involve both heavy CPU computation and network I/O loads, such as real-time log ingestion and data cleaning.</li><li>• Large-scale data processing at high speed, such as log aggregation and event stream processing.</li></ul> <p><b>Precautions</b></p> <ul style="list-style-type: none"><li>• Set the number of task threads (pipeline.workers) to equal the number of vCPUs. This optimizes CPU utilization while avoiding resource contention.</li><li>• Pay attention to warnings on I/O performance bottlenecks. If the load of network plugins (such as beats and http) is heavy, make sure there is sufficient network bandwidth.</li></ul> |

| CPU Architecture | Node Flavor                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                  | General computing-plus - AC | <p><b>Core advantages</b></p> <ul style="list-style-type: none"><li>• Dedicated CPU: no resource contention between different instances, stable performance at a relatively low cost, ideal for high-priority tasks.</li><li>• Low latency: guaranteed efficiency for CPU-intensive plugins.</li></ul> <p><b>Application scenarios</b></p> <ul style="list-style-type: none"><li>• High CPU-load tasks: real-time log parsing and complex field extraction (grok and ruby)</li><li>• Mission-critical service processes: reliable performance required (such as financial transaction log processing)</li><li>• Multi-thread processing: high-concurrency tasks</li></ul> <p><b>Precautions</b></p> <p>If tasks mainly involve high network I/O, you are advised to use ultra-high I/O disks.</p>                                                                |
|                  | General computing           | <p><b>Core advantages</b></p> <p>Balanced configuration: default specifications, suitable for medium-scale data processing tasks.</p> <p><b>Application scenarios</b></p> <ul style="list-style-type: none"><li>• Medium-scale log processing, such as enterprise log ingestion and monitoring data aggregation.</li><li>• Low CPU-load tasks: mainly network I/O (such as file and kafka plugins).</li><li>• Standard deployment: meets the needs of general use cases, no need for special tuning.</li></ul> <p><b>Memory usage evaluation</b></p> <p>Estimate the needed memory capacity using this formula: <b>Average size of each piece of data processed by Logstash x (pipeline.workers x pipeline.batch.size)</b></p> <p>Example: If the average data size is 1 KB, pipeline.workers = 4, and pipeline.batch.size = 1000, the memory size is ~4 MB.</p> |



| CPU Architecture | Node Flavor               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                  | Memory-optimized          | <p><b>Core advantages</b></p> <ul style="list-style-type: none"><li>• Large memory capacity: suitable for memory-intensive tasks, reduced disk I/O pressure.</li><li>• Optimized memory queues: data cached in memory queues, more efficient data processing.</li></ul> <p><b>Application scenarios</b></p> <ul style="list-style-type: none"><li>• Large-scale log aggregation: log analytics platform, security information and event management (SIEM), etc.</li><li>• Complex data transformation: tasks that involve temporary storage of large amounts of data (for example, the aggregate plugin).</li></ul> <p><b>Precautions</b></p> <ul style="list-style-type: none"><li>• Monitor the memory usage in real time to avoid out of memory (OOM).</li><li>• This cost is high. Use this flavor for memory-intensive tasks that have high priorities.</li></ul> |
| Kunpeng          | Kunpeng general computing | <p><b>Core advantages</b></p> <ul style="list-style-type: none"><li>• Cost-effective: In general, the Arm architecture is more cost-effective than x86, with lower power consumption.</li><li>• Memory-efficient: Use this flavor for workloads that are cost-sensitive and require large memory capacities.</li></ul> <p><b>Application scenarios</b></p> <ul style="list-style-type: none"><li>• Cost-sensitive workloads, such as log ingestion for small and medium enterprises and testing environments.</li><li>• Arm ecosystem compatibility: use when your workloads must be compatible with Kunpeng servers.</li></ul> <p><b>Precautions</b></p> <p>Make sure your Logstash plugins (including Java virtual machine and third-party plugins) are compatible with the Arm architecture.</p>                                                                    |

## 4.3 Creating a Logstash Cluster (New Version)

Logstash clusters created in CSS can be used for efficient ingestion, parsing, processing, and transmission of vast quantities of data. They support real-time log analytics, data integration and transformation, and cross-system data transfers while providing capabilities for data monitoring, alarming, as well as in-depth analysis and visualization.

You can create a Logstash cluster in either of the following ways:

- Method 1: Create a Logstash cluster on the CSS management console. This topic provides the operation guide for the default UI (that is, the new UI). For the operation guide for the old UI, see [Creating a Logstash Cluster \(Old Version\)](#).

### NOTE

To improve efficiency in creating clusters, CSS has made several enhancements on the cluster creation page in terms of ease-of-use. Being more intuitive, the new UI streamlines cluster creation procedures.

- Method 2: Create a Logstash cluster using CSS APIs. For details, see [Creating a Cluster](#).


## Prerequisites

- You have planned the Logstash cluster configuration by following the guidelines in [Logstash Cluster Planning Suggestions](#), including the cluster size, node configuration, and storage capacity requirements.
- You have created a VPC, subnet, and security group beforehand, or you can create them when creating the cluster. For more information, see [Virtual Private Cloud User Guide](#).

## Opening the Create Cluster Page

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the upper right corner, click **Create Cluster**. The new-version UI for creating a cluster is displayed by default.

**Figure 4-1** Create Cluster (new version)



## Cluster Configuration

Select the cluster type and version.

**Table 4-5** Cluster Configuration

| Parameter       | Description                                       |
|-----------------|---------------------------------------------------|
| Cluster Type    | Select <b>Logstash</b> .                          |
| Cluster Version | Select a cluster version from the drop-down list. |

## Basic Settings

Select the region, billing mode, and AZ of the cluster.

**Table 4-6** Basic settings

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                       |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Region    | Select the region where the cluster is located. A region is the location of a physical data center. Regions are defined based on their geographical location and network latency. For lower network latency and quicker resource access, select the nearest region.                                                                                                               |
| AZ        | Select one or more AZs associated with the cluster region. An AZ is a physical region where resources use independent power supplies and networks. AZs are physically isolated but interconnected through an internal network.<br><br>A maximum of two AZs can be configured. For details about the use of multiple AZs, see <a href="#">Suggestions on Multi-AZ Deployment</a> . |

| Parameter    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Billing Mode | <p>Billing mode for the cluster, which can be <b>Yearly/Monthly</b> or <b>Pay-per-use</b>.</p> <ul style="list-style-type: none"><li>• <b>Yearly/Monthly</b>: You prepay for a yearly or monthly subscription.<br/>This mode is cost-effective when the duration of your service usage is predictable. It is recommended for long-term users. For more information, see <a href="#">Yearly/Monthly</a>.<br/>A yearly/monthly subscription cannot be deleted. If you no longer need it, unsubscribe from it.</li><li>• <b>Pay-per-use</b> (postpaid): You will be billed hourly by actual duration of use. Any partial hour of usage will be rounded up to one hour.<br/>This mode is ideal for short-term or unpredictable workloads. For more information, see <a href="#">Pay-per-Use</a>.<br/>Pay-per-use resources can be provisioned and deleted at any time.</li></ul> |

## Configuring Logstash Nodes

Configure the nodes of the Logstash cluster.

**Table 4-7** Configuring Logstash nodes

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CPU Architecture    | Select the CPU architecture of the Logstash nodes. <b>x86</b> and <b>Kunpeng</b> nodes are supported. The architectures actually supported may vary depending on the regional environment.                                                                                                                                                                                                                                                                                                                                                                                                            |
| Node Specifications | <p>Select the specifications of the Logstash nodes. Click <b>Available</b>. On the displayed page, select a flavor that suits your needs.</p> <p>The nodes of a Logstash cluster support EVS disks only. For more information about different node specifications, see <a href="#">ECS Types</a>.</p> <p>In the node flavor list, <b>vCPUs   Memory</b> indicate the number of vCPUs and memory capacity available for each flavor, and <b>Recommended Storage</b> indicates the supported storage capacity range.</p> <p>The node flavors available may vary depending on the region you select.</p> |

| Parameter                      | Description                                                                                                                                                                                                                                                                                                |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Node Storage Type and Capacity | Select the storage type and capacity of the Logstash nodes. <ul style="list-style-type: none"><li>Available EVS disk types vary depending on your region. For more on EVS disk performance, see <a href="#">Disk Types and Performance</a>.</li><li>The node storage capacity is fixed at 40 GB.</li></ul> |
| Nodes                          | Set the number of nodes for the Logstash cluster.<br>The number of nodes determines the data migration speed. Select the number of Logstash nodes based on service needs.<br>Value range: 1 to 100.                                                                                                        |

## Network Settings

Configure the VPC, subnet, IP addresses, and security group of the cluster.

**Table 4-8** Configuring network settings

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VPC       | <p>Select a VPC for the cluster for proper network isolation.</p> <p>You can select a shared VPC, which is indicated by <b>Shared VPC</b> in the VPC name. To create a CSS cluster in a shared VPC, the <b>default vpc subnet statement</b> permission is required.</p> <p>VPC</p>  <p>VPC sharing allows you to centrally manage resources across multiple accounts, helping to improve resource management efficiency and reduce O&amp;M costs. For more information about shared VPCs, see <a href="#">VPC Sharing</a>.</p> <p>If none of the existing VPCs meets your requirements, click <b>Create VPC</b> to create one. (Make sure you have the permission to create VPCs.) For details, see <a href="#">Creating a VPC and Subnet</a>.</p> |

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Subnet         | <p>Select a subnet for the cluster. A subnet improves network security by providing exclusive network resources that are isolated from other networks.</p> <p>Select a subnet in the current VPC. Or you can select a subnet under a shared VPC. The owner of a VPC can share the VPC to specified users.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| IPv4 Address   | <p>Assign IPv4 addresses to cluster nodes.</p> <ul style="list-style-type: none"><li>• <b>Assign automatically:</b> Private IPv4 addresses will be automatically assigned to cluster nodes.</li><li>• <b>Assign manually:</b> Manually assign private IPv4 addresses to cluster nodes. Before assigning an IP address, click <b>View In-Use IP Address</b> to check IP addresses that are already in use.</li></ul> <p>If you manually specify fewer IP addresses than the total number of cluster nodes, the system will automatically assign random private IPv4 addresses to the remaining nodes.</p> <p>Default value: <b>Assign automatically</b></p>                                                                                           |
| Security Group | <p>Select a security group for the cluster. A security group serves as a virtual firewall that provides access control policies for clusters.</p> <p>To select a security group that meets your requirements, click <b>View Security Group</b> to go to the security group list, where you can check the details of each security group.</p> <p>The selected security group must allow port 9600 in the inbound direction. Otherwise, the cluster may be inaccessible to external services.</p> <p>If none of the existing security groups meets your requirements, click <b>Create Security Group</b> to create one. (Make sure you have the permission to create security groups.) For details, see <a href="#">Creating a Security Group</a>.</p> |

## Cluster Management

Configure the cluster name, enterprise project, and more.

**Table 4-9** Cluster management information

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Name       | User-defined cluster name.<br>The cluster name must start with a letter and can contain 4 to 32 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Add Description    | Add a description for the cluster for easy recognition.<br>Click <b>+ Add Description</b> and then enter a description of the cluster. The value can contain 0 to 128 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Enterprise Project | Associate the cluster with an enterprise project.<br>An enterprise project groups cloud resources, so you can manage resources and members by project. The default project is <b>default</b> .<br>If enterprise projects are enabled, you can select an enterprise project from the drop-down list. You can click <b>View Enterprise Projects</b> to go to the Enterprise Project Management Service console and check enterprise project details. For more information, see <a href="#">Accessing the Enterprise Center</a> .                                                                                                                                                                                                |
| Tags               | Adding tags to clusters helps you identify and manage your cluster resources.<br>Click <b>+ Add</b> . You can select existing tag keys and values from the drop-down lists or enter new ones.<br>If you want to use the same tag to identify multiple cloud resources for better resource grouping, we recommend that you predefine tags in Tag Management Service (TMS). For details, see <a href="#">Predefined Tags</a> .<br>Each cluster can have a maximum of 20 tags.<br>If your organization has configured tag policies for CSS, add tags to clusters based on these policies. If a tag does not comply with the tag policies, cluster creation may fail. Contact the administrator to learn more about tag policies. |

## Quantity

Configure the cluster's subscription duration and auto-renewal option. This parameter is available only when **Billing Mode** is set to **Yearly/Monthly**.

**Table 4-10** Configuring the subscription duration

| Parameter         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Required Duration | Set the duration of your subscription in the range of 1 month to 3 years. If you plan to use the cluster for more than nine months, you should choose a yearly subscription for a better price.                                                                                                                                                                                                                                                                                                   |
| Auto-renew        | <p>Whether to automatically renew the subscription. By default, auto-renewal is disabled. To enable it, select <b>Auto-renew</b>.</p> <p>When <b>Auto-renew</b> is selected, a yearly/monthly subscription is automatically renewed upon expiry.</p> <ul style="list-style-type: none"><li>Monthly subscription: auto-renews for 1 month each time.</li><li>Yearly subscription: auto-renews for 1 year each time.</li></ul> <p>For more information, see <a href="#">Auto-Renewal Rules</a>.</p> |

## Submit

1. In the **Configuration Summary** panel on the right, review the cluster configuration details.  
Any mandatory fields that are missing are displayed in red, and you have to set them in the configuration area.
2. Click **Create Now**.  
Cluster creation time depends on the number of nodes. Typically, this process takes less than 60 minutes, though clusters with a large number of nodes may require additional time.
3. Return to the cluster list and check the newly created cluster.  
If the cluster is created successfully, **Cluster Status** changes to **Available**.

## Related Documents

- For details about how to check basic cluster information, such as the cluster status, private IP address, and node information, see [Viewing Logstash Cluster Information](#).
- For how to create a Logstash configuration file to define data transmission rules, see [Creating a Logstash Configuration File](#).

## 4.4 Creating a Logstash Cluster (Old Version)

You can create a Logstash cluster in either of the following ways:

- Method 1: Create a Logstash cluster on the CSS management console. This topic provides the operation guide for the old UI. For the operation guide for the new UI, see [Creating an OpenSearch Cluster \(New Version\)](#).



