Alibaba Cloud Elasticsearch

Instances

Issue: 20190912

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. **You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.**
2. **No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company, or individual in any form or by any means without the prior written consent of Alibaba Cloud.**
3. **The content of this document may be changed due to product version upgrades , adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.**
4. **This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults " and "as available" basis. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity , applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, special, or punitive damages, including lost profits arising from the use**

or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.

5. By law, all the content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of the Alibaba Cloud website, product programs, or content shall be used, modified , reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates . The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates).
6. Please contact Alibaba Cloud directly if you discover any errors in this document.

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 Danger: Resetting will result in the loss of user configuration data.
	This warning information indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 Warning: Restarting will cause business interruption. About 10 minutes are required to restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus , page names, and other UI elements.	Click OK.
`Courier font`	It is used for commands.	Run the `cd / d C :/ windows` command to enter the Windows system folder.
`Italics`	It is used for parameters and variables.	`bae log list --instanceid` *`Instance_ID`*
[] or [a|b]	It indicates that it is a optional value, and only one item can be selected.	`ipconfig` *`[-all|-t]`*

Style	Description	Example
{} or {a|b}	It indicates that it is a required value, and only one item can be selected.	`swich {stand | slave}`

Contents

1 Instance management

This topic describes the instance management feature of Alibaba Cloud Elasticsearch, including cluster monitoring, instance restart, refresh, and task list.

Manage instances

Alibaba Cloud Elasticsearch supports the Cluster monitoring, Restart instances, Refresh, and Tasks features for you to manage instances.

Cluster monitoring

Alibaba Cloud Elasticsearch supports cluster monitoring and sending alerts to users through SMS messages. You can customize the threshold for triggering alerts. For more information, see CloudMonitor alerts for Elasticsearch.

Restart instances

Alibaba Cloud Elasticsearch allows you to use the restart and force restart methods to restart instances. Follow these guidelines to select an appropriate restart method:

- `Prerequisi tes` : Before you restart an instance, make sure that the status of the Elasticsearch instance is Active (green flag), the instance has at least one index replica, and the resource usage is not high. You can go to the #unique_5 page to check the resource usage. Ensure that the Node CPU Usage (%) is 80% or lower, the Node Heep Memory Usage (%) is around 50%, and the Node Workload Within One Minute does not exceed the number of cores of the current data node.

  Restart: If the Elasticsearch instance is restarted by this method, it can continuously provide services during the restart process. However, the instance must meet the requirements in Prerequisites. The restart process is time-consuming.

  Notice:

  - Before you restart the instance, make sure that the status of the instance is Active (green flag). Otherwise, you have to use the force restart method to restart the instance.

- The CPU and memory usage of the Elasticsearch instance will experience a usage spike during the restart process. This may affect the stability of your service for a short period of time.
- The time that the restart process takes depends on the amount of data stored on the instance, the number of nodes, and the number of indexes and replicas. Elasticsearch cannot estimate the total amount of time required to restart an instance. However, you can check the progress of the restart process in Tasks.

· Force restart: If an Elasticsearch instance is restarted by this method, the services running on the instance may become unstable during the restart process. The restart process takes only a short period of time.

 Notice:
When the disk usage exceeds 85%, the status of the Elasticsearch instance may change to a yellow or red flag. If a yellow or red flag is displayed, you cannot use the restart method to restart the instance. You can only forcibly restart the instance.

- When a yellow or red flag is displayed, we recommend that you do not perform these operations on the instance: upgrade nodes, upgrade disk space, restart, reset password, and other operations that may change the configuration of the instance. Perform these operations only after the status of the instance changes to a green flag.
- If you update the configuration of an Elasticsearch instance with a yellow or red flag and the instance contains two or more nodes, the instance will be constantly in the Initializing state. You can submit a ticket to contact the Alibaba Cloud Elasticsearch Technical Support to resolve this issue.

Refresh

You can use this feature to manually refresh the information displayed in the console . For example, if the console fails to display the status of the Elasticsearch instance that you have just created, use the refresh feature to update the status.

Tasks

You can click the Tasks icon to view the progress of tasks, such as the instance creation or restart progress.

- No task is running on the current instance.

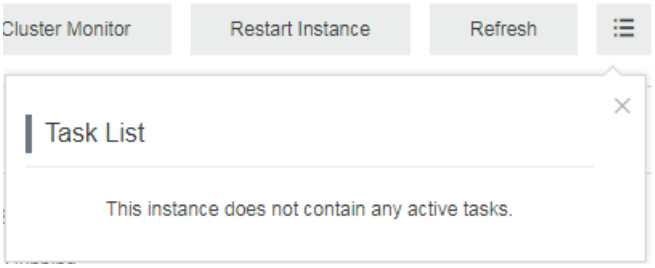

- Tasks that are running on the current instance.

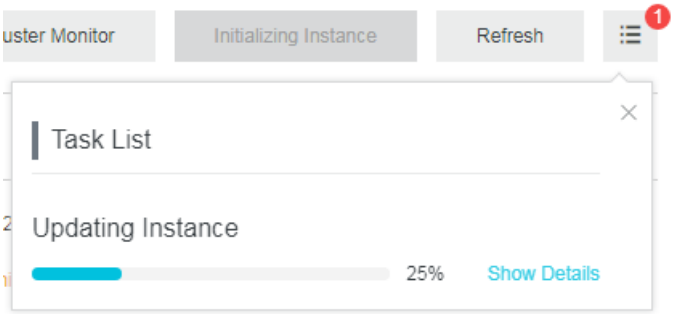

- Show detailed information about a running task.

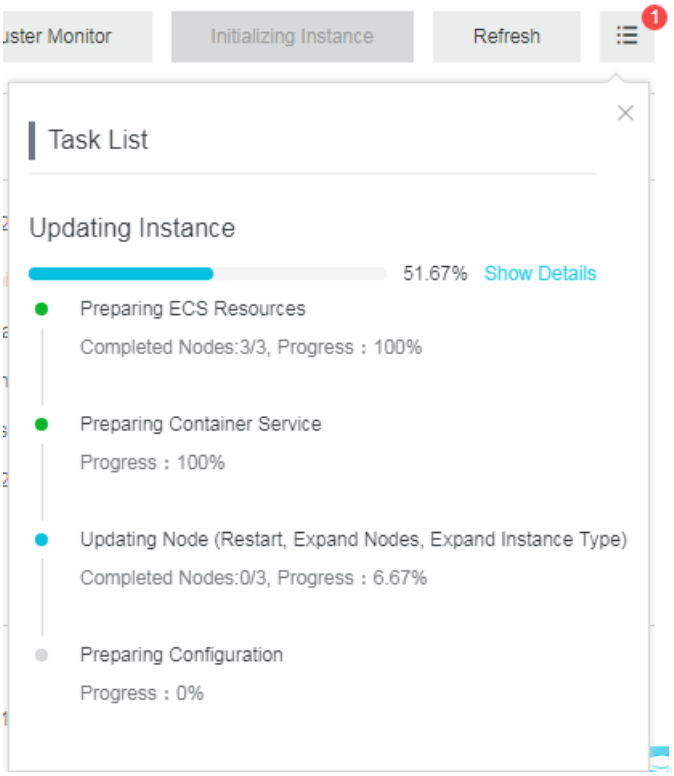

2 Basic Information

2.1 Basic information

Elasticsearch subscription instances

The following figure shows the information of an Alibaba Cloud Elasticsearch instance that uses the subscription billing method. For parameter descriptions, see the following sections and #unique_10.