**NOTE**

To improve efficiency in creating clusters, CSS has made several enhancements on the cluster creation page in terms of ease-of-use. Being more intuitive, the new UI streamlines cluster creation procedures. You are advised to use the new UI to create Logstash clusters.

- Method 2: Create a Logstash cluster using CSS APIs. For details, see [Creating a Cluster](#).


## Prerequisites

You have planned the Logstash clusters that need to be created by following the instructions in [Logstash Cluster Planning Suggestions](#).

## Creating a Cluster

1. Log in to the [CSS management console](#).
2. On the **Dashboard** page, click **Create Cluster** in the upper right corner. The **Create Cluster** page is displayed.  
Alternatively, choose **Clusters** > **Logstash** in the navigation tree on the left. Click **Create Cluster** in the upper right corner.
3. Click **Back to Old Version** on the right to go to the old UI for cluster creation.

**Figure 4-2** Create Cluster (old version)



4. On the **Basic Configuration** page, configure basic information and resources for the Logstash cluster.

**Table 4-11** Basic configuration of the Logstash cluster

| Parameter    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Billing Mode | Select <b>Yearly/Monthly</b> or <b>Pay-per-use</b> . <ul style="list-style-type: none"><li>• <b>Yearly/monthly</b>: You pay for the cluster by year or month, in advance. The service duration ranges from one month to three years. If you plan to use the cluster for more than nine months, you should choose a yearly subscription for a better price.</li><li>• <b>Pay-per-use</b>: You will be billed hourly by actual duration of use. Any partial hour of usage will be rounded up to one hour.</li></ul> |

| Parameter             | Description                                                                                                                                                                                                         |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Required Duration     | The duration for which the purchased EIP will be used. The duration must be specified if the <b>Billing Mode</b> is set to <b>Yearly/Monthly</b> .<br>Configure automatic renewal if necessary.                     |
| Region                | Select the region where the cluster is located.<br>ECSs in different regions cannot communicate with each other over an intranet. For lower network latency and quicker resource access, select the nearest region. |
| AZ                    | Select AZs associated with the cluster region.<br>A maximum of two AZs can be configured. For details about the use of multiple AZs, see <a href="#">Suggestions on Multi-AZ Deployment</a> .                       |
| Type                  | Select <b>Logstash</b> .                                                                                                                                                                                            |
| Version               | Select a cluster version from the drop-down list.                                                                                                                                                                   |
| Name                  | User-defined cluster name.<br>The cluster name must start with a letter and can contain 4 to 32 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.                                     |
| Cluster Description   | Add a description for the cluster for easy recognition.<br>The value can contain 0 to 128 characters.                                                                                                               |
| Nodes                 | Number of nodes in a cluster, in the range 1 to 100.                                                                                                                                                                |
| CPU Architecture      | <b>x86</b> and <b>Kunpeng</b> are supported. The supported types depend on the actual regional environment.                                                                                                         |
| Node Specifications   | Specifications of nodes in a cluster. You can select node specifications based on your needs. Each cluster supports only one specification. For details, see <a href="#">ECS Types</a> .                            |
| Node Storage Type     | Options include Common I/O, High I/O, Ultra-high I/O, and Extreme SSD.<br><b>NOTE</b><br>If the type of storage in use is not supported, the storage type is not displayed.                                         |
| Node Storage Capacity | It can only be set to 40 GB for a Logstash cluster.                                                                                                                                                                 |

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enterprise Project | <p>When creating a CSS cluster, you can bind an enterprise project to the cluster if you have enabled the enterprise project function.</p> <p>Select an enterprise project from the <b>Enterprise Project</b> drop-down list, or click <b>View Enterprise Project</b> to go to the <b>Enterprise Project Management Service</b> page and check existing enterprise projects.</p> |

5. Click **Next: Network**.
6. On the **Network** page, configure the network and security group settings for the Logstash cluster.

**Table 4-12** Parameter description

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VPC            | <p>Specify a VPC to isolate the cluster's network.</p> <p>Click <b>View VPC</b> to go to the VPC management console and check the created VPCs or VPCs shared with the current account. You can select a shared VPC when creating a cluster, in which case, the <b>default vpc subnet statement</b> permission is required.</p> <p>VPC sharing allows you to centrally manage resources across multiple accounts, helping to improve resource management efficiency and reduce O&amp;M costs. For more information about shared VPCs, see <a href="#">VPC Sharing</a>.</p> <p>If no VPC is available, contact the CSS administrator to create a new VPC. For details, see <a href="#">Creating a VPC and Subnet</a>.</p> <p><b>NOTE</b><br/>The VPC must contain CIDRs. Otherwise, cluster creation will fail. By default, a created VPC contains CIDRs.</p> |
| Subnet         | <p>A subnet provides dedicated network resources that are isolated from other networks, improving network security.</p> <p>Select a subnet needed by the cluster in the current VPC. You may select a subnet in a shared VPC.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Security Group | <p>A security group serves as a virtual firewall that provides access control policies for clusters.</p> <p>Select a security group for the cluster. Click <b>View Security Group</b> to go to the security group list, where you can view details about security groups.</p> <p><b>NOTE</b><br/>To enable cluster access, ensure that port 9600 is allowed by the security group.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

**Figure 4-3** Configuring network specifications

The screenshot shows a configuration interface with three sections:

- VPC:** A dropdown menu with 'vpc-' selected. To the right is a 'View VPC' link. Below it is a descriptive text: 'A VPC is a logically isolated virtual network that facilitates secure internal network management and configurations. Clusters in the same VPC can communicate with each other, but those in different VPCs cannot communicate with each other by default.'
- Subnet:** A dropdown menu with 'subnet-' selected. To the right is a search icon. Below it is a descriptive text: 'A subnet is a range of IP addresses in a VPC.'
- Security Group:** A dropdown menu with 'sg-' selected. To the right is a 'View Security Group' link. Below it is a descriptive text: 'A security group implements access control for associated clusters, providing an additional layer of security. You can add inbound and outbound rules that apply to all clusters associated with a security group.'

7. Click **Next: Advanced Settings**.
8. On the **Advanced Settings** page, configure advanced settings for the Logstash cluster.

If you select **Custom**, you can set tags.

- **Default:** Tags are disabled by default.
- **Custom:** Configure **Tags** for the cluster if necessary.

### Tags

Adding tags to clusters helps you identify and manage your cluster resources. You can configure your own tags or use tags predefined by the Tag Management Service (TMS). If you want to use the same tag to identify multiple cloud resources for better resource grouping, we recommend that you predefine tags in Tag Management Service (TMS). For details, see [Predefined Tags Overview](#).

### NOTE

If your organization has configured tag policies for CSS, add cluster tags based on these policies. Tags that do not comply with predefined tag policies will cause a cluster creation failure. Contact the administrator to learn about the tag policies.


9. Click **Next: Confirm Configuration**. Confirm the settings, and click **Create Now** to create the cluster.
10. Return to the cluster list and check the newly created cluster. If the cluster is created successfully, **Cluster Status** changes to **Available**. Cluster creation time depends on the number of nodes. Typically, this process takes less than 60 minutes, though clusters with a large number of nodes may require additional time.  
  
If cluster creation fails, try creating the cluster again by rectifying the errors returned.

## 4.5 Configuring Networking for a Logstash Cluster

### 4.5.1 Configuring Routes for a Logstash Cluster

By default, a CSS Logstash cluster (server) cannot access data sources that are located in a different VPC—it cannot pull data from or push data to the data sources. To enable it, you need to configure routes for the Logstash cluster that point to the data sources.

**Figure 4-4** Configuring routes for a Logstash cluster



1. Establish a VPC peering connection to connect two VPCs.
2. On CSS, configure routes for the Logstash cluster that point to the data sources, enabling the cluster to pull data from or push data to the data sources.
3. Verify network connectivity between the data sources and Logstash.

Configuring routes to enable cluster communication offers the following advantages:

- Improved cross-VPC communication performance: A route table specifies the traffic forwarding paths, so that traffic does not have to pass through the public network.
- Secure communication: Private networks can communicate through VPC peering connections, which is more secure than transmitting data over the public network.
- Flexible access control policies: You can customize access permissions for specific subnets.
- Support for a hybrid cloud architecture: Services can communicate across VPCs in different environments.

## Constraints

Adding cluster routes will alter network connections, possibly causing disruptions in the network. Exercise caution.

## Prerequisites

- The Logstash cluster is Available, and port 9600 is allowed in the inbound direction by the cluster's security group.

- The data sources are available, and the security group used by the server of the data sources allows the necessary ports the inbound direction. For example, with an Elasticsearch cluster, port 9200 must be allowed.

## Establishing a VPC Peering Connection

Before configuring routes for a Logstash cluster, establish a connection between the cluster's VPC and that of the data sources, that is, a VPC peering connection.

- If the Logstash cluster and the data sources belong to the same account, see [Creating a VPC Peering Connection to Connect Two VPCs in the Same Account](#).
- If the Logstash cluster and the data sources belong to different accounts, see [Creating a VPC Peering Connection to Connect Two VPCs in Different Accounts](#).

Set the peer VPC to the VPC where the Logstash cluster is located and the local VPC to the VPC where the data sources are located.

When adding routes for a VPC peering connection, configure the destination addresses based on service requirements.

If the Logstash cluster and the data sources share the same security group, there is no need to configure security group rules for the instances in the VPC at either end of the VPC peering connection.

## Configuring Routes for a Logstash Cluster

On CSS, configure routes for the Logstash cluster that point to the data sources, enabling the cluster to pull data from or push data to the data sources.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Overview** tab. In the **Configuration** area, click **Add Route** next to **Cluster Route**.

**Figure 4-5** Adding a route



5. In the displayed dialog box, configure the route information.

**Table 4-13** Adding a route

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP Address  | Enter the IP address (or a CIDR block) of a data source, for example, <b>10.10.1.0</b> .<br>If the IP addresses of data sources are scattered, configure multiple routes for them.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Subnet Mask | Enter the subnet mask used with the IP address.<br>The subnet mask must align with the network part of the IP address. That is: <ul style="list-style-type: none"><li>• The network bits (the 1s) of the subnet mask must encompass those of the IP address.</li><li>• The host bits (the 0s) of the subnet mask must accommodate the IP address's host bits.</li></ul> The subnet mask should not cover too wide a range, avoiding unnecessary IP address exposure or waste of the address space. 255.255.255.0 (that is, /24) is recommended, which applies to most internal networking scenarios. |

6. Click **OK** to add the route.

After the route is added, click **View Route** to check the route information. You can click **Add Route** to add another route; or click **Delete Route** to delete one—In the **Delete Route** dialog box, select a route, enter **DELETE** manually, and then click **OK**.

## Testing Connectivity

Test the connectivity between the Logstash cluster and the data sources.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Select the **Configuration Center** tab, and click **Test Connectivity**.
5. In the **Test Connectivity** dialog box, enter the IP address and port number of a data source, and click **Test**.

**Figure 4-6** Testing connectivity

Test Connectivity

IP address Port

Add More You can still add 9 more.

Test

Cancel Test

If **Available** is displayed, the network between the Logstash cluster and the data source is connected.

 **NOTE**

You can test a maximum of 10 IP addresses at a time. Click **Add** to add multiple IP addresses and click **Test** at the bottom to test connectivity to all these IP addresses at the same time.

## 4.5.2 Configuring Public Network Access for a Logstash Cluster

CSS Logstash is a data ingestion service that ingests data from designated data sources on the public network and writes it to designated storage systems. The data is subsequently used for purposes such as big data analytics and log management. An NAT gateway and return routes need to be configured to ensure proper data ingestion, storage, and processing across networks.

- NAT gateway: Configure an NAT gateway for the Logstash cluster to enable the cluster to initiate connection requests and connect to Internet data sources.
- Return routes: They direct Logstash nodes to connect to the IP addresses of data sources on the public network and back, ensuring proper data ingestion and transmission.

### Constraints

- An NAT gateway and EIP are required to enable public network access for a Logstash cluster. They must be purchased additionally and be in the same region as the Logstash cluster.
- The SNAT rule added to the NAT gateway applies only to resources in the specified subnet. Make sure you use the subnet of the Logstash cluster.

## Configuring Public Network Access for a Logstash Cluster

**Step 1** Obtain the VPC and subnet of the Logstash cluster for configuring an NAT gateway later.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Overview** tab. In the **Configuration** area, record the cluster's VPC and subnet information.

**Step 2** Create a public NAT gateway for the Logstash cluster. For details, see [Creating a Public NAT Gateway](#).

[Table 4-14](#) describes the key parameters. Set other parameters based on service requirements.



**Table 4-14** Configuring a public NAT gateway

| Parameter      | Description                                                                                                                                          |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Region         | Use the region of the Logstash cluster.                                                                                                              |
| Specifications | Select a value based on the needed bandwidth. The maximum number of connections supported by SNAT varies depending on the specifications you choose. |
| VPC            | Select the VPC recorded in <a href="#">Step 1</a> .                                                                                                  |
| Subnet         | Select the subnet recorded in <a href="#">Step 1</a> .                                                                                               |

**Step 3** Add SNAT rules on the NAT gateway to allow the Logstash cluster to access the public network through a shared EIP on the gateway. For details, see [Adding an SNAT Rule](#).

[Table 4-15](#) describes the key parameters. Set other parameters based on service requirements.

**Table 4-15** Adding an SNAT rule

| Parameter              | Description                                                                                                                                                                                                                                                                                                                                                           |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Scenario               | Select <b>VPC</b> .                                                                                                                                                                                                                                                                                                                                                   |
| CIDR Block             | Select <b>Existing</b> , and then select the VPC recorded in <a href="#">Step 1</a> .                                                                                                                                                                                                                                                                                 |
| Public IP Address Type | Select <b>EIP</b> , and then select an IP address that meets your requirements.<br>If no EIP is available, click <b>View EIP</b> to go to the EIP console and <a href="#">buy an EIP</a> .<br>You are advised to configure the EIP as follows: <ul style="list-style-type: none"><li>• <b>Billing Mode: Pay-per-use</b></li><li>• <b>Billed By: Traffic</b></li></ul> |

----End

## Configuring Routes for a Logstash Cluster

Configure a return route for the Logstash cluster to make sure the Logstash nodes can connect to the IP addresses of data sources on the public network. The return route must cover all these IP addresses.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Click the **Overview** tab. In the **Configuration** area, click **Add Route** next to **Cluster Route**.

**Figure 4-7** Adding a route

5. In the displayed dialog box, configure the route information.

**Table 4-16** Adding a route

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP Address  | Enter the IP address (or a CIDR block) of a data source, for example, <b>10.10.1.0</b> .<br>If the IP addresses of data sources are scattered, configure multiple routes for them.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Subnet Mask | Enter the subnet mask used with the IP address.<br>The subnet mask must align with the network part of the IP address. That is: <ul style="list-style-type: none"><li>• The network bits (the 1s) of the subnet mask must encompass those of the IP address.</li><li>• The host bits (the 0s) of the subnet mask must accommodate the IP address's host bits.</li></ul> The subnet mask should not cover too wide a range, avoiding unnecessary IP address exposure or waste of the address space. 255.255.255.0 (that is, /24) is recommended, which applies to most internal networking scenarios. |

6. Click **OK** to add the route.

After the route is added, click **View Route** to check the route information. You can click **Add Route** to add another route; or click **Delete Route** to delete one—In the **Delete Route** dialog box, select a route, enter **DELETE** manually, and then click **OK**.

## Testing Connectivity

Test the connectivity between the Logstash cluster and the data sources.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Select the **Configuration Center** tab, and click **Test Connectivity**.
5. In the **Test Connectivity** dialog box, enter the IP address and port number of a data source, and click **Test**.

**Figure 4-8** Testing connectivity

If **Available** is displayed, the network between the Logstash cluster and the data source is connected.

#### NOTE

You can test a maximum of 10 IP addresses at a time. Click **Add** to add multiple IP addresses and click **Test** at the bottom to test connectivity to all these IP addresses at the same time.

## 4.6 Configuring a Logstash Migration Task

### 4.6.1 Creating a Logstash Configuration File

Logstash clusters use configuration files to define plans to migrate data from different data sources (input) to different destinations (output).

#### Constraints

- The size of each configuration file cannot exceed 100 KB.
- A Logstash cluster can have a maximum of 50 configuration files. There can be a maximum of five configuration files that are in the **Verifying** state at the same time.

#### Prerequisites

You have obtained information about servers or clusters at the data source and destination, including IP addresses, usernames, and passwords.

#### Creating a Configuration File

1. Go to the **Configuration Center** page.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > Logstash**.
  - c. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  - d. Click the **Configuration Center** tab.
2. On the Configuration Center page, click **Create** in the upper-right corner. On the **Create Configuration File** page that is displayed, edit the configuration file.

**Table 4-17** Creating a configuration file

| Parameter                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                       | User-defined configuration file name.<br>It can contain only letters, digits, hyphens (-), and underscores (_), and must start with a letter. The minimum length is 4 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Configuration File Content | You can configure the configuration file content by modifying a template, and develop it from scratch.<br>To use a configuration file template, expand <b>System Templates</b> or <b>Custom Templates</b> , select a template based on its description, and click <b>Apply</b> in the <b>Operation</b> column to copy the content of the template to the <b>Configuration File Content</b> area, where you can further modify relevant items. <ul style="list-style-type: none"><li>System templates are preset configuration file templates of CSS. For details, see <a href="#">Logstash Configuration File Templates</a>.</li><li>Custom templates are templates added by users.</li></ul> <b>NOTE</b><br>The size of each configuration file cannot exceed 100 KB. |
| Hidden Content             | For items that you enter in this box, the corresponding strings will be replaced with *** in the configurations.<br>Enter sensitive strings that you want to hide, and press <b>Enter</b> .<br>You can enter a maximum of 20 strings, each with a maximum length of 512 bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

3. Click **Next** to configure Logstash pipeline parameters.

**Table 4-18** Pipeline parameters

| Parameter           | Description                                                                                                                                                                                                    |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| pipeline.workers    | Number of worker threads that will execute the Filters and Outputs stages of the pipeline in parallel.<br>Value range: 1–20<br>Default value: Number of CPU cores                                              |
| pipeline.batch.size | Maximum number of events that a worker thread collects from inputs before attempting to execute its filters and outputs. A larger value is more effective but increases memory overhead.<br>Default value: 125 |

| Parameter               | Description                                                                                                                                                                                                                                               |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| pipeline.batch.delay    | Maximum wait time for each new event before scheduling small batches to the pipeline worker thread and creating a pipeline event batch.<br>Default value: 50<br>Unit: milliseconds.                                                                       |
| queue.type              | An internal queue for event buffering. <ul style="list-style-type: none"><li>• <b>memory</b> indicates a traditional memory-based queue.</li><li>• <b>persisted</b> indicates a disk-based ACKed persistent queue.</li></ul> Default value: <b>memory</b> |
| queue.checkpoint.writes | Maximum number of written events before a forced checkpoint. This parameter needs to be set only when <b>queue.type</b> is set to <b>persisted</b> .<br>Default value: 1024                                                                               |
| queue.max_bytes         | Total capacity of each persistent queue in MB. Ensure that the disk capacity is greater than this value. This parameter needs to be set only when <b>queue.type</b> is set to <b>persisted</b> .<br>Default value: 1024<br>Unit: MB                       |

4. After the configuration is complete, click **Create**.

On the **Configuration Center** page, you can check the created configuration file. If its status changes to **Available**, it has been successfully created.

## Managing Configuration Files

You can edit existing configuration files, set a configuration file as a custom template, and delete configuration files.

**Table 4-19** Managing configuration files

| Function                       | Constraints                                                                | Operation Guide                                                                                                                                                                                                                                                                      |
|--------------------------------|----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Modifying a configuration file | A configuration file that has an ongoing pipeline task cannot be modified. | In the configuration file list, find the row that contains the configuration file you want to edit, and click <b>Edit</b> in the <b>Operation</b> column. On the displayed page, you can modify <b>Configuration File Content</b> , <b>Hidden Content</b> , and pipeline parameters. |

| Function                                   | Constraints                                                               | Operation Guide                                                                                                                                                                                                                                                                                                                              |
|--------------------------------------------|---------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Setting a configuration file as a template | N/A                                                                       | In the configuration file list, click <b>Add to Custom Template</b> in the <b>Operation</b> column. In the displayed <b>Add to Custom Template</b> dialog box, set the template name, description, and configuration file content, and click <b>OK</b> . It can then be used as a template for configuring configuration files subsequently. |
| Exporting configuration files              | N/A                                                                       | Click  in the upper right corner above the configuration file list to export all configuration files for local backup.                                                                                                                                     |
| Deleting a configuration file              | A configuration file that has an ongoing pipeline task cannot be deleted. | In the configuration list, find the row that contains the configuration file you want to delete, and click <b>Delete</b> in the <b>Operation</b> column. In the displayed dialog box, click <b>OK</b> .                                                                                                                                      |

## 4.6.2 Starting a Logstash Migration Task

### Scenario

To start a migration task with a Logstash cluster, perform the following steps:

1. Test the network connectivity between the data source and destination and the Logstash cluster. For details, see [Testing Connectivity](#).
2. Start a Logstash migration task by starting a configuration file.
  - If there is no **Running** task in the pipeline list, go to [Starting a Configuration File](#).
  - If there is a **Running** task, go to [Hot Starting a Configuration File](#).
3. Monitor the migration task through metrics, operation records, and run logs. For details, see [Monitoring Migration Tasks](#).

### Constraints

- There can be up to 50 **Running** configuration files in the pipeline list of a Logstash cluster.
- Configuration files that use the **logstash stdin** plugin cannot use the **Hot Start** option.
- If hot-starting a configuration file fails and the Logstash process exits unexpectedly, the cluster will try to restart the original Logstash process. Exercise caution when using the **Hot Start** option.

## Testing Connectivity

Before using a Logstash cluster to migrate data, test the network connectivity between the data source and destination and the Logstash cluster.

1. Go to the **Configuration Center** page.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > Logstash**.
  - c. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  - d. Click the **Configuration Center** tab.
2. On the **Configuration Center** page, click **Test Connectivity**.
3. In the **Test Connectivity** dialog box, enter the IP address and port number of the data source and destination, and click **Test**.

You can test a maximum of 10 IP addresses at a time. You can click **Add** to add more IP addresses and click **Test** at the bottom to test connectivity to multiple IP addresses at a time.

**Figure 4-9** Test Connectivity



If **Available** is displayed, the network is connected. If the network is disconnected, configure routes for the Logstash cluster to connect the clusters. For details, see [Configuring Routes for a Logstash Cluster](#).

## Starting a Configuration File

If there is no **Running** task in the pipeline list, start the Logstash migration task by starting a configuration file. You can select multiple configuration files to start them all at once. A maximum of 50 configuration files can be selected at a time.

1. In the configuration file list, select configuration files you want to start, and click **Start** in the upper left corner.
2. In the **Start Logstash** dialog box, select **Keepalive** if necessary.

When **Keepalive** is enabled, a daemon process will be configured on each node. If the Logstash service becomes faulty, the daemon process will try to rectify the fault and restart the service. You are advised to enable Keepalive for services running long-term. Do not enable it for services running only short-term, or your migration tasks may fail due to a lack of source data.
3. Click **OK** to start the configuration file and hence the Logstash migration task. You can view the started configuration file in the pipeline list.

## Hot Starting a Configuration File

If there are **Running** tasks in the pipeline list, start the Logstash migration task by hot-starting a configuration file. Only one configuration file can be hot-started at a time. That is, only one pipeline can be added at a time.

1. In the configuration file list, select a configuration file, and click **Hot Start** in the upper left corner.
2. In the **Hot Start Logstash Configuration** dialog box, the Keepalived setting is kept consistent with other tasks in the pipeline list by default and cannot be changed.
3. Click **OK** to hot-start the configuration file and hence the Logstash migration task.

You can view the hot-started configuration file in the pipeline list.

## Monitoring Migration Tasks

After a configuration file is started, a task is generated in the pipeline list, where you can monitor the task.

- In the pipeline list, click **Metric Monitoring** in the **Operation** column to go to the Cloud Eye console and view monitoring information.

When the **Events** data of a pipeline changes dynamically, the monitoring data changes accordingly. When a pipeline task is being started or stopped, or the **Events** data is stable, the monitoring data remains unchanged.

For details about the monitoring metrics supported, see [Logstash Pipeline Monitoring Metrics](#). For how to configure monitoring alarms, see [Using Cloud Eye to Monitor Logstash Clusters](#).

- You can also click **View Operation Record** or **View Running Log** above the pipeline list to view configuration file operation records or the run logs of a migration task.

## 4.6.3 Stopping Logstash Migration Tasks

### Scenario

You can stop Logstash migration tasks using either of the following methods:

- **Stop All:** Stop all tasks in the pipeline list. After the tasks are stopped successfully, their **Status** changes to **Stopped**.
- **Hot Stop:** In the pipeline list, select one task, and click **Hot Stop**. Only one pipeline task can be hot-stopped at a time. A successfully hot-stopped pipeline task is removed from the pipeline list right away.

#### NOTE

Stopping a pipeline will interrupt the corresponding Logstash data migration task. Exercise caution.

### Prerequisites

There are **Running** pipeline tasks in the configuration center of the Logstash cluster.



## Stopping All Migration Tasks

1. Go to the **Configuration Center** page.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > Logstash**.
  - c. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  - d. Click the **Configuration Center** tab.
2. To stop all data migration tasks, click **Stop All** above the pipeline list.
3. In the displayed dialog box, click **OK**.

If all tasks are stopped successfully, **Status** changes to **Stopped** for all tasks in the pipeline list, and data migration is stopped for all of them.

## Hot Stopping a Migration Task

1. Go to the **Configuration Center** page.
  - a. Log in to the [CSS management console](#).
  - b. In the navigation pane on the left, choose **Clusters > Logstash**.
  - c. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
  - d. Click the **Configuration Center** tab.
2. To stop the data migration task of a single configuration file in the pipeline list, select the target pipeline, and click **Hot Stop** above.
3. In the displayed dialog box, click **OK**.

If the hot stop is successful, the task is removed from the pipeline list and data migration is stopped.