- Name: By default, the name of an Alibaba Cloud Elasticsearch instance is the same as its ID. You can edit the name of the instance. You can also search instances by name.
- Internal Network Address: You can use the IP address of a VPC-connected ECS instance to access an Alibaba Cloud Elasticsearch instance.

  Notice:
  If you access an Alibaba Cloud Elasticsearch instance through the Internet, data security is not guaranteed. To protect your data, we recommend that you purchase an ECS instance that is connected to the same VPC network as your Elasticsearch instance. You can then use an internal network address to access the Elasticsearch instance.
- Internal Network Port: The following ports are supported:
  - Port `9200` for HTTP and HTTPS.
  - Port `9300` for TCP. Only Alibaba Cloud Elasticsearch 5.5.3 with Commercial Feature supports this port.

    Note:
    You cannot use the transport client to access Alibaba Cloud Elasticsearch 6.3.2 with Commercial Feature and Alibaba Cloud Elasticsearch 6.7.0 with Commercial Feature through port `9300`.
- Public Network Access: You can use public network addresses to access Alibaba Cloud Elasticsearch instances.

- Public Network Port: The following ports are supported:
  - Port `9200` for HTTP and HTTPS.
  - Port `9300` for TCP. Only Alibaba Cloud Elasticsearch 5.5.3 with Commercial Feature supports this port.

    Note:

    ■ You cannot use the transport client to access Alibaba Cloud Elasticsearch 6.3.2 with Commercial Feature and Alibaba Cloud Elasticsearch 6.7.0 with Commercial Feature through port `9300` .

    ■ To access an Elasticsearch instance through the Internet, you must configure the #unique_11/unique_11_Connect_42_section_ux5_yct_zgb. By default, the public network access feature forbids all IP addresses.
- Protocol: By default, HTTP is selected. You can click Edit to change the protocol. Currently, you can choose HTTP or HTTPS. For more information, see #unique_12/unique_12_Connect_42_section_i7x_sqt_enx.
- Renew: You can click Renew on the right side of Basic Information to renew the instance. You can renew your subscription one or more months. The minimum renewal period is one month.

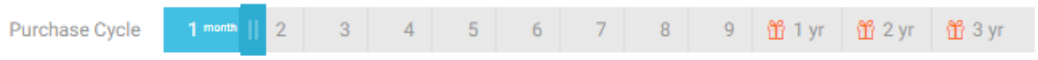

Elasticsearch pay-as-you-go instances

The following figure shows the basic information of an Alibaba Cloud Elasticsearch instance that uses the pay-as-you-go billing method. For parameter descriptions, see Elasticsearch subscription instances and #unique_10.

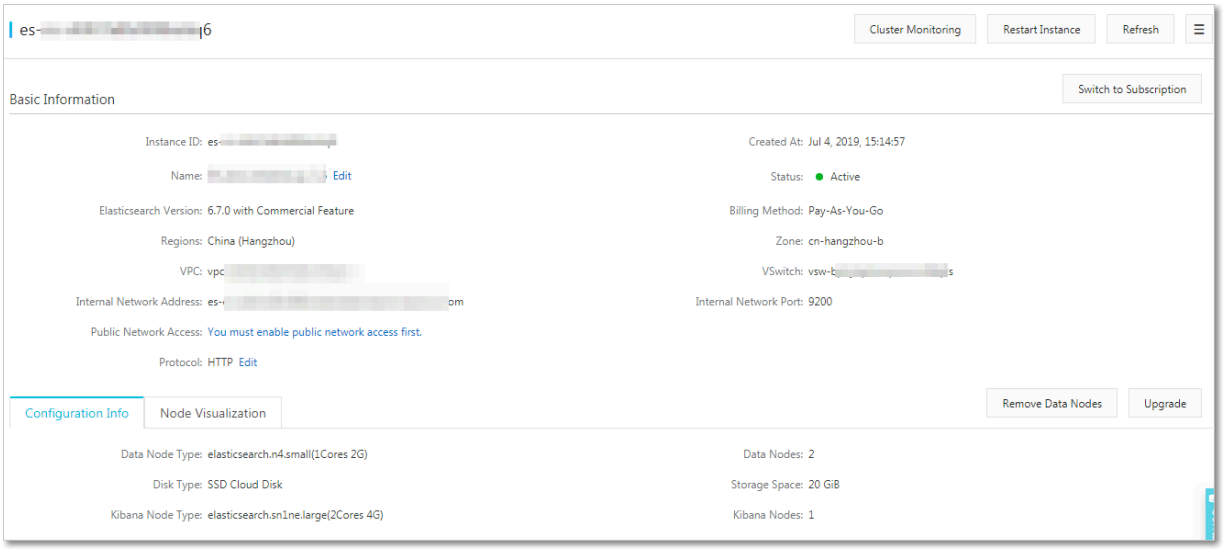

You can switch an Alibaba Cloud Elasticsearch instance from pay-as-you-go to subscription. To perform this task, click Switch to Subscription on the right side of Basic Information, and follow the instructions to switch the billing method.

Configuration information

For more information about parameter descriptions, see #unique_13.