## 4.6.4 Logstash Configuration File Templates

### Introduction to the Templates

**Table 4-20** lists commonly used Logstash configuration file templates.

**Table 4-20** Logstash configuration file templates

| Name          | Description                                                     | Details                                |
|---------------|-----------------------------------------------------------------|----------------------------------------|
| redis         | Imports data from a Redis database to an Elasticsearch cluster. | <a href="#">Redis Template</a>         |
| elasticsearch | Migrates data between Elasticsearch clusters.                   | <a href="#">Elasticsearch Template</a> |
| jdbc          | Imports data from JDBC to an Elasticsearch cluster.             | <a href="#">JDBC Template</a>          |
| kafka         | Imports data from Kafka to an Elasticsearch cluster.            | <a href="#">Kafka Template</a>         |

| Name  | Description                                          | Details                        |
|-------|------------------------------------------------------|--------------------------------|
| beats | Imports data from Beats to an Elasticsearch cluster. | <a href="#">DIS Template</a>   |
| dis   | Imports data from DIS to an Elasticsearch cluster.   | <a href="#">Beats Template</a> |

## Redis Template

Imports data from a Redis database to an Elasticsearch cluster.

```
input {
 redis {
 data_type => "pattern_channel" # data_type, one of ["list", "channel", "pattern_channel"]
 key => "lgs-*" # The name of a Redis list or channel.
 host => "xxx.xxx.xxx.xxx"
 port => 6379
 }
}

filter {
 # Delete some fields added by Logstash.
 mutate {
 remove_field => ["@timestamp", "@version"]
 }
}

output {
 elasticsearch {
 hosts => ["http://xxx.xxx.xxx.xxx:9200", "http://xxx.xxx.xxx.xxx:9200", "http://xxx.xxx.xxx.xxx:9200"] #
 }
 host
 # user => "xxx" # user name, just for security cluster
 # password => "xxx" # password, just for security cluster
 index => "xxxxxx" # destination_index
}
```

**Table 4-21** Configuration items

| Configuration Item | Mandatory | Description                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| data_type          | Yes       | Data source type. The options are <b>list</b> , <b>channel</b> , and <b>pattern_channel</b> . <ul style="list-style-type: none"><li>• If the parameter is set to <b>list</b>, the BLPOP key is used.</li><li>• If the parameter is set to <b>channel</b>, the SUBSCRIBE key is used.</li><li>• If the parameter is set to <b>pattern_channel</b>, the PSUBSCRIBE key is used.</li></ul> |
| key                | Yes       | Redis list or channel name                                                                                                                                                                                                                                                                                                                                                              |
| host               | Yes       | IP address of the Redis server                                                                                                                                                                                                                                                                                                                                                          |

| Configuration Item | Mandatory | Description                                                                                                                                                   |
|--------------------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| port               | No        | Number of the port to be connected.<br>Default value: <b>6379</b> .                                                                                           |
| hosts              | Yes       | Address for accessing the Elasticsearch cluster.                                                                                                              |
| user               | No        | Username for accessing the Elasticsearch cluster. Generally, the value is <b>admin</b> .<br>This parameter is required only for a security cluster.           |
| password           | No        | Password for accessing the Elasticsearch cluster. The password is set when the cluster is created.<br>This parameter is required only for a security cluster. |
| index              | Yes       | Index to which the data is to be migrated.<br>Only one index can be configured.                                                                               |

For more information, see the Logstash document [Redis input plugin](#).

## Elasticsearch Template

Migrates data between Elasticsearch clusters.

```
input {
 elasticsearch {
 # Source ES cluster IP address. When SSL is enabled, please use 7.10.0 and you donot need to add a
 # protocol, otherwise an error is reported.
 hosts => ["http://xxx.xxx.xxx.xxx:9200", "http://xxx.xxx.xxx.xxx:9200"]
 # user => "xxx" # user name, just for security cluster
 # password => "xxx" # password, just for security cluster
 index => "xxx,xxx,xxx" # List of indexes to be migrated, separate with commas (,)
 docinfo => true
 # source ES certificate, the cluster on the cloud this value remains the same. Enter the corresponding
 # path when using a custom certificate. Self-built Logstash cluster, You can download the certs file from the
 # ES cluster details page, enter the corresponding path here.
 # ca_file => "/rds/datastore/logstash/v7.10.0/package/logstash-7.10.0/extend/certs" # for 7.10.0
 # ssl => true # Set to true when SSL is enabled.
 }
}

filter {
 # Delete some fields added by Logstash.
 mutate {
 remove_field => ["@timestamp", "@version"]
 }
}

output {
 elasticsearch {
 # Destination ES cluster IP address. When SSL is enabled, please use 7.10.0 and you donot need to
 # add a protocol.
 hosts => ["http://xxx.xxx.xxx.xxx:9200", "http://xxx.xxx.xxx.xxx:9200"]
 # user => "xxx" # user name, just for security cluster
 # password => "xxx" # password, just for security cluster
 }
}
```

```
index => "%{[@metadata][_index]}" # destination _index, this config ensure that the index is
consistent with that on the source. You can also specify the index name.
document_type => "%{[@metadata][_type]}" # destination _type, this config ensure that the
document_type is consistent with that on the source.
document_id => "%{[@metadata][_id]}" # destination _id, if you do not need to retain the original
_id, delete this line to improve performance.
source ES certificate, the cluster on the cloud this value remains the same. Enter the corresponding
path when using a custom certificate. Self-built Logstash cluster, You can download the certs file from the
ES cluster details page, enter the corresponding path here.
cacert => "/rds/datastore/logstash/v7.10.0/package/logstash-7.10.0/extend/certs" # for 7.10.0
ssl => true # Set to true when SSL is enabled.
ssl_certificate_verification => false # Set to false to ignore the validation server certificate when SSL
is enabled.
}
```

**Table 4-22** Configuration items

| Configuration Item | Mandatory | Description                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| hosts              | Yes       | Address for accessing the destination Elasticsearch cluster to which data is imported.                                                                                                                                                                                                                                                                                                                                        |
| user               | No        | Username for accessing the Elasticsearch cluster. Generally, the value is <b>admin</b> .<br>This parameter is required only for a security cluster.                                                                                                                                                                                                                                                                           |
| password           | No        | Password for accessing the Elasticsearch cluster. The password is set when the cluster is created.<br>This parameter is required only for a security cluster.                                                                                                                                                                                                                                                                 |
| index              | Yes       | Index from which data is to be migrated.                                                                                                                                                                                                                                                                                                                                                                                      |
| docinfo            | No        | Document information.<br>Value: <b>true</b> or <b>false</b> .<br>If this parameter is specified, include Elasticsearch document information in the event, such as the index, type, and ID.                                                                                                                                                                                                                                    |
| ca_file            | No        | Default value: <b>/rds/datastore/logstash/v7.10.0/package/logstash-7.10.0/extend/certs</b> .<br>For a cloud-based Logstash cluster, retain the default value, or when using a custom certificate, enter the path of the custom certificate.<br>For an in-house built Logstash cluster, you can download the certificate on the details page of the SSL-enabled Elasticsearch cluster. Enter the path of the certificate here. |
| ssl                | No        | Set this parameter to <b>true</b> if SSL is enabled for the source Elasticsearch cluster.                                                                                                                                                                                                                                                                                                                                     |

| Configuration Item           | Mandatory | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| hosts                        | Yes       | Address for accessing the source Elasticsearch cluster where data comes from.                                                                                                                                                                                                                                                                                                                                                                                                   |
| user                         | No        | Username for accessing the Elasticsearch cluster. Generally, the value is <b>admin</b> .<br>This parameter is required only for a security cluster.                                                                                                                                                                                                                                                                                                                             |
| password                     | No        | Password for accessing the Elasticsearch cluster. The password is set when the cluster is created.<br>This parameter is required only for a security cluster.                                                                                                                                                                                                                                                                                                                   |
| index                        | Yes       | Index to which the data is to be migrated. Only one index can be configured.                                                                                                                                                                                                                                                                                                                                                                                                    |
| document_type                | No        | <ul style="list-style-type: none"><li>This parameter is valid only when <b>docinfo</b> is set to <b>true</b>.</li><li>If <b>docinfo</b> is set to <b>false</b>, delete this parameter from the configuration file.</li></ul>                                                                                                                                                                                                                                                    |
| document_id                  | No        | <ul style="list-style-type: none"><li>This parameter is valid only when <b>docinfo</b> is set to <b>true</b>.</li><li>If <b>docinfo</b> is set to <b>false</b>, delete this parameter from the configuration file.</li></ul>                                                                                                                                                                                                                                                    |
| cacert                       | No        | Default value: <b>/rds/datastore/logstash/v7.10.0/package/logstash-7.10.0/extend/certs</b> . <ul style="list-style-type: none"><li>For a cloud-based Logstash cluster, retain the default value, or when using a custom certificate, enter the path of the custom certificate.</li><li>For an in-house built Logstash cluster, you can download the certificate on the details page of the SSL-enabled Elasticsearch cluster. Enter the path of the certificate here.</li></ul> |
| ssl                          | No        | Set this parameter to <b>true</b> if SSL is enabled for the destination Elasticsearch cluster.                                                                                                                                                                                                                                                                                                                                                                                  |
| ssl_certificate_verification | No        | Set this parameter to <b>false</b> to enable SSL and ignore the server certificate verification.                                                                                                                                                                                                                                                                                                                                                                                |

For more information, see the Logstash document [Elasticsearch input plugin](#).

## JDBC Template

Imports data from JDBC to an Elasticsearch cluster.

```
input {
 jdbc {
 # for 7.10.0, jdbc_driver_library => "/rds/datastore/logstash/v7.10.0/package/logstash-7.10.0/extend/
jars/mariadb-java-client-2.7.0.jar"
 jdbc_driver_library => "xxxxxxxxxx"
 jdbc_driver_class => "org.mariadb.jdbc.Driver"
 jdbc_connection_string => "jdbc:mariadb://xxx.xxx.xxx.xxx:xxx/data_base_name"
 jdbc_user => "xxxx"
 jdbc_password => "xxxx"
 statement => "SELECT * from table_name" # This SQL statement determines the data to be input
 }
}

filter {
 # Delete some fields added by Logstash.
 mutate {
 remove_field => ["@timestamp", "@version"]
 }
}

output {
 elasticsearch {
 hosts => ["http://xxx.xxx.xxx.xxx:9200", "http://xxx.xxx.xxx.xxx:9200", "http://xxx.xxx.xxx.xxx:9200"] #
host
 # user => "xxxx" # user name, just for security cluster
 # password => "xxxx" # password, just for security cluster
 index => "xxxxxx" # destination_index
 }
}
```

**Table 4-23** Configuration items

| Configuration Item     | Mandator y | Description                                                                                                                                                                                                                                                                                                                                             |
|------------------------|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| jdbc_driver_library    | Yes        | Path of the JDBC driver library. <ul style="list-style-type: none"><li>If the Logstash cluster version is 7.10.0, set this parameter to <b>/rds/datastore/logstash/v7.10.0/package/logstash-7.10.0/extend/jars/mariadb-java-client-2.7.0.jar</b>.</li></ul> Currently, only these existing drivers are supported. User-defined upload is not supported. |
| jdbc_driver_class      | Yes        | JDBC driver class to be loaded, for example, <b>org.mariadb.jdbc.Driver</b> .                                                                                                                                                                                                                                                                           |
| jdbc_connection_string | Yes        | JDBC connection string                                                                                                                                                                                                                                                                                                                                  |
| jdbc_user              | Yes        | JDBC username                                                                                                                                                                                                                                                                                                                                           |
| jdbc_password          | Yes        | JDBC password                                                                                                                                                                                                                                                                                                                                           |
| statement              | Yes        | SQL statement of the input data.                                                                                                                                                                                                                                                                                                                        |
| hosts                  | Yes        | Address for accessing the source Elasticsearch cluster where data comes from.                                                                                                                                                                                                                                                                           |

| Configuration Item | Mandatory | Description                                                                                                                                                   |
|--------------------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| user               | No        | Username for accessing the Elasticsearch cluster. Generally, the value is <b>admin</b> .<br>This parameter is required only for a security cluster.           |
| password           | No        | Password for accessing the Elasticsearch cluster. The password is set when the cluster is created.<br>This parameter is required only for a security cluster. |
| index              | Yes       | Index to which the data is to be migrated.<br>Only one index can be configured.                                                                               |

For more information, see the Logstash document [Jdbc input plugin](#).

## Kafka Template

Imports data from Kafka to an Elasticsearch cluster.

```
input {
 kafka {
 bootstrap_servers=>"xxx.xxx.xxx.xxx:xxx"
 topics=>["xxxxxxx"]
 group_id=>"kafka_es_test"
 auto_offset_reset=>"earliest"
 }
}

filter {
 # Delete some fields added by Logstash.
 mutate {
 remove_field => ["@timestamp", "@version"]
 }
}

output {
 elasticsearch {
 hosts => ["http://xxx.xxx.xxx.xxx:9200", "http://xxx.xxx.xxx.xxx:9200", "http://xxx.xxx.xxx.xxx:9200"] #
 host
 # user => "xxx" # user name, just for security cluster
 # password => "xxx" # password, just for security cluster
 index => "xxxxxx" # destination_index
 }
}
```

**Table 4-24** Configuration items

| Configuration Item | Mandatory | Description                                      |
|--------------------|-----------|--------------------------------------------------|
| bootstrap_servers  | Yes       | IP address and port number of the Kafka instance |
| topics             | Yes       | List of topics to be subscribed to               |

| Configuration Item | Mandatory | Description                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| group_id           | Yes       | Identifier of the group to which the consumer belongs.                                                                                                                                                                                                                                                                                                                                                                                  |
| auto_offset_reset  | Yes       | Initial offset in Kafka. <ul style="list-style-type: none"><li>• <b>earliest</b>: Automatically reset the offset to the earliest offset.</li><li>• <b>latest</b>: Automatically reset the offset to the latest offset.</li><li>• <b>none</b>: Report an exception to the consumer if the previous offset for the consumer group is not found.</li><li>• <b>anything else</b>: The system throws an exception to the consumer.</li></ul> |
| hosts              | Yes       | Address for accessing the source Elasticsearch cluster where data comes from.                                                                                                                                                                                                                                                                                                                                                           |
| user               | No        | Username for accessing the Elasticsearch cluster. Generally, the value is <b>admin</b> .<br>This parameter is required only for a security cluster.                                                                                                                                                                                                                                                                                     |
| password           | No        | Password for accessing the Elasticsearch cluster. The password is set when the cluster is created.<br>This parameter is required only for a security cluster.                                                                                                                                                                                                                                                                           |
| index              | Yes       | Index to which the data is to be migrated.<br>Only one index can be configured.                                                                                                                                                                                                                                                                                                                                                         |

For more information, see the Logstash document [Kafka input plugin](#).

## DIS Template

Imports data from DIS to an Elasticsearch cluster.

```
input {
 dis {
 # For details, see https://support.huaweicloud.com/usermanual-dis/dis_01_0099.html.
 streams => ["YOUR_DIS_STREAM_NAME"]
 endpoint => "https://dis.xxxxx.myhuaweicloud.com" # xxxxx should be replaced by your region name
 ak => "YOUR_ACCESS_KEY_ID"
 sk => "YOUR_SECRET_KEY_ID"
 region => "YOUR_Region"
 project_id => "YOUR_PROJECT_ID"
 group_id => "YOUR_APP_ID"
 client_id => "YOUR_CLIENT_ID"
 auto_offset_reset => "earliest"
 }
}
```



```
filter {
 # Delete some fields added by Logstash.
 mutate {
 remove_field => ["@timestamp", "@version"]
 }
}

output {
 elasticsearch {
 hosts => ["http://xxx.xxx.xxx.xxx:9200", "http://xxx.xxx.xxx.xxx:9200", "http://xxx.xxx.xxx.xxx:9200"] #
 host
 # user => "xxx" # user name, just for security cluster
 # password => "xxx" # password, just for security cluster
 index => "xxxxxx" # destination _index
}
}
```

**Table 4-25** Configuration items

| Configuration Item | Mandatory | Description                                                                                                                                                                                                                                                                              |
|--------------------|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| streams            | Yes       | Name of the DIS stream.<br>The entered DIS stream name must be the same as the stream name specified when you are creating a DIS stream on the DIS console.                                                                                                                              |
| endpoint           | Yes       | Data API address of the region where DIS resides.                                                                                                                                                                                                                                        |
| ak                 | Yes       | User's access key (AK).<br>For details, see <a href="#">Checking Authentication Information</a> .                                                                                                                                                                                        |
| sk                 | Yes       | User's secret key (SK).<br>For details, see <a href="#">Checking Authentication Information</a> .                                                                                                                                                                                        |
| region             | Yes       | Region where DIS is supported.                                                                                                                                                                                                                                                           |
| project_id         | Yes       | Project ID of the region.<br>For details, see <a href="#">Checking Authentication Information</a> .                                                                                                                                                                                      |
| group_id           | Yes       | DIS App name, used to identify a consumer group. The value can be any character string.                                                                                                                                                                                                  |
| client_id          | No        | Client ID, which identifies a consumer in a consumer group.<br>If multiple pipelines or Logstash instances are started for consumption, set this parameter to different values. For example, the value of instance 1 is <b>client1</b> , and the value of instance 2 is <b>client2</b> . |

| Configuration Item | Mandatory | Description                                                                                                                                                                                                                                                      |
|--------------------|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| auto_offset_reset  | No        | Position where data starts to be consumed from the stream. The options are as follows: <ul style="list-style-type: none"><li>• <b>earliest</b>: Data is consumed from the earliest one.</li><li>• <b>latest</b>: Data is consumed from the latest one.</li></ul> |
| hosts              | Yes       | Address for accessing the source Elasticsearch cluster where data comes from.                                                                                                                                                                                    |
| user               | No        | Username for accessing the Elasticsearch cluster. Generally, the value is <b>admin</b> . This parameter is required only for a security cluster.                                                                                                                 |
| password           | No        | Password for accessing the Elasticsearch cluster. The password is set when the cluster is created. This parameter is required only for a security cluster.                                                                                                       |
| index              | Yes       | Index to which the data is to be migrated. Only one index can be configured.                                                                                                                                                                                     |

## Beats Template

Imports data from Beats to an Elasticsearch cluster.

```
input {
 beats {
 port => 5044 # port
 }
}

filter {
 # Delete some fields added by Logstash.
 mutate {
 remove_field => ["@timestamp", "@version"]
 }
}

output {
 elasticsearch {
 hosts => ["http://xxx.xxx.xxx.xxx:9200", "http://xxx.xxx.xxx.xxx:9200", "http://xxx.xxx.xxx.xxx:9200"] #
 }
 host {
 # user => "xxx" # user name, just for security cluster
 # password => "xxx" # password, just for security cluster
 index => "xxxxxx" # destination_index
 }
}
```

**Table 4-26** Configuration items

| Configuration Item | Mandatory | Description                                                                                                                                                   |
|--------------------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| port               | Yes       | The port number <b>5044</b> is used for connecting and indexing Elasticsearch through Beats.                                                                  |
| hosts              | Yes       | Address for accessing the source Elasticsearch cluster where data comes from.                                                                                 |
| user               | No        | Username for accessing the Elasticsearch cluster. Generally, the value is <b>admin</b> .<br>This parameter is required only for a security cluster.           |
| password           | No        | Password for accessing the Elasticsearch cluster. The password is set when the cluster is created.<br>This parameter is required only for a security cluster. |
| index              | Yes       | Index to which the data is to be migrated.<br>Only one index can be configured.                                                                               |

For more information, see the Logstash document [Beats input plugin](#).

## 4.7 Scaling a Logstash Cluster

### 4.7.1 Scaling Out a Logstash Cluster

If the workloads on the data plane of a Logstash cluster increase, you can scale out the cluster by adding more nodes.

You may scale out a Logstash cluster for the following purposes:


- To ensure continuous data availability, more nodes may be added to a Logstash cluster to enhance fault tolerance.
- In the face of fast data growth, more nodes may be added to a Logstash cluster to increase data processing capacity.
- More nodes may be added to a Logstash cluster to execute more complex data processing tasks, such as data enrichment and ETL.

Scale-out process: Add more nodes of the appropriate specifications to the Logstash cluster and modify cluster configuration.

#### Prerequisites

- The cluster status is **Available**, and there are no ongoing tasks.
- Your CSS resource quotas are sufficient for the capacity expansion you are about to perform. You can check available resources on the **Modify Configuration** page.

Figure 4-10 Checking available resources



Impact on Billing

For a pay-per-use cluster, you can see its new price when confirming the scale-out or scale-up on the console. After the scale-out or scale-up is complete, the new price will apply. For pricing details, see [Cloud Search Service Price Calculator](#).

For a yearly/monthly cluster, a new order will be placed after you confirm the scale-out or scale-up. Then you can check the new price.

Change Impact

Scaling out a Logstash cluster does not affect services. However, once started, a scale-out task cannot be stopped until it succeeds or fails.

Scale-out Duration

The scale-out operation takes 10 to 30 minutes, depending on the cluster's scheduling capabilities.


Adding New Nodes

- 1. Log in to the [CSS management console](#).
- 2. In the navigation pane on the left, choose **Clusters > Logstash**.
- 3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
- 4. Select the **Scale Cluster** tab and click **Scale out** to set parameters.

Table 4-27 Adding new nodes

| Parameter | Description                                              |
|-----------|----------------------------------------------------------|
| Action    | Select <b>Scale out</b> .                                |
| Resources | Quantities of resources added.                           |
| Nodes     | Increase the number of nodes in the <b>Nodes</b> column. |

Figure 4-11 Adding more nodes



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

## Related Documents

You can also expand the capacity of a Logstash cluster by upgrading node specifications. For details, see [Changing the Node Specifications of a Logstash Cluster](#).

### 4.7.2 Scaling In a Logstash Cluster

If a cluster can process current traffic without fully utilizing its resources, you can scale in the cluster to cut costs.

Scaling in a Logstash cluster means to randomly remove some of its nodes, so that it is less costly to keep it running.

Scale-in process: Random remove cluster nodes and modify the cluster configuration.

## Prerequisites

The cluster status is **Available**, and there are no ongoing tasks.

## Impact on Billing

For a pay-per-use cluster, you can see its new price when confirming the scale-in on the console. After the scale-out or scale-up is complete, the new price will apply. For pricing details, see [Cloud Search Service Price Calculator](#).

For a yearly/monthly cluster, a refund request will be triggered after you confirm the scale-in. Then you can check the new price.

## Constraints

- To reduce the impact on services, you are advised to perform scale-in during off-peak hours.
- In a cross-AZ cluster, the difference between the numbers of the same-type nodes in different AZs cannot exceed 1.

## Change Impact

Before the change, learn about possible impacts and operation suggestions, and develop a plan to minimize these impacts.

- Impact on cluster throughput  
Taking Logstash nodes offline usually does not interrupt services. However, jobs that are running on these will be stopped, causing the cluster throughput to decrease.
- Characteristics of this process

Once started, a scaling task cannot be stopped until it succeeds or fails.

## Scale-in Duration

The following formula can be used to estimate how long a scale-in operation will take:

**Scale-in duration (min) = 5 (min) x Number of nodes to be removed**

where, 5 minutes is the duration per node. It is an empirical value.

## Reducing the Nodes of a Logstash Cluster


1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. Select the **Scale Cluster** tab and click **Scale in** to set parameters.

**Table 4-28** Removing nodes randomly

| Parameter | Description                      |
|-----------|----------------------------------|
| Action    | Select <b>Scale in</b> .         |
| Resources | Quantities of resources reduced. |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agency    | <p>When a node is deleted, NICs are released. This means you need to have VPC permissions. Select an IAM agency to grant the current account the permission to access and use VPC.</p> <ul style="list-style-type: none"><li>This parameter is available only when the new IAM plane is connected.</li><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-upgrade-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the permissions associated with the <b>VPC Administrator</b> role or the <b>VPC FullAccess</b> system policy deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<div>"vpc:subnets:get",<br/>"vpc:ports:*"</div></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<div>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</div></li><li>To use an IAM agency, the following minimum permissions are required:<div>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</div></li></ul> |
| Nodes     | Reduce the number of nodes in the <b>Nodes</b> column.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Figure 4-12 Reducing nodes



5. Click **Next**.

6. Confirm the information and click **Submit**.
7. Click **Back to Cluster List** to go back to the **Clusters** page. **Task Status** is **Scaling in**. When **Cluster Status** changes to **Available**, the cluster has been successfully scaled in.

## Related Documents

To optimize costs, you may also reduce the capacity of a Logstash cluster by downgrading node specifications. For details, see [Changing the Node Specifications of a Logstash Cluster](#).

# 4.8 Managing Logstash Clusters

## 4.8.1 Viewing Logstash Cluster Information

On the CSS management console, you can check a Logstash cluster's status, version, nodes, configurations and settings, and resource allocation on the cluster list and the cluster information page.

### Introduction to Cluster Lists

A cluster list displays all clusters of the same type under the current account. If there are a large number of clusters, the cluster list will consist of multiple pages, showing clusters of all states.

Clusters are listed in chronological order by default in the cluster list, with the most recent clusters displayed at the top.

You can click **Export** to download the full cluster list.

**Table 4-29** Cluster list parameters

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name/ID        | Name and ID of a cluster. You can click a cluster name to switch to the <b>Cluster Information</b> page. The cluster ID is automatically generated by the system and uniquely identifies a cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Cluster Status | <p>Current status of a cluster.</p> <ul style="list-style-type: none"><li>• Available: The cluster is running properly.</li><li>• Creating: The cluster is being created.</li><li>• Processing: The cluster is being restarted, scaled out, scaled in, or its node specifications are being changed.</li><li>• Abnormal: The cluster failed to be created or is unavailable.</li></ul> <p><b>WARNING</b><br/>A cluster that is unavailable can be deleted, but cannot be scaled out. For such a cluster, you can check its monitoring data or restart it to try to fix the problem. If the fault persists, contact technical support.</p> |



| Parameter               | Description                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Task Status             | The status of tasks such as cluster restarting, scaling out, and scaling in, or ongoing pipeline tasks.                                                                                                                                                                                                                                                                         |
| Version                 | Cluster version and image version.<br>For example, if the version is <b>7.10.0(7.10.0_x.x.x_x.x.x)</b> , <b>7.10.0</b> is the cluster version and <b>7.10.0_x.x.x_x.x.x</b> enclosed in parentheses is the image version.<br><b>CAUTION</b><br>For a cluster created a long time ago, its image version may not be available because it is likely not recorded in the database. |
| Private Network Address | Private network address and port number of a cluster. You can use it to access the cluster. If the cluster has multiple nodes, the private network addresses and port numbers of all nodes are displayed.                                                                                                                                                                       |
| Enterprise Project      | Enterprise project that a cluster belongs to.                                                                                                                                                                                                                                                                                                                                   |
| Cluster Description     | Provides supplementary information about the cluster.                                                                                                                                                                                                                                                                                                                           |
| Created                 | Cluster creation time.                                                                                                                                                                                                                                                                                                                                                          |
| Billing Mode            | Billing mode of a cluster. It can be <b>Pay-per-use</b> or <b>Yearly/Monthly</b> .                                                                                                                                                                                                                                                                                              |
| Operation               | Operations that can be performed on a cluster, such as checking metrics, restart, and deletion. If an operation is unavailable, the menu button is grayed out.                                                                                                                                                                                                                  |

## Viewing Basic Cluster Information

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. On the **Overview** tab, check the cluster's basic information, configuration, and node information.

**Table 4-30** Cluster information

| Type                | Parameter | Description                                                                                                                                                                                   |
|---------------------|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Information | Name      | Cluster name. The name can be customized.<br><br>You can click  on the right to change the cluster name. |

| Type          | Parameter                     | Description                                                                                                                                                                                                                                                                                                                                                   |
|---------------|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               | ID                            | Unique ID of a cluster, which is automatically generated by the system.<br>Each cluster in the same region has a unique ID.                                                                                                                                                                                                                                   |
|               | Cluster Description           | Provides supplementary information about the cluster. You can click  on the right to modify the description.                                                                                                                                                               |
|               | Version                       | Cluster version.<br>For example, if the version is <b>7.10.0(7.10.0_x.x.x_x.x.x)</b> , <b>7.10.0</b> is the cluster version and <b>7.10.0_x.x.x_x.x.x</b> enclosed in parentheses is the image version.<br><b>CAUTION</b><br>For a cluster created a long time ago, its image version may not be available because it is likely not recorded in the database. |
|               | Cluster Status                | Current status of a cluster                                                                                                                                                                                                                                                                                                                                   |
|               | Task Status                   | Current task status of a cluster. If there is no ongoing task, -- 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 a cluster belongs                                                                                                                                                                                                                                                                                                                                |
|               | Subnet                        | Subnet to which the cluster belongs                                                                                                                                                                                                                                                                                                                           |

| Type             | Parameter            | Description                                                                                                                                                                                                                                                                                                                                            |
|------------------|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                  | Security Group       | Security group to which a cluster belongs.<br><br>To modify the security group of a cluster, click <b>Change Security Group</b> on the right.<br><b>WARNING</b><br>The selected security group must allow port 9600 in the inbound direction. Otherwise, the cluster may be inaccessible to external services.                                         |
|                  | Cluster Routing      | Cluster route information. You can view, add, or modify routes for a cluster. For details, see <a href="#">Configuring Routes for a Logstash Cluster</a> .                                                                                                                                                                                             |
|                  | Enterprise Project   | Enterprise project that a cluster belongs to.<br><br>You can click a project name to check the basic project information.                                                                                                                                                                                                                              |
|                  | Private IPv4 Address | Private IPv4 address and port number of a cluster. If the cluster has only one node, the IP address and port number of only one node are displayed, for example, <b>10.62.179.32:9600</b> . If the cluster has multiple nodes, the IP addresses and port numbers of all nodes are displayed, for example, <b>10.62.179.32:9600,10.62.179.33:9600</b> . |
| Node Information | 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                                                                                                                                                                                                                                                                                                                           |

## 4.8.2 Setting Tags for a Logstash Cluster

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

Tags can be added during or after cluster creation. If you want to use the same tag to identify multiple cloud resources for better resource grouping, we recommend that you predefine tags in Tag Management Service (TMS). For details, see [Predefined Tags Overview](#).