Node visualization

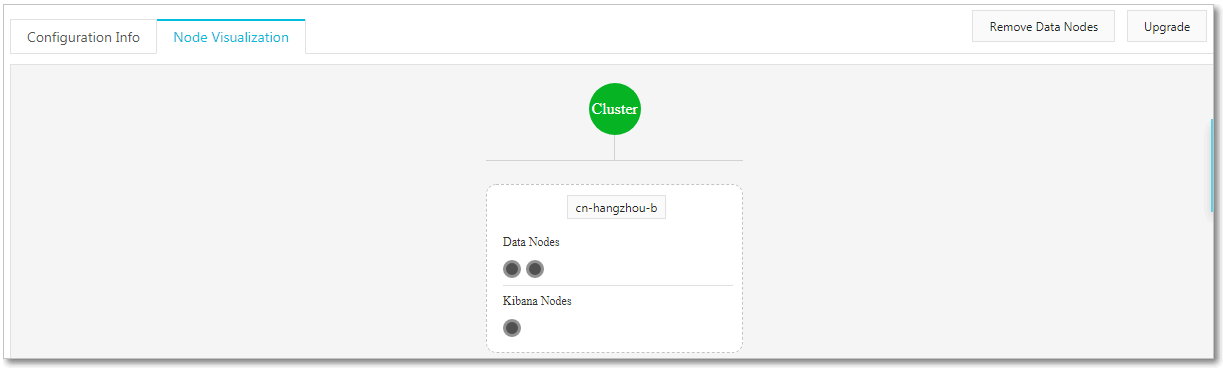

Remove data nodes

Currently, you can downgrade data nodes for Elasticsearch pay-as-you-go instances and Elasticsearch instances deployed in one zone. Elasticsearch subscription instances and instances deployed across zones are not supported. This function only allows you to remove data nodes from an Alibaba Cloud Elasticsearch instance. You

cannot downgrade the specification or disk space of dedicated master nodes, client nodes, and Kibaba nodes. For more information, see #unique_14.

Upgrade

You can upgrade the instance specification, number of nodes, dedicated master node specification, and storage space per data node for an Elasticsearch instance. For more information, see #unique_15.

2.2 Downgrade data nodes

You can only downgrade data nodes in an Alibaba Cloud Elasticsearch instance that uses the Pay-As-You-Go billing method and is deployed in one zone. You cannot downgrade data nodes in an instance that uses the Subscription billing method or that is deployed across zones. Currently, Alibaba Cloud Elasicsearch only supports removing data nodes from an Alibaba Cloud Elasticsearch instance. The specification and disk capacity of dedicated master nodes, client nodes, and Kibana nodes cannot be downgraded.

Procedure

1. Log on to the Alibaba Cloud Elasticsearch console, locate the Elasticsearch instance that you need to downgrade data nodes for, and click the instance ID.
2. On the Basic Information tab page, click Downgrade Data Nodes.

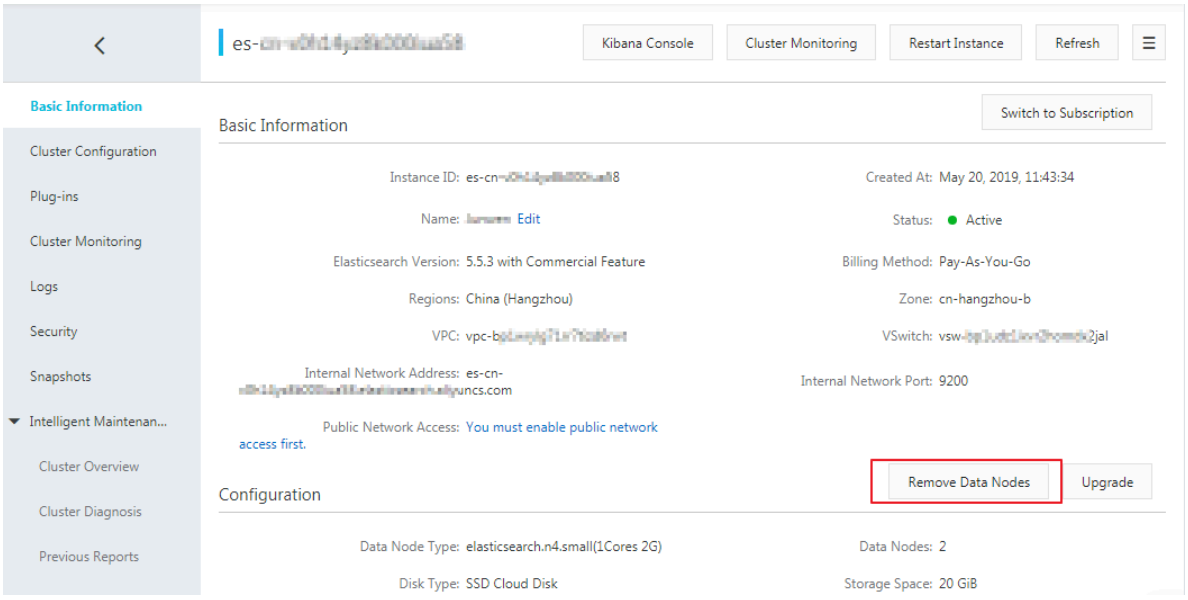

3. On the Downgrade Data Nodes page, select Data Node, and then specify the data nodes to be downgraded.

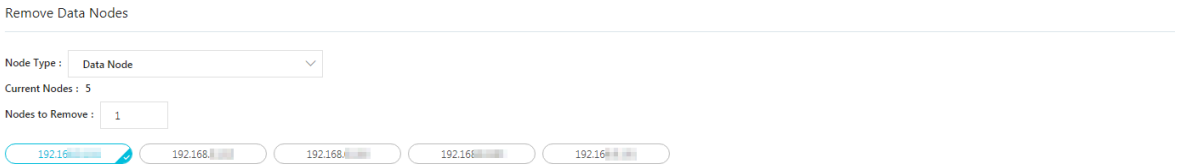

Note:

For data security, make sure that no data is stored on these data nodes. If the data nodes still contain data, click Data Migration Tool to migrate the data. After the data migration process is complete, no index data is stored on the data nodes. New index data is not written into these data nodes.

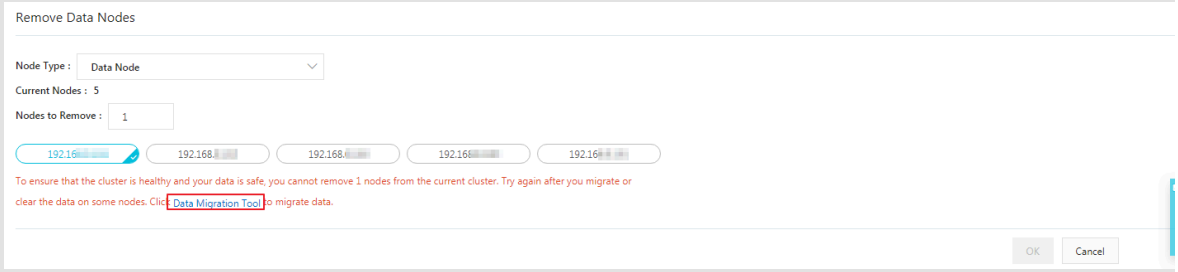

You can choose the smart migration or custom migration method to migrate the data:

- Smart migration

  The system will automatically select the data nodes to be downgraded for you. You must select the check box to agree to the terms of data migration, and then click OK.