## Constraints

- Each cluster can have a maximum of 20 tags.
- If your organization has configured tag policies for CSS, add tags to clusters based on these policies. If a tag does not comply with existing tag policies, cluster creation may fail. Contact the administrator to learn more about these policies.

## Managing Tags

In CSS, you can add, delete, modify, and query cluster tags.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Select the **Tags** tab, where 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 **Edit Tag** in the upper left corner. In the displayed dialog box, click **Add Tag**, enter the tag key and value, and click **OK**.
  - Modify  
Click **Edit Tag** in the upper left corner. In the displayed dialog box, modify the key and value of a tag, and click **OK**.
  - Delete  
Click **Edit Tag** in the upper left corner. In the displayed dialog box, click **Delete** in the row that contains a tag, and click **OK**.

## Searching for Clusters by Tag

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, click the search box, and select a tag key and value to search for the target cluster.

You can select a tag key and value from the drop-down list. The system returns a list of clusters that exactly match the tag key and value. If you select multiple tags, only clusters including all of the selected tags will be returned.

The filter can include a maximum of 10 tags.

### 4.8.3 Associating a Logstash Cluster with an Enterprise Project

In CSS, you can associate your Elasticsearch cluster with 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 these users and user groups.

You can associate your cluster with an enterprise project either when creating the cluster or afterward.

Prerequisites

To use the enterprise project function, you need to assign permissions to the corresponding account. You can [submit a service ticket](#) to apply for the permissions.

You have [created an enterprise project](#).

Managing Enterprise Projects


- 1. Log in to the [CSS management console](#).
- 2. In the navigation pane on the left, choose **Clusters > Logstash**.
- 3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
- 4. On the **Overview** tab, find **Enterprise Project** in the **Configuration** area, and click the enterprise project name next to it. The **Project Management** page is displayed.

Figure 4-13 Enterprise Project



- 5. On the **Resources** tab, select the region of the current cluster, and select **CloudSearchService** as **Service**. Then click **Search**. The CSS clusters are displayed in the resource list below.

Figure 4-14 Filtering CSS clusters



- 6. Select the cluster whose enterprise project you want to modify and click **Remove**.
- 7. In the **Remove Resource** dialog box, specify **Mode** and select **Destination Enterprise Project**. Then click **OK**.

**Figure 4-15 Remove Resource**

**Remove Resource**

**Note**

1. Removing independent resources allows you to remove multiple types of resources from an enterprise project at the same time. To ensure correct consumption record of an enterprise project, remove EVS disks and EIPs from the enterprise project that contains their associated ECSs. EVS disks and EIPs will only be removed from a project together with the associated ECS when they are in the same enterprise project as the ECS.
2. The packages of some cloud services are associated with enterprise projects. Fees are deducted for only the packages purchased when the cloud service resources and packages belong to the same enterprise project.
3. If you migrate cloud service resources to another enterprise project, namely, the cloud service resources and packages do not belong to the same enterprise project, fees are deducted for packages and resources separately.

Mode: **Independent resources** | ECSs and ECS associated resources

Destination Enterprise Project:

1 of the 1 selected resources can be removed.

| Resource Name           | Project | Region | Service                | Resource Type |
|-------------------------|---------|--------|------------------------|---------------|
| css-b69600d9c1dc1a33... |         |        | CloudSearchService CSS | Cluster       |

8. After the change is complete, return to the Logstash cluster list on the CSS console and check the cluster's new enterprise project.

**Figure 4-16 Enterprise Project**

| Name/ID                           | Clust...    | Task... | Version | Private IP... | Secur... | Enterprise Project |
|-----------------------------------|-------------|---------|---------|---------------|----------|--------------------|
| css-... <b>NEW</b><br>60a7b7a3... | ○ Availa... | --      | ...     | 10... 1...    | Disabled | d...t              |

## 4.8.4 Changing the Node Specifications of a Logstash Cluster

If the workloads on the data plane of a Logstash cluster change, you can scale the cluster vertically by changing its node specifications.

Typically, you increase node specifications instead of decreasing them. Common scenarios include:

- In the face of fast data growth, node specifications can be increased for a Logstash cluster to increase data processing and storage capacities, thus enhancing overall performance.
- During peak hours, node specifications can be increased for a Logstash cluster to handle the extra load.

Decreasing node specifications will reduce the cluster's data processing and storage capacities. Exercise caution.

The process of changing node specifications is as follows:


1. Bring a node offline.
2. Change the node specifications.
3. Restart the node.
4. After the node recovers, proceed to another node and repeat the steps above. This goes on until all nodes are changed.

The node specifications are changed one node at a time. This is to ensure that there are sufficient resources to keep services running.

## Prerequisites

- The cluster status is **Available**, and there are no ongoing tasks.
- Your CSS resource quotas are sufficient for the capacity expansion you are about to perform. You can check available resources on the **Modify Configuration** page.

**Figure 4-17** Checking available resources



## Impact on Billing

For a pay-per-use cluster, you can see its new price when confirming the node specifications or storage type change on the console. After the scale-out or scale-up is complete, the new price will apply. For pricing details, see [Cloud Search Service Price Calculator](#).

For a yearly/monthly cluster, you can also see the new price when you confirm the change. In the case of an upgrade, a new order is triggered when you confirm the change; and in the case of a downgrade, a refund request is triggered when you do so.

## Constraints

Node specifications can be changed only for clusters that have at least two nodes.

## Change Impact

Before changing a cluster's node specifications, it is essential to assess the potential impacts and review operational recommendations. This enables proper scheduling of the change, minimizing service interruptions.

- Impact on cluster throughput  
Taking nodes offline one at a time usually does not interrupt services. However, jobs that are running on these will be stopped, causing the cluster throughput to decrease. To minimize this impact, perform this operation when the Logstash cluster has no ongoing tasks or after temporarily suspending all its tasks.
- Characteristics of this process  
Once started, a change task cannot be stopped until it succeeds or fails.

## Change Duration

The following formulas can be used to estimate how long the node specifications change will take:

**Change duration (min) = 10 (min) x Number of nodes**

where, 10 minutes is the duration per node. It is an empirical value.


## Changing Node Specifications

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, find the target cluster, and choose **More > Modify Configuration** in the **Operation** column. The **Modify Configuration** page is displayed.
4. Select the **Scale Cluster** tab, and click **Change specifications** to set parameters.

**Table 4-31** Changing node specifications

| Parameter | Description                                                                                                                                                                                                                            |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Action    | Select <b>Change specifications</b> .                                                                                                                                                                                                  |
| Resources | Shows the change of resources for this operation.                                                                                                                                                                                      |
| Nodes     | Configure the changes you want to make. <ol style="list-style-type: none"><li>1. Select <b>Logstash Nodes</b> in the <b>Node Type</b> column.</li><li>2. In the <b>Node Specifications</b> column, select a new node flavor.</li></ol> |

**Figure 4-18** Changing cluster specifications



5. Click **Next**.
6. Confirm the information and click **Submit**.
7. In the dialog box that is displayed, choose whether to select **Cluster status check**, and click **OK** to start the specifications change.

Cluster status check: The cluster status is checked by default before a node specifications change. The specifications of nodes are changed one by one to ensure success and data security. If the cluster is overloaded and services are faulty, the request for a node specifications change may fail to be delivered. In this case, you can disable cluster status check. Note that if you ignore cluster status check before the node specifications change, the cluster may become faulty and services may be interrupted without your knowledge. Exercise caution.

8. Click **Back to Cluster List** to go back to the **Clusters** page. The **Cluster Status** is **Configuration modified**. When **Cluster Status** changes to **Available**, the cluster node specifications have been successfully changed.



## 4.8.5 Managing CA Certificates for a Logstash Cluster

When connecting a Logstash cluster to an Elasticsearch cluster of CSS, you can skip this section and use a default CA certificate. To connect it to other data sources that require security authentication, you need to use a customize certificate by following the procedures described in this topic.

### Scenario

- When a Logstash cluster connects to an Elasticsearch cluster that uses the secure mode and the HTTPS protocol, the default CA certificate preconfigured in CSS can be used to trust the server. Logstash clusters support [Viewing Default Certificates](#).
- If the Logstash cluster connects to other data sources that require security authentication, you need to upload a custom certificate to trust the server. For details, see [Uploading a Custom Certificate](#).

### Viewing Default Certificates

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Certificate Management > Default Certificates**.
5. On the **Default Certificates** tab, check the default CA certificate.

### Uploading a Custom Certificate

1. Prepare the custom certificate.
  - The certificate name contains 4 to 32 characters, and must start with a letter, and end with **.cer/.crt/.rsa/.jks/.pem/.p10/.pfx/.p12/.csr/.der/.keystore**. The value can contain letters, digits, hyphens (-), underscores (\_), and periods (.). Other special characters are not allowed.
  - Up to 50 certificates can be uploaded.
  - The certificate file size cannot exceed 1 MB.
2. Upload the custom certificate to an OBS bucket. For details, see [Uploading an Object](#).
3. Log in to the [CSS management console](#).
4. In the navigation pane on the left, choose **Clusters > Logstash**.
5. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
6. Choose **Certificate Management > Custom Certificates**.
7. On the **Custom Certificates** page, click **Upload Certificate**. In the displayed dialog box, configure the certificate.
  - **OBS Bucket**: Select the OBS bucket that holds the custom certificate.
  - **Certificate Object**: Click **Select**. In the **Select Certificate Object** dialog box, select the custom certificate that has been uploaded to the OBS bucket in advance and click **OK**.

8. Click **OK** to upload the custom certificate.  
After the custom certificate is uploaded, it is displayed in the certificate list.
9. To delete a custom certificate that you no longer need, click **Delete** in the **Operation** column, and then click **OK** to confirm the deletion.

## 4.8.6 Forcibly Restarting a Logstash Cluster

If a Logstash cluster is faulty due to long-term running or other unknown reasons, you can forcibly restart the cluster to restore it.

### Constraints

- A cluster becomes unavailable during restart. Exercise caution.
- For a cluster in working state, the Logstash process is stopped during restart.
  - If Keepalive is disabled for pipelines, restarting the cluster will stop all running pipelines.
  - If Keepalived is enabled, Logstash will try restoring these pipelines. If the Logstash process is restarted within 10 minutes, the pipelines will be restored to running. Otherwise, the restoration fails.

### Forcibly Restarting a Cluster

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, find the target cluster, and choose **More > Force Restart** in the **Operation** column.
4. In the **Force Restart Cluster VMs** dialog box, confirm the restart information and click **OK**.

During the cluster 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.

## 4.8.7 Deleting a Logstash Cluster

After using a Logstash cluster for data migration, you can delete the cluster if you will not be using it again for a long time.

### Constraints

Deleting a cluster will clear all its data. Exercise caution.

### Deleting a Pay-per-Use Cluster

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, find the target cluster, and choose **More > Delete** in the **Operation** column.
4. In the displayed dialog box, manually type in **DELETE**, and click **OK**.

## Deleting a Yearly/Monthly Cluster

To delete a yearly/monthly cluster, unsubscribe from it. After unsubscription, the cluster's resources and data will be permanently deleted. If you want to retain data, manually back up data before submitting the unsubscription request.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, find the target cluster, and choose **More > Unsubscribe/Release** in the **Operation** column.
4. In the displayed dialog box, manually type in **YES**, and click **OK**.  
On the **Unsubscribe from Resource** page, confirm the resource information and refund amount.
5. Select a reason for unsubscription, select the acknowledgement check boxes, and click **Confirm**.

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

### NOTE

If the cluster is in **Available** state, an unsubscription order will be generated for refund, and the cluster will be deleted. If the cluster has expired or has been frozen, it will be deleted directly. For more information, see [Billing Center](#).

## 4.9 Logstash Cluster Monitoring and Log Management

### 4.9.1 Monitoring Metrics for Logstash Clusters in Cloud Eye

You can use Cloud Eye to monitor cluster metrics of CSS in real time and quickly handle exceptions.

#### Function

This section describes CSS metrics that can be monitored by Cloud Eye as well as their namespaces and dimensions. You can query metrics and alerts generated by CSS on the [Cloud Eye console](#) or using an [API](#)

- If the configuration center of a Logstash cluster does not have any records of operations on the pipeline list, the monitoring records of this Logstash cluster are empty.
- When the **Events** data of a pipeline changes dynamically, the monitoring data changes accordingly. When a pipeline task is being started or stopped, or the **Events** data is stable, the monitoring data remains unchanged.

#### Namespaces

SYS.ES

CSS.CUSTOM

Cluster Monitoring Metrics

 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 4-32 CSS metrics

| Metric ID              | Metric                     | Description                                                         | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|------------------------|----------------------------|---------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| max_jvm_heap_usage     | Max. JVM Heap Usage        | Maximum JVM heap usage of nodes in a CSS cluster                    | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| max_jvm_young_gc_time  | Max. JVM Young GC Duration | Maximum accumulated JVM Young GC duration of nodes in a CSS cluster | ≥ 0         | ms    | N/A           | CSS cluster                  | 1 minute                  |
| max_jvm_young_gc_count | Max. JVM Young GC Count    | Maximum accumulated JVM Young GC count of nodes in a CSS cluster    | ≥ 0         | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_jvm_old_gc_time    | Max. JVM Old GC Duration   | Maximum accumulated JVM Old GC duration of nodes in a CSS cluster   | ≥ 0         | ms    | N/A           | CSS cluster                  | 1 minute                  |
| max_jvm_old_gc_count   | Max. JVM Old GC Count      | Maximum accumulated JVM Old GC count of nodes in a CSS cluster.     | ≥ 0         | Count | N/A           | CSS cluster                  | 1 minute                  |
| max_cpu_usage          | Max. CPU Usage             | Maximum node CPU usage in a CSS cluster                             | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID            | Metric                                 | Description                                                                    | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|----------------------|----------------------------------------|--------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| max_load_average     | Max. Node Load                         | Maximum number of average queued tasks per minute on nodes in a cluster        | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_cpu_usage        | Avg. CPU Usage                         | Average node CPU usage in a CSS cluster.                                       | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| avg_load_average     | Avg. Node Load                         | Average number of queued tasks per minute on nodes in a CSS cluster            | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_jvm_heap_usage   | Avg. JVM Heap Usage                    | Average node JVM heap usage in a CSS cluster                                   | 0~100       | %     | N/A           | CSS cluster                  | 1 minute                  |
| avg_jvm_old_gc_count | Avg. GCs of Old-Generation JVM         | Average number of old-generation garbage collections of nodes in a CSS cluster | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_jvm_old_gc_time  | Avg. GC Duration of Old-Generation JVM | Average old-generation garbage collection duration of nodes in a CSS cluster   | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |

| Metric ID              | Metric                                          | Description                                                                                             | Value Range | Unit  | Number System | Monitored Object (Dimension) | Monitoring Interval (Raw) |
|------------------------|-------------------------------------------------|---------------------------------------------------------------------------------------------------------|-------------|-------|---------------|------------------------------|---------------------------|
| avg_jvm_young_gc_count | Avg. GCs of Young-Generation JVM                | Average number of young-generation garbage collections of nodes in a CSS cluster                        | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| avg_jvm_young_gc_time  | Avg. GC Duration of Young-Generation JVM        | Average young-generation garbage collection duration of nodes in a CSS cluster                          | $\geq 0$    | ms    | N/A           | CSS cluster                  | 1 minute                  |
| sum_events_in          | Total Records Passed Through the Input Plug-in  | Total number of records that have passed through the <b>input</b> plugin on all the nodes in a cluster  | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_events_filtered    | Total Records Passed Through the Filter Plug-in | Total number of records that have passed through the <b>filter</b> plugin on all the nodes in a cluster | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |
| sum_events_out         | Total Records Passed Through the Out Plug-in    | Total number of records that have passed through the <b>out</b> plugin on all the nodes in a cluster    | $\geq 0$    | Count | N/A           | CSS cluster                  | 1 minute                  |

Node Monitoring Metrics

Table 4-33 CSS node monitoring metrics

| Metric ID        | Metric                                  | Description                                         | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|------------------|-----------------------------------------|-----------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| jvm_heap_usage   | JVM Heap Usage                          | JVM heap memory usage of a node.                    | 0~100       | %     | N/A           | CSS cluster - cloud service node | 1 minute                  |
| cpu_usage        | CPU Usage                               | CPU usage.                                          | 0~100       | %     | N/A           | CSS cluster - cloud service node | 1 minute                  |
| load_average     | Average Load                            | Average number of queued tasks per minute on a node | ≥ 0         | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| jvm_old_gc_count | Total GCs of Old-Generation JVM         | Number of old-generation garbage collection times   | ≥ 0         | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| jvm_old_gc_time  | Total GC Duration of Old-Generation JVM | Time spent on old-generation garbage collection     | ≥ 0         | ms    | N/A           | CSS cluster - cloud service node | 1 minute                  |

| Metric ID          | Metric                                            | Description                                                                                 | Value Range | Unit  | Number System | Monitored Object (Dimension)     | Monitoring Interval (Raw) |
|--------------------|---------------------------------------------------|---------------------------------------------------------------------------------------------|-------------|-------|---------------|----------------------------------|---------------------------|
| jvm_young_gc_count | Total GCs of Young-Generation JVM                 | Number of young-generation garbage collection times                                         | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| jvm_young_gc_time  | GC Duration of Young-Generation JVM               | Time spent on young-generation garbage collection                                           | $\geq 0$    | ms    | N/A           | CSS cluster - cloud service node | 1 minute                  |
| events_in          | Records Passed Through the Input Plug-in on Node  | Number of data records that have passed through the <b>input</b> plugin on the current node | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| events_filtered    | Records Passed Through the Filter Plug-in on Node | Number of records that have passed through the <b>filter</b> plugin on the current node     | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |
| events_out         | Records Passed Through the Out Plug-in on Node    | Number of records that have passed through the <b>out</b> plugin on the current node        | $\geq 0$    | Count | N/A           | CSS cluster - cloud service node | 1 minute                  |



Logstash Pipeline Monitoring Metrics

Table 4-34 Logstash pipeline monitoring metrics

| Metric ID                         | Metric                                                 | Description                                                                                                          | Value Range | Unit  | Number System | Monitored Object (Dimension)                                                   | Monitoring Interval (Raw) |
|-----------------------------------|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|--------------------------------------------------------------------------------|---------------------------|
| logstash_pipeline_events_in       | Records Passed Through the Input Plug-in into Pipeline | Number of records that have passed through the <b>input</b> plug-in during the current pipeline monitoring period    | ≥ 0         | Count | N/A           | CSS cluster - CSS node - Logstash pipeline; or CSS cluster - Logstash pipeline | 1 minute                  |
| logstash_pipeline_events_filtered | Records Filtered for Pipeline                          | Number of records that have passed through the <b>filtered</b> plug-in during the current pipeline monitoring period | ≥ 0         | Count | N/A           | CSS cluster - CSS node - Logstash pipeline; or CSS cluster - Logstash pipeline | 1 minute                  |

| Metric ID                    | Metric                                                 | Description                                                                                                     | Value Range | Unit  | Number System | Monitored Object (Dimension)                                                   | Monitoring Interval (Raw) |
|------------------------------|--------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|-------------|-------|---------------|--------------------------------------------------------------------------------|---------------------------|
| logstash_pipeline_events_out | Records Passed Through the Out Plug-in Out of Pipeline | Number of records that have passed through the <b>out</b> plug-in during the current pipeline monitoring period | $\geq 0$    | Count | N/A           | CSS cluster - CSS node - Logstash pipeline; or CSS cluster - Logstash pipeline | 1 minute                  |

## Dimension

**Table 4-35** Dimension description

| Key                                  | Value                                     |
|--------------------------------------|-------------------------------------------|
| cluster_id                           | CSS cluster ID                            |
| cluster_id,instance_id               | CSS cluster node ID                       |
| cluster_id,instance_id,pipeline_name | Logstash pipeline name, node, CSS cluster |
| cluster_id,pipeline_name             | Logstash pipeline name, CSS cluster       |

### 4.9.2 Using Cloud Eye to Monitor Logstash Clusters

You can use Cloud Eye to monitor created CSS clusters. After configuring 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 via email or HTTP/HTTPS.
2. **Configuring Monitoring Metrics:** Configure monitoring metrics for a cluster or a node in the cluster.
3. **Viewing Cluster Monitoring Information:** Monitor the trends and changes of the monitoring metrics over specific periods.

## 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.

## Prerequisites

- The cluster is in the **Available** or **Processing** state.
- The cluster has been running properly for more than 10 minutes.

## 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 to filter alarm rules.

If no alarm rules are available, create some by referring to [Creating an Alarm Rule](#). [Table 4-36](#) and [Table 4-37](#) describes the key parameters. For more parameters, customize them as needed.

**Table 4-36** Alarm rule configuration (clusters and cluster nodes)

| Parameter      | Description                                                                                                                                                                                                                                                      |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Alarm Type     | Select <b>Metric</b> .                                                                                                                                                                                                                                           |
| Cloud Service  | Select <b>Cloud Search Service</b> .                                                                                                                                                                                                                             |
| Resource Layer | Select the layer by which you are going to configure alarm rules. <ul style="list-style-type: none"><li>• <b>CSS Clusters:</b> Configure alarm rules by cluster.</li><li>• <b>CSS Clusters - CSS Instances:</b> Configure alarm rules by cluster node.</li></ul> |

**Table 4-37** Alarm rule configuration (Logstash pipelines)

| Parameter  | Description            |
|------------|------------------------|
| Alarm Type | Select <b>Metric</b> . |

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                        |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cloud Service  | Select <b>CSS.CUSTOM</b> .                                                                                                                                                                                                                                                                                                                         |
| Resource Layer | Select the layer by which you are going to configure alarm rules. <ul style="list-style-type: none"><li>• <b>cluster_id,instance_id,pipeline_name</b>: Configure alarm rules by Logstash pipeline name of CSS cluster nodes.</li><li>• <b>cluster_id,pipeline_name</b>: Configure alarm rules by Logstash pipeline name of CSS clusters.</li></ul> |

## Configuring Monitoring Metrics


1. Create a monitoring dashboard by referring to [Creating a Dashboard](#). If an available monitoring dashboard has been created, skip this step.
2. Add CSS monitoring graphs by referring to [Adding a Graph](#).

## Viewing Cluster Monitoring Information

In the cluster list, check the monitoring information of clusters and cluster nodes.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, find the target cluster, and choose **More > Monitoring Metrics** in the **Operation** column.
4. On the **Monitoring Metrics** page, check the cluster and node metrics.

**Figure 4-19** Monitoring Metrics



- Click the **Instances** tab to check cluster metrics.
  - **Instances**: Select a cluster to check its metrics. The current cluster is selected by default. Multiple clusters can be selected so that you may compare between them.
  - You can also set monitoring periods so that you can compare metrics between different periods for the same cluster.

- Add Graph Group: You can add a custom group in addition to the default group. When you move the pointer over a group name on the left, the icons for changing the group name and deleting the group are displayed on the right. You can modify the group.
- Select Metric: In the displayed **Select Metric** dialog box, set **All Metrics** and **Top N Metrics**.
- Click the **CSS Instances** tab to check the metrics of cluster nodes. Only the node monitoring information of the first cluster selected on the **Instances** tab is displayed.
  - CSS Instances: Select the cluster nodes whose metrics and information you want to check. By default, the first node of the current cluster is selected. Multiple nodes can be selected so that you may compare between them.
  - You can also set monitoring periods so that you can compare metrics between different periods for the same cluster.
  - Add Graph Group: You can add a custom group in addition to the default group. When you move the pointer over a group name on the left, the icons for changing the group name and deleting the group are displayed on the right. You can modify the group.
  - Select Metric: In the displayed **Select Metric** dialog box, set **All Metrics**.

### 4.9.3 Managing Logstash Cluster Logs

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

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

#### Impact on Billing

When log backup is enabled, the generated log backups are stored in OBS buckets, which will result in additional costs. For details, see [Object Storage Service Billing](#).

#### Prerequisites

The OBS bucket used for storing log backups has been created. The OBS bucket must meet the following requirements:

- **Storage Class: Standard.**
- **Region:** the same as that of the cluster.

## Querying Logs

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Logs > Log Search**. The **Log Search** page is displayed.  
You can search log records by node or keyword. For a detailed description of each type of logs, see [Log Introduction](#).

When a log file reaches 128 MB or when the time reaches 00:00 UTC, the system automatically compresses and archives it. Only unarchived logs appear on the log search page, while archived logs remain accessible through the log backup function.

## Backing Up Logs

Cluster logs can be backed up to OBS buckets, where you can download them for in-depth analysis at any time.

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, choose **Clusters > Logstash**.
3. In the cluster list, click the name of the target cluster. The cluster information page is displayed.
4. Choose **Logs > Log Backup**. The **Log Backup** page is displayed.
5. Enable log backup.  
Perform the following steps to enable log backup. If it is already enabled, skip this step.
  - a. Click **Enable Backup**. In the displayed dialog box, configure necessary settings.

**Table 4-38** Log backup settings

| Parameter  | Description                                                                                                                                                                                                                                         |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBS Bucket | Select an OBS bucket for storing log backups from the drop-down list box.<br><br>If no OBS buckets meet your requirements, click <b>Create Bucket</b> to go to the OBS console and create one. For details, see <a href="#">Creating a Bucket</a> . |

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                                       |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup Path | <p>Set the log storage location in the OBS bucket.</p> <p>The backup path cannot:</p> <ul style="list-style-type: none"><li>• Contain the following characters: \:*?"&lt;&gt; '{}.</li><li>• Start with a slash (/).</li><li>• Start or end with a period (.).</li><li>• Contain more than two consecutive slashes (/) or periods (.).</li><li>• Exceed 512 characters.</li></ul> |

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IAM Agency | <p>To back up data to an OBS bucket, you must have the write permission to it. By configuring an IAM agency, you can authorize CSS to access its OBS resources through an associated account.</p> <ul style="list-style-type: none"><li>If you are configuring an agency for the first time, click <b>Automatically Create IAM Agency</b> to create <b>css-obs-agency</b>.</li><li>If there is an IAM agency automatically created earlier, you can click <b>One-click authorization</b> to have the OBS Administrator permissions deleted automatically, and have the following custom policies added automatically instead to implement more refined permissions control.<br/><pre>"obs:bucket:GetBucketLocation",<br/>"obs:object:GetObjectVersion",<br/>"obs:object:GetObject",<br/>"obs:object:DeleteObject",<br/>"obs:bucket:HeadBucket",<br/>"obs:bucket:GetBucketStoragePolicy",<br/>"obs:object:DeleteObjectVersion",<br/>"obs:bucket:ListBucketVersions",<br/>"obs:bucket:ListBucket",<br/>"obs:object:PutObject"</pre></li><li>When OBS buckets use SSE-KMS encryption, the IAM agency must be granted KMS permissions. You can click <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b> to have the following custom policies created automatically.<br/><pre>"kms:cmk:create",<br/>"kms:dek:create",<br/>"kms:cmk:get",<br/>"kms:dek:decrypt",<br/>"kms:cmk:list"</pre></li><li>To use <b>Automatically Create IAM Agency</b> and <b>One-click authorization</b>, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:roles:listRoles",<br/>"iam:agencies:getAgency",<br/>"iam:agencies:createAgency",<br/>"iam:permissions:listRolesForAgency",<br/>"iam:permissions:grantRoleToAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:revokeRoleFromAgency",<br/>"iam:roles:createRole"</pre></li><li>To use an IAM agency, the following minimum permissions are required:<br/><pre>"iam:agencies:listAgencies",<br/>"iam:agencies:getAgency",<br/>"iam:permissions:listRolesForAgencyOnProject",<br/>"iam:permissions:listRolesForAgency"</pre></li></ul> |




| Parameter | Description                                                                                                                                             |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
|           | <b>WARNING</b><br>The agency name can contain only letters (case-sensitive), digits, underscores (_), and hyphens (-). Otherwise, the backup will fail. |

Determine whether to enable **Automatic backup** based on service needs.