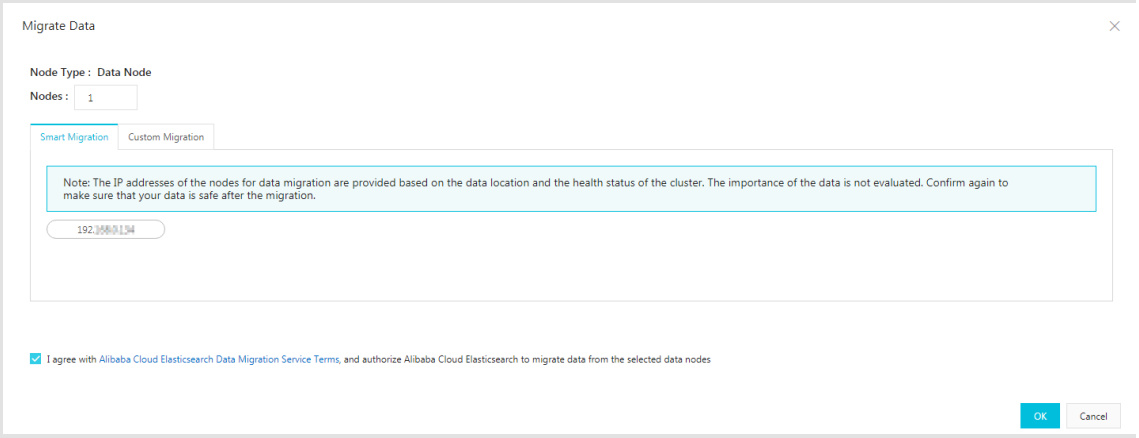

- Custom migration

  You need to manually specify the data nodes to be downgraded on the Custom Migration page, select the check box to agree to the terms of data migration, and click OK.

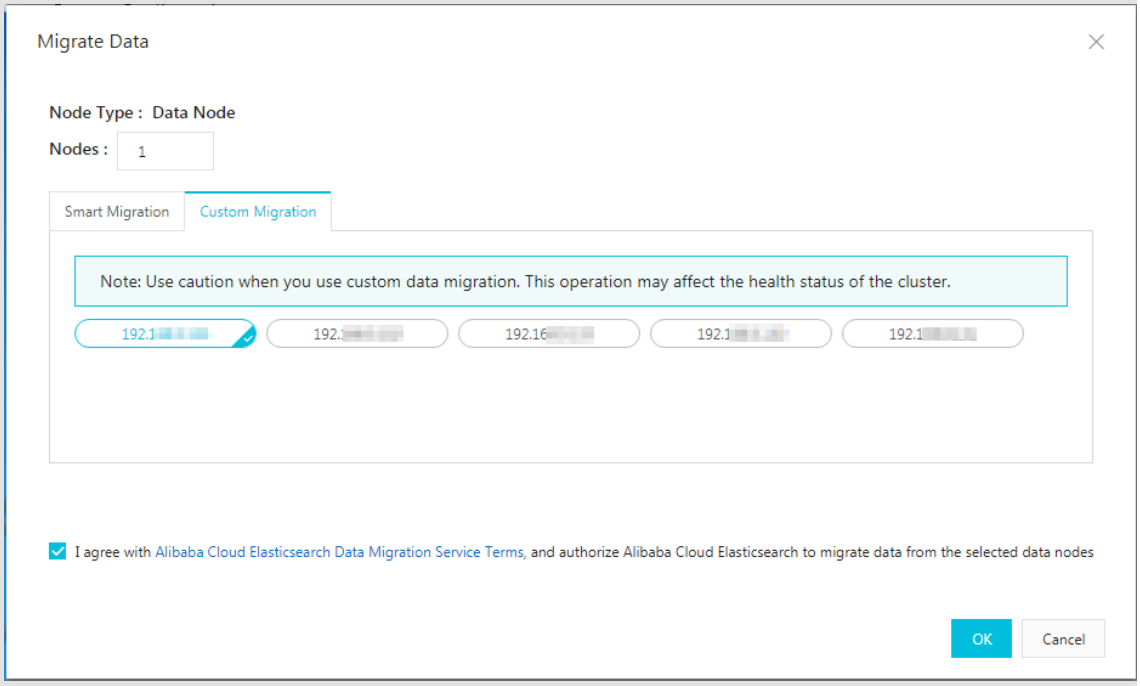

View the downgrade or data migration progress

You can click the tasks list icon in the upper-right corner of the page to view the progress of the downgrade or data migration process.

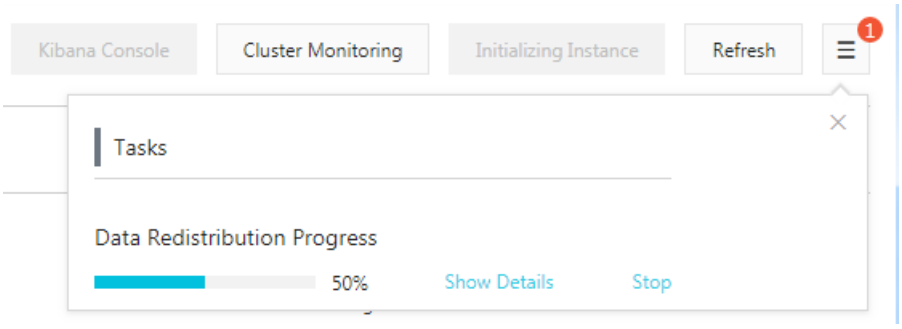

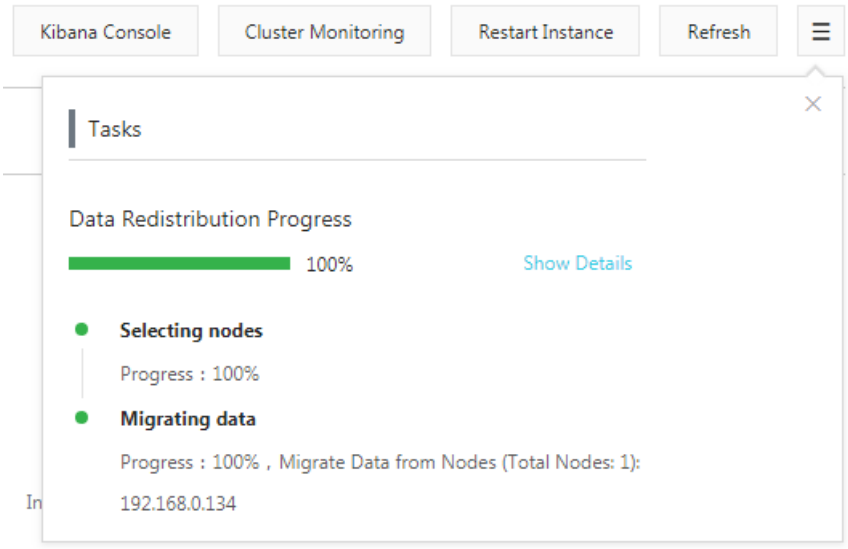

Migration rollback

During the migration process, you can stop the migration task to roll back the migration.

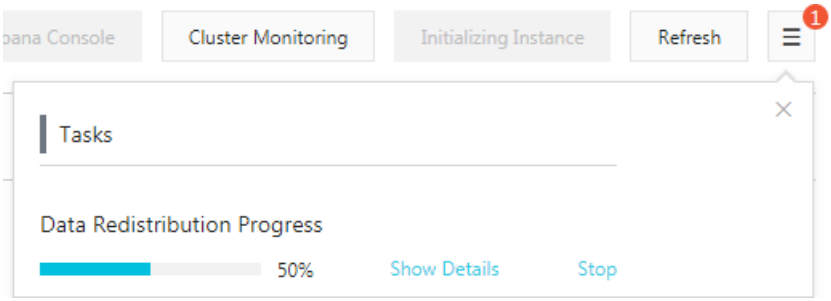

Handle data migration failures

The data migration process is time-consuming. Any cluster status or data changes may result in data migration failures. You can check the Tasks list in the upper-right corner to locate the cause. You can perform the following operations when the data migration task fails or after the task is complete:

1. Query the IP addresses of the data nodes

   You can go to the tasks list or call the Elasticsearch API to query the IP addresses of the data nodes where the data is migrated:

```
// Call the following operation to query the cluster configurat ion
 PUT _cluster / settings

// Sample response
{
  " transient ": {
    " cluster ": {
      " routing ": {
```

```
        " allocation ": {
          " exclude ": {
            " _ip ": " 192 . 168 . ***. ***, 192 . 168 . ***. ***,
 192 . 168 . ***. ***"
          }
        }
      }
    }
  }
}
```

2. Roll back data nodes

   You can call the following operation to roll back data nodes:

```
//  To   roll   back   the   required   data   nodes ,  specify
 the   IP   addresses   of   the   data   nodes   that   you   do
 not   want   to   roll   back   in   the   API   request .
 PUT   _cluster / settings
{
  " transient ": {
    " cluster ": {
      " routing ": {
        " allocation ": {
          " exclude ": {
            " _ip ": " 192 . 168 . ***. ***, 192 . 168 . ***. ***"
          }
        }
      }
    }
  }
}

//  Roll   back   all   data   nodes
 PUT   _cluster / settings
{
  " transient ": {
    " cluster ": {
      " routing ": {
        " allocation ": {
          " exclude ": {
            " _ip ":  null
          }
        }
      }
    }
  }
}
```

3. Verify the rollback result

   You can call the `GET  _cluster / settings` operation to confirm the IP addresses of the data nodes. At the same time, you can check whether shards are reallocated to the data nodes to determine the progress of the rollback task.

   To check the status of the data migration or rollback task, call the `GET  _cat / shards ?  v` operation.

Error messages

Error messages and solutions

During the data migration or downgrade process, the system may prompt the following error messages:

- This operation may cause a shard distribution error or insufficient storage, CPU, or memory resources.

  Cause and solution: after the data migration or downgrade task is complete, the cluster does not have sufficient storage, memory, or CPU resources to store the system data or handle the workload. Call the `GET _cat/indices?v` operation to check whether the number of index replicas in the cluster exceeds the number of data nodes after the cluster is scaled. You also need to check whether the storage, memory, or CPU resources are sufficient to store the existing data or handle the requests.
- The cluster is running tasks or in an error status. Try again later.

  Cause and solution: call the `GET _cluster/health` operation to check the health status of the cluster or go to the Intelligent Maintenance page to verify the cause.
- The nodes in the cluster contain data. You must migrate the data first.
- The number of nodes that you reserve must be greater than two and greater than half of the existing nodes.

  Cause and solution: to ensure the reliability of the cluster, the number of reserved nodes must be greater than 2. To ensure the stability of the cluster, the number of data nodes specified for data migration or downgrading must be no greater than half of the existing data nodes.
- The current Elasticsearch cluster configuration does not support this operation. Check the Elasticsearch cluster configuration first.

  Cause and solution: call the `GET _cluster/settings` operation to query the cluster configuration and check whether the cluster configuration contains settings that forbid data allocation.

auto_expand_replicas

Some users may use the permission management function supported by X-Pack. In former Elasticsearch versions, this function applies the `"index.auto_expan`

d_replicas " : " 0 - all " setting to indexes .security and .security-6 by default. This causes data migration or downgrading failures. We recommend that you modify the auto_expand_replicas option as follows:

```
//  Query   the   index   configurat  ion
 GET  . security / _settings

//  Returned   results
{
  ". security - 6 " : {
    " settings " : {
      " index " : {
        " number_of_  shards " : " 1 ",
        " auto_expan  d_replicas " : " 0 - all ",
        " provided_n  ame " : ". security - 6 ",
        " format " : " 6 ",
        " creation_d  ate " : " 1555142250  367 ",
        " priority " : " 1000 ",
        " number_of_  replicas " : " 9 ",
        " uuid " : " 9t2hotc7S5  OpPuKEIJ ****",
        " version " : {
          " created " : " 6070099 "
        }
      }
    }
  }
}

//  Use   one   of   the   following   methods   to   modify   the
 auto_expan  d_replicas   setting
 PUT  . security / _settings
{
  " index " : {
    " auto_expan  d_replicas " : " 0 - 1 "
  }
}

 PUT  . security / _settings
{
  " index " : {
    " auto_expan  d_replicas " : " false ",
    " number_of_  replicas " :  1 ,
  }
}

//  Set   the   number   of   replicas   based   on   the   actual
 needs .  The   number   of   replicas   must   be   greater   than
 1   and   less   than   or   equal   to   the   number   of   the
 available   data   nodes .
```

2.3 Cluster upgrade

This topic describes the procedure, guidelines, and restrictions of upgrading an Alibaba Cloud Elasticsearch instance.

Alibaba Cloud Elasticsearch allows you to upgrade the instance specification, number of nodes, dedicated master node specification, number of client nodes, client node

specification, number of warm nodes, warm node specification, warm node storage space, and storage space per data node of an Elasticsearch instance.

> **Note:**
> You may not be able to upgrade some of the cluster properties due to certain restrictions. For more information, see Configuration upgrade.

Log on to the Alibaba Cloud Elasticsearch console, select Instance ID > Basic Information, and then click Upgrade to navigate to the Update page.