- To enable automatic, periodic log backup, select **Automatic backup**, and set the backup policy. For more information, see [6](#).
  - If manual log backup is sufficient, deselect **Automatic backup**.
- b. Click **OK** to enable log backup.
- The configuration information will be displayed on the **Log Backup** tab.
6. Back up logs. Two options are available: automatic or manual.
- Automatic log backup: Logs are backed up periodically based on the preset policy.
- a. On the **Log Backup** page, click **Modify Settings** on the right.

**Figure 4-20** Modifying Settings



- b. In the **Modify Settings** dialog box, select **Automatic backup**, and configure the automatic backup policy.

**Table 4-39** Automatic backup settings

| Parameter         | Description                                                                                                                                                                     |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time Zone         | Select a time zone for the backup start time.                                                                                                                                   |
| Backup Start Time | Specify the start time of auto backup.<br>Select a value from the drop-down list. The value range is from <b>00:00</b> to <b>23:00</b> . The backup always happens on the hour. |


- c. Click **OK** to enable automatic log backup.
- Automatic log backup tasks are displayed in the task list. When **Task Status** in the task list changes to **Succeeded**, the backup is successful.
- If log backup fails, click **Failed Tasks** learn the failure causes. A maximum of 10 failed tasks can be displayed. When log backup is disabled or the cluster is deleted, the failure records are also cleared.

- d. If automatic log backup is not required, click **Modify Settings**. In the displayed dialog box, deselect **Automatic backup**, and click **OK**. In the displayed confirmation dialog box, click **OK**.
- Disabling log backup does not automatically delete log backups automatically generated earlier. Instead, you need to manually delete them on the OBS console.

Manual log backup: Back up logs right away.

- a. On the **Log Backup** page, click **Back Up Manually** under **Log Backup Tasks**.

Figure 4-21 Manually backing up logs



- b. In the displayed dialog box, confirm the log backup path and click **OK**.  
The log backup task is displayed in the task list. When **Task Status** in the task list changes to **Succeeded**, the backup is successful.  
If log backup fails, click **Failed Tasks** learn the failure causes. A maximum of 10 failed tasks can be displayed. When log backup is disabled or the cluster is deleted, the failure records are also cleared.
7. Check the backed-up log files.  
Logs are backed up incrementally. After the backup is successful, you can access the target OBS bucket to obtain the full log files by clicking **Log Path**.

Figure 4-22 OBS bucket address




Table 4-40 lists the log types, where **clustername** indicates the cluster name.

Table 4-40 Log types

| Log Name                           | Description          |
|------------------------------------|----------------------|
| <i>clustername_deprecation.log</i> | Deprecation log file |
| <i>clustername-plain.log</i>       | Run log file         |

8. If the log backup function is no longer needed, you can disable it.

On the **Log Backup** page, click **Disable Backup**. In the displayed dialog box, click **OK**. Disabling log backup does not automatically delete existing log backups. Instead, you need to manually delete them on the OBS console.

## Log Introduction

**Table 4-41** Introduction to different log types

| Log Type         | Description                                                                                                                                                                                                                                                                                                  | Purpose                                                                 |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| Run logs         | Run logs record a cluster's node and pipeline status, such as source-destination connectivity, pipeline creation or modification, and pipeline running errors.                                                                                                                                               | Check run logs to troubleshoot pipeline errors.                         |
| Deprecation logs | Deprecation logs record deprecation warnings. Deprecation warnings are written to this log when you use APIs, configurations, or functions that are marked for removal in future versions. You cannot check deprecation logs on the console. To check them, you need to back them up to an OBS bucket first. | Check for APIs or features that are about to expire in future versions. |

- **Run log description**

Run logs record a cluster's node and pipeline status. For example, the log record below indicates that the destination cluster could not be reached. You need to check whether the cluster address is correct and whether the cluster status is normal.

**Figure 4-23** A sample of run logs

```
[2025-07-04T07:33:39.742Z][WARN][logstash.outputs.elasticsearch][random2es] Attempted to resurrect connection to dead ES instance, but got an error: { "url": "http://192.168.0.171:9200/", "error_type": "Logstash::Outputs::ElasticSearch::HttpClient::Pool::HostUnreachableError", "error": "Elasticsearch Unreachable [http://192.168.0.171:9200/][Maricore::SocketException] Connection timed out (Connection timed out)" }
[2025-07-04T07:33:46.840Z][WARN][logstash.outputs.elasticsearch][random2es] Attempted to resurrect connection to dead ES instance, but got an error: { "url": "http://192.168.0.228:9200/", "error_type": "Logstash::Outputs::ElasticSearch::HttpClient::Pool::HostUnreachableError", "error": "Elasticsearch Unreachable [http://192.168.0.228:9200/][Maricore::SocketException] Connection timed out (Connection timed out)" }
```

Log content:

- 1. Log generation time
- 2. Log level, which can be DEBUG, INFO, WARN, or ERROR
- 3. Log-generating module
- 4. Name of the log-generating node
- 5. Log content

- **Deprecation log description**

Deprecation logs record deprecation warnings.

**Figure 4-24** A sample of deprecation logs

<sup>1</sup>  
[2025-06-10T08:03:57.944][WARN ][deprecation.logstash.outputs.elasticsearch] <sup>2</sup>  
<sup>3</sup> Relying on default value of 'pipeline.ecs\_compatibility', <sup>4</sup>  
which may change in a future major release of Logstash. To avoid unexpected changes when upgrading Logstash, please explicitly declare your  
desired ECS Compatibility mode.

Log content:

- 1. Log generation time
- 2. Log level, which can only be DEPRECATION.
- 3. Log-generating module
- 4. Log content. The log record shown in the figure above indicates that the ECS compatibility mode was not explicitly declared for a plugin when Logstash was started.

## 4.10 Viewing Logstash Cluster Audit Logs

With CTS, you can record operations associated with CSS for subsequent query, audit, and backtracking purposes.

### Prerequisites

CTS has been enabled.

### Key Operations Recorded by CTS

**Table 4-42** Key operations recorded by CTS

| Operation                                                                  | Resource Type | Event Name                     |
|----------------------------------------------------------------------------|---------------|--------------------------------|
| Creating a cluster                                                         | cluster       | createCluster                  |
| Changing the billing mode of a cluster from pay-per-use to yearly/ monthly | cluster       | updateOndemandClusterTo-Period |
| Querying cluster details                                                   | cluster       | showClusterDetail              |
| Deleting a cluster                                                         | cluster       | deleteCluster                  |
| Changing a cluster name                                                    | cluster       | updateClusterName              |
| Querying the cluster list                                                  | cluster       | listClusters                   |
| Expanding cluster capacity                                                 | cluster       | updateExtendCluster            |
| Adding instances and expanding storage                                     | cluster       | updateExtendInstanceStorage    |
| Obtaining the instance specifications list                                 | cluster       | listFlavors                    |

| Operation                                    | Resource Type | Event Name              |
|----------------------------------------------|---------------|-------------------------|
| Querying all tags                            | cluster       | listClustersTags        |
| Querying the tags of a specified cluster     | cluster       | showClusterTag          |
| Adding tags for a cluster                    | cluster       | createClustersTags      |
| Deleting the tags of a cluster               | cluster       | deleteClustersTags      |
| Adding or deleting cluster tags in batches   | cluster       | updateBatchClustersTags |
| Scaling in a cluster                         | cluster       | updateShrinkNodes       |
| Changing the flavor of a specified node type | cluster       | updateFlavorByType      |
| Changing the security group                  | cluster       | changeSecurityGroup     |
| Creating a V2 cluster                        | cluster       | createClusterV2         |
| Restarting a cluster                         | cluster       | rebootCluster           |
| Creating a configuration file                | cluster       | createCnf               |
| Updating a configuration file                | cluster       | updateCnf               |
| Starting a data migration pipeline           | cluster       | startPipeline           |
| Stopping a data migration pipeline           | cluster       | stopPipeline            |
| Adding a custom template                     | cluster       | addFavorite             |
| Testing connectivity                         | cluster       | startConnectivityTest   |
| Querying the template list                   | cluster       | listTemplates           |
| Querying the configuration file list         | cluster       | listConfs               |
| Querying the pipeline list                   | cluster       | listPipelines           |
| Querying operation records                   | cluster       | listActions             |
| Querying the content of a configuration file | cluster       | showGetConfDetail       |

| Operation                              | Resource Type | Event Name               |
|----------------------------------------|---------------|--------------------------|
| Deleting a configuration file          | cluster       | deleteConf               |
| Deleting a custom template             | cluster       | deleteTemplate           |
| Hot-starting a data migration pipeline | cluster       | startHotPipeline         |
| Hot-stopping a data migration pipeline | cluster       | stopHotPipeline          |
| Updating cluster routes                | cluster       | updateRoute              |
| Obtaining a cluster route              | cluster       | getRoutes                |
| Uploading a certificate file           | cluster       | uploadCerts              |
| Querying the certificate list          | cluster       | listCerts                |
| Querying certificate information       | cluster       | showCertsDetail          |
| Deleting a certificate file            | cluster       | deleteCerts              |
| Enabling logging                       | cluster       | startLogs                |
| Disabling logging                      | cluster       | stopLogs                 |
| Querying the job list                  | cluster       | listLogsJob              |
| Querying basic log configurations      | cluster       | showGetLogSetting        |
| Modifying basic log configurations     | cluster       | updateLogSetting         |
| Enabling automatic log backup          | cluster       | startLogAutoBackupPolicy |
| Disabling automatic log backup         | cluster       | stopLogAutoBackupPolicy  |
| Backing up logs                        | cluster       | createLogBackup          |
| Querying logs                          | cluster       | showLogBackup            |

## Querying Real-Time Traces

After a management tracker is created on the CTS console, the system starts recording operations performed on cloud service resources. After a data tracker is

created, the system starts recording operations performed on data in OBS buckets. CTS retains operation records generated in the latest seven days.

To view or export operation records of the last seven days on the CTS console, see [Querying Real-Time Traces](#).

# 5 CSS Resource Monitoring

---

On the **Monitoring** and **Dashboard** pages of the CSS management console, you can check the resource monitoring metrics and alarms of CSS.

## Checking the Monitoring Dashboards

CSS's out-of-the-box monitoring dashboards offer a one-stop resource monitoring solution. By integrating cloud resource usage data and core service metrics, the monitoring dashboards quickly build unified, intuitive monitoring views for key services, enabling more effective resource monitoring.

The monitoring dashboards depend on the following:

- Metric collection: Collects key resource metrics for cloud services.
- Preset dashboards: Display prebuilt monitoring dashboards for key cloud services. There is no need to manually configure metrics.
- Data visualization: Intuitively shows resource usage, service status, and trends using charts.

To check CSS's monitoring dashboards, perform the following steps:

1. Log in to the [CSS management console](#).
2. In the navigation pane, choose **Monitoring** to view the default monitoring dashboards.

---

### CAUTION

- Monitoring dashboards are not yet available in all regions. For regions where they are unavailable, a message will be displayed indicating no support.
  - Custom metrics cannot be configured. Only preset dashboards are provided. To use a custom dashboard, click **More** in the upper-right corner to go to the Cloud Eye console and configure the dashboard.
  - Cross-region resource monitoring is not supported. You can only monitor resources in the current region.
- 

Monitoring dashboards cover the following core metrics:



- Resource metrics: cluster health status, average indexing speed, average query speed, maximum CPU usage, maximum JVM heap usage, maximum node loads, number of shards, maximum disk usage, total number of tasks in the write queue, and total number of tasks in the search queue.
- Data metrics: current, maximum, minimum, average, and sum values.
- Trend analysis: week-on-week (this week vs last week) and day-on-day (today vs yesterday) comparisons.


For how to use the monitoring dashboards, see [Overview of Dashboards](#).

## Checking Resource Alarms

Check CSS alarms on the **Dashboard** page.

1. Log in to the [CSS management console](#).
2. In the navigation tree on the left, choose **Dashboard**. Check CSS alarms in the **Alarm Management** area.

**Figure 5-1 Alarm Management**



- Click **All**. On the displayed **Alarm Records** page, check detailed alarm information.
  - Hover the mouse pointer over a resource name to check alarm information.
3. Click **Create Alarm Rule**. On the **Alarm Configuration** page, set alarm rules as needed.

If the preset alarm rules cannot meet your requirements, click **Create Alarm Rule** to create new rules. For details, see [Creating an Alarm Rule](#).

## Checking Resource Monitoring Metrics

Check CSS resource monitoring metrics on the **Dashboard** page.

1. Log in to the [CSS management console](#).
2. In the navigation tree on the left, choose **Dashboard**. Check resource monitoring metrics in the **Monitoring** area.

**Figure 5-2** Monitoring metrics



- Click **Configure**. In the displayed dialog box, set the metrics to be displayed, **Rollup** (aggregation method), and **Chart Type**.
- Click **More** to go to the Cloud Eye console to check more details.