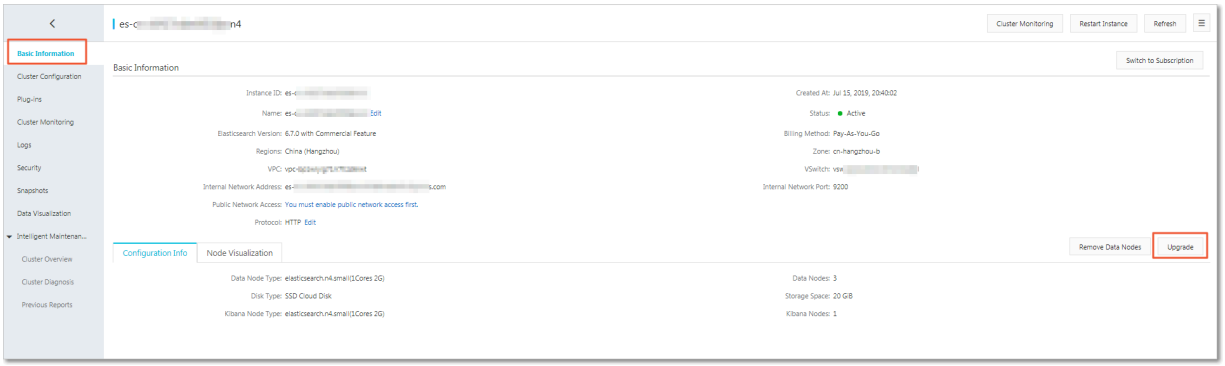

The Update page includes the Current Config and Configuration Upgrade information. For more information, see Current configuration and Configuration upgrade.

Current configuration

The Current Config section shows the configuration of the current Alibaba Cloud Elasticsearch instance. You can reference the information when you upgrade the instance.

Precautions

Before you upgrade an Elasticsearch instance, pay close attention to the following precautions:

- If you need to upgrade the instance due to business requirements, make an assessment before you upgrade the cluster.
- For each upgrade operation, you can only change one of the upgradable cluster properties.
- Typically, Elasticsearch needs to restart your Elasticsearch instance for the upgrade to take effect. For an Elasticsearch instance with dedicated master nodes, if you change the number of nodes, the instance will not be restarted.

- If the status of your Elasticsearch instance is unhealthy (showing a yellow or red flag), then you must select Force Update to upgrade the instance. Force update may affect your businesses.
- You cannot change the disk type of nodes by upgrading the instance. You can only change the storage space per node.
- Alibaba Cloud Elasticsearch allows you to upgrade the specification of the Kibana node. Fees are charged for upgrading the Kibana node.
- Alibaba Cloud Elasticsearch subscription instances currently do not support downgrading. For example, you cannot remove nodes from clusters, scale in the disk space, or downgrade the node specifications.
- You can downgrade Alibaba Cloud Elasticsearch pay-as-you-go instances by scaling in the number of data nodes. The number of data nodes that you can scale in is restricted. Currently, you cannot perform other downgrade operations. For example, you cannot scale in the disk space or downgrade the node specification.
- After you change the configuration of the instance, you can check the amount of your order on the Update page.
- After you submit the order, your Elasticsearch instance will be billed based on the new configuration.

Configuration upgrade

> **Notice:**
> Before you upgrade the configuration of an Elasticsearch instance, make sure that you have read the precautions in Precautions.

You can follow the instructions on the configuration upgrade page to change the configuration of the instance to meet your business requirements. For more information about the parameters, see #unique_19.

Some of the parameters are described as follows:

- Specification family and instance type

  The Specification Family cannot be changed. If the Specification Family is set to a local disk type, then the Instance Type cannot be changed.

- Dedicated master nodes

  On the Update page, click Yes on the right side of Dedicated Master Node to purchase dedicated master nodes. You can upgrade the specification of the purchased dedicated master nodes. By default, three dedicated master nodes are purchased. Each dedicated master node has 2 cores, 8 GB of memory, and a cloud disk of 20 GiB. After you upgrade the dedicated master nodes, the Elasticsearch instance will be billed based on the new configuration.

   Note:
  If you have purchased 1-core 2 GB dedicated master nodes, then you can repurchase dedicated master nodes of higher specifications on the Update page. The Elasticsearch instance will be billed based on the new configuration. If your dedicated master nodes are free nodes provided by Elasticsearch, then after you upgrade these nodes, we will start charging these nodes.

- Client nodes

  On the Update page, click Yes on the right side of Client Node to purchase client nodes. You can upgrade the specification of the purchased client nodes. By default, two client nodes are purchased. Each client node has 2 cores, 8 GB of memory, and a cloud disk of 20 GiB. After you upgrade the client nodes, the Elasticsearch instance will be billed based on the new configuration.

- Warm nodes

  On the Update page, click Yes on the right side of Warm Node to purchase warm nodes. You can upgrade the specification of the purchased warm nodes. By default, two warm nodes are purchased. Each warm node has 2 cores, 8 GB of memory, and a cloud disk of 500 GiB. After you upgrade the warm nodes, the Elasticsearch instance will be billed based on the new configuration.

- Kibana node

  On the Update page, click Yes on the right side of Kibana Node to purchase a Kibana node. You can upgrade the specification of the purchased Kibana node. By default, the Kibana node has two cores and 4 GB of memory.

  > Notice:
  > After you purchase an Alibaba Cloud Elasticsearch instance, Elasticsearch provides you a free Kibana node with 1 core and 2 GB of memory. After you upgrade the Kibana node, the Elasticsearch instance will be billed based on the new configuration.

- Force update

  If the status of your Elasticsearch instance is unhealthy (showing a red or yellow flag), then your businesses have been severely affected. You must upgrade the instance immediately. You can select Force Update to ignore the status of the Elasticsearch instance and forcibly upgrade the instance. The upgrade process only takes a short period of time.

  > Notice:
  > - The Elasticsearch instance needs to restart to complete the force update process.
  > - During the force update process, the services running on the Elasticsearch instance may become unstable.
  > - If you do not select Force Update, the restart method is used to upgrade the instance by default. For more information, see #unique_20/unique_20_Connect_42_section_p5n_ccm_zgb.
  > - If the status of your Alibaba Cloud Elasticsearch instance is not healthy (a red or yellow flag), then the system will automatically select Force Update for you. Elasticsearch will not use the restart method to upgrade the instance.

- Node storage

  The storage space of nodes is measured in GiB. A standard SSD disk can provide up to 2,048 GiB (2 TiB) of storage space.

  You can scale out an ultra disk to up 2 TiB. When you purchase an ultra disk, you can set the storage space to up to 5,120 GiB (5 TiB). Ultra disks larger than 2,048 GiB include 2,560 GiB, 3,072 GiB, 3,584 GiB, 4,096 GiB, 4,608 GiB, and 5,120 GiB.

3 Elasticsearch cluster configuration

3.1 Elasticsearch cluster configuration

Word splitting

This feature uses the synonym dictionary. New indexes will use the updated synonym dictionary. For more information, see #unique_23.

Note:

- After you upload and submit a synonym dictionary file, the Alibaba Cloud Elasticsearch instance will not restart immediately. It takes some time for the new configuration to take effect.
- If an index that is created before the uploaded synonym dictionary file takes effect needs to use synonyms, you must recreate the indexes and configure synonyms.

Write one synonym expression in each row and save the code as a `UTF - 8` encoded `. txt` file. Examples:

```
corn , maize => maize , corn
begin , start => start , begin
```

Configuration procedure:

1. Upload and save a synonym dictionary file in the Alibaba Cloud Elasticsearch console. Make sure that the uploaded file takes effect.
2. When you create an index and configure the `settings`, you need to specify the `" synonyms_p ath ": " analysis / your_dict_ name . txt "` path. Add a `mapping` for this index to configure synonyms for the specified field.
3. Confirm the synonyms and upload a file for testing.

YML configurations

The YML Configurations page displays the settings of the current Alibaba Cloud Elasticsearch instance.

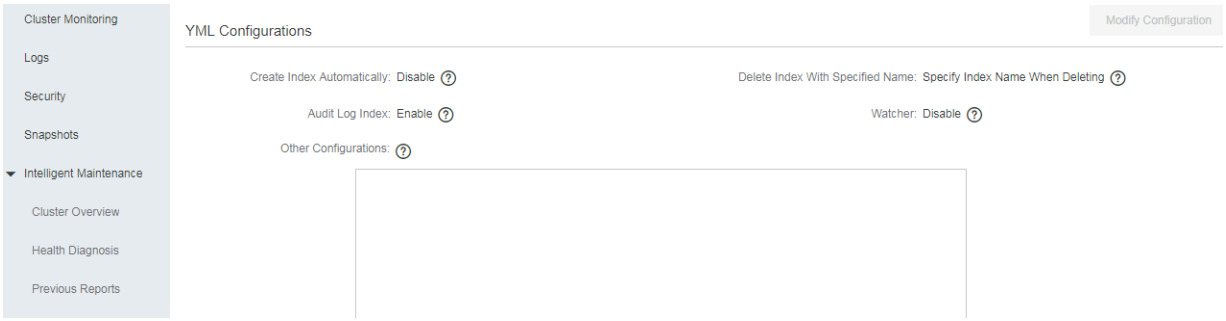

Modify YML configurations

After you modify the YML Configurations, you must restart the Alibaba Cloud Elasticsearch instance for the new configuration to take effect.

Note:
After you modify the YML Configurations, select This operation requires a restart of the instance. Exercise with caution. at the bottom of the page and click OK. The Alibaba Cloud Elasticsearch instance automatically restarts.

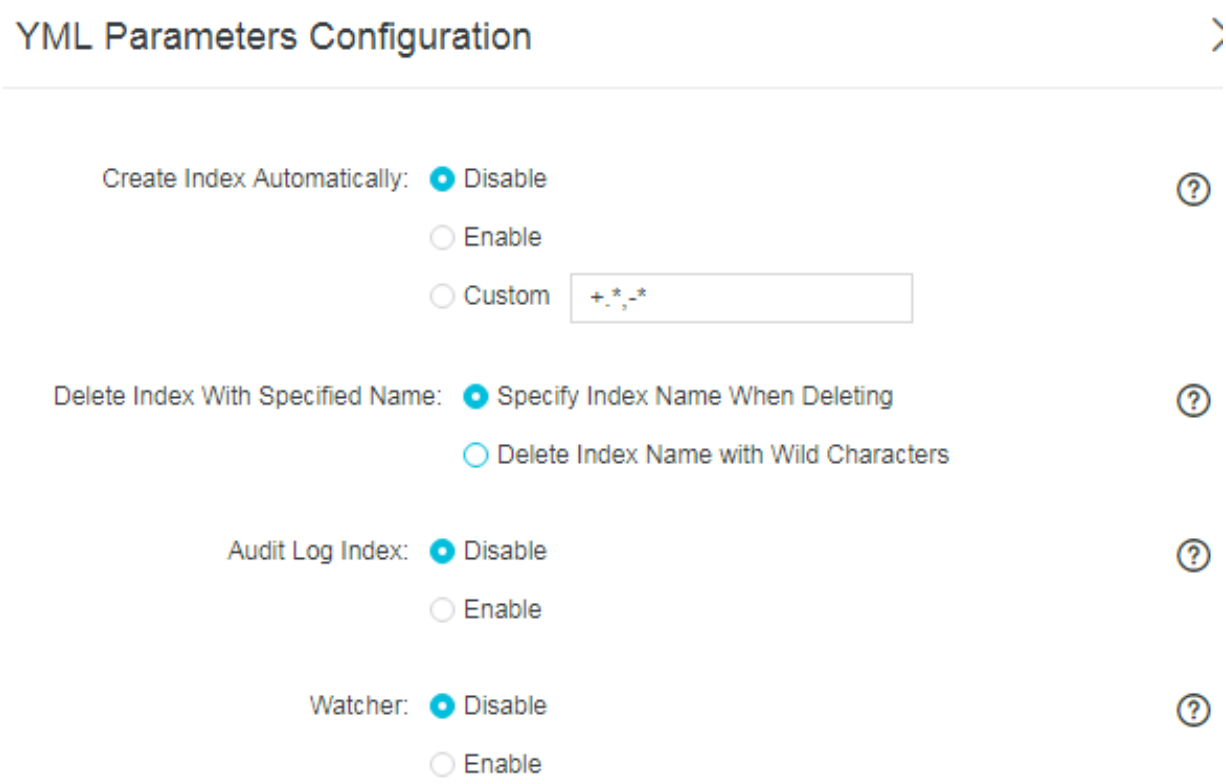

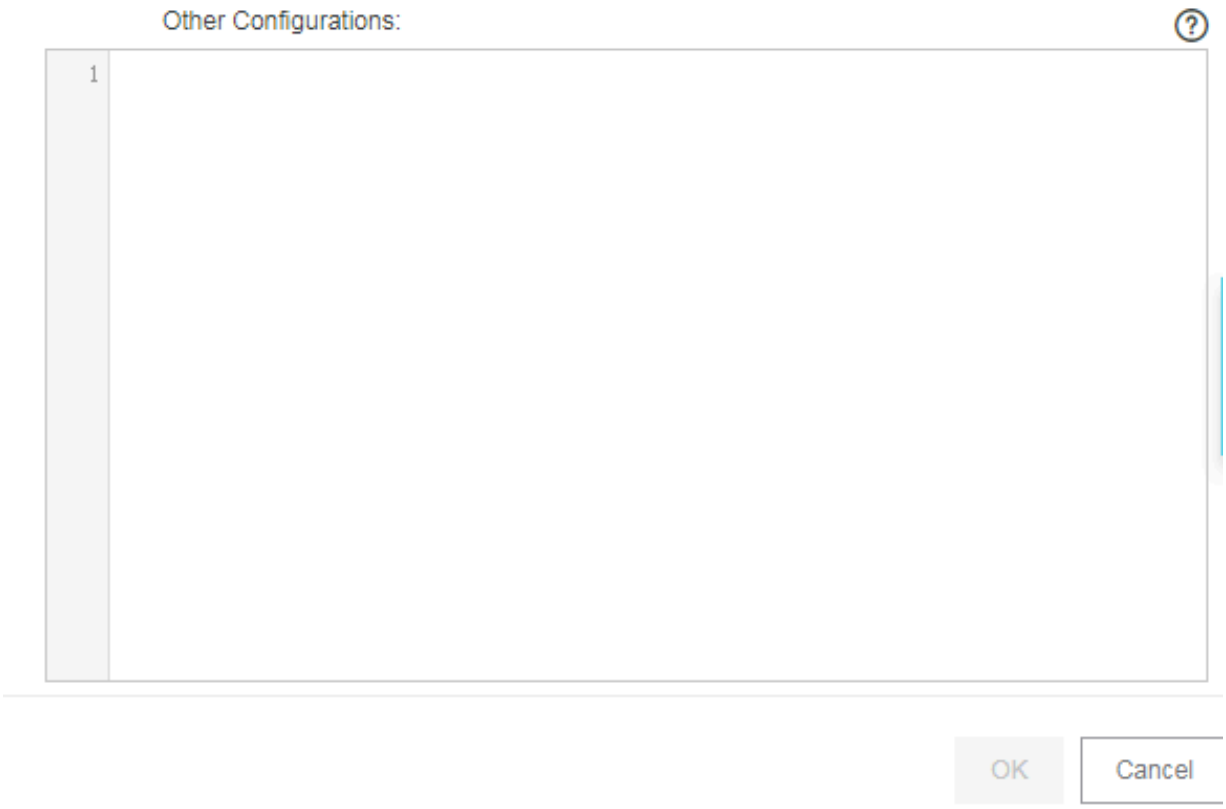

- Create Index Automatically: if you enable this feature, it allows the system to automatically create new indexes if a new file is uploaded to the Alibaba Cloud Elasticsearch instance and no indexes have been created on the file. We recommend that you disable this feature. Indexes created by this feature may not meet your requirements.
- Delete Index With Specified Name: this feature indicates whether you are required to specify the name of the index that you need to delete. If you select Delete Index Name with Wild Characters, you can delete multiple indexes by using a wildcard character. Indexes that are deleted cannot be restored. Proceed with caution.
- Audit Log Index: if you enable this feature, index logs are created and stored when you create, delete, modify, or view an Alibaba Cloud Elasticsearch instance. These logs consume disk space and affect the performance. We recommend that you disable this feature. Proceed with caution.
- Watcher: if you enable this feature, it allows you to use the X-Pack Watcher feature. Make sure that you regularly clear the `. watcher - history *` index. This index consumes large amounts of disk space.

- Other Configurations: the following parameters are supported. For more information, see #unique_24.

  Note:
  Excluding the parameters that have an Alibaba Cloud Elasticsearch version specified, the remaining parameters can only be applied to Elasticsearch V5.5.3 and V6.3.2.

  - http.cors.enabled
  - http.cors.allow-origin
  - http.cors.max-age
  - http.cors.allow-methods
  - http.cors.allow-headers
  - http.cors.allow-credentials
  - reindex.remote.whitelist
  - action.auto_create_index
  - action.destructive_requires_name
  - thread_pool.bulk.queue_size (Elasticsearch V5.5.3 with X-Pack)
  - thread_pool.write.queue_size (Elasticsearch V6.3.2 with X-Pack)
  - thread_pool.search.queue_size

3.2 Configure synonyms

Description

Note:

- After you upload a synonym dictionary file to an Alibaba Cloud Elasticsearch instance, you do not need to restart the nodes in the instance. The system will update the synonym dictionary file to all nodes. Depending on the number of nodes, this process may be time-consuming.
- For example, index ‘index-aliyun’ is using the synonym dictionary file ‘aliyun.txt’. You have uploaded a new synonym dictionary file to overwrite the existing dictionary file. However, index ‘index-aliyun’ cannot automatically load the updated dictionary file. If you want the index to load the updated dictionary file, disable the index and then re-enable the index. We recommend that you

rebuild the index after you update the dictionary file as a best practice. Otherwise , this may cause an issue that only the newly created data is using the updated dictionary file.

You can use a filter to configure synonyms. The sample code is as follows:

```
 PUT  / test_index
{
    " settings ": {
        " index " : {
            " analysis " : {
                " analyzer " : {
                    " synonym " : {
                        " tokenizer " : " whitespace ",

                        " filter " : [" synonym "]
                        }
                    },
                    " filter " : {
                        " synonym " : {
                            " type " : " synonym ",

                             " synonyms_p  ath " : " analysis /
 synonym . txt ",
                             " tokenizer " : " whitespace "

                        }
                    }
                }
            }
        }
}
```

- `filter` : configure a `synonym` token filter that contains the path `analysis / synonym . txt` . This path is relative to the location of config.
- `tokenizer` : the tokenizer that tokenizes synonyms. It is set to `whitespace` by default. Additional settings:
  - `ignore_cas  e` : the default value is false.
  - `expand` : the default value is true.

Two synonym formats are supported: Solr and WordNet.

- **Solr synonyms**

  The following is a sample format of the file:

```
#  Blank   lines   and   lines   starting   with   pound   are
 comments .
#  Explicit   mappings   match   any   token   sequence   on   the
 LHS   of  "=>"
#  and   replace   with   all   alternativ  es   on   the   RHS .
 These   types   of   mappings
#  ignore   the   expand   parameter   in   the   schema .
```

```
#  Examples :
 i - pod ,  i   pod  =>  ipod ,
 sea   biscuit ,  sea   biscit  =>  seabiscuit
#  Equivalent   synonyms   may   be   separated   with   commas
 and   give
#  no   explicit   mapping .   In   this   case   the   mapping
 behavior   will
#  be   taken   from   the   expand   parameter   in   the   schema
 .   This   allows
#  the   same   synonym   file   to   be   used   in   different
 synonym   handling   strategies .
#  Examples :
 ipod ,  i - pod ,  i   pod
 foozball  ,  foosball
 universe  ,  cosmos
 lol ,  laughing   out   loud
#  If   expand == true , " ipod ,  i - pod ,  i   pod "  is
 equivalent
#  to   the   explicit   mapping :
 ipod ,  i - pod ,  i   pod  =>  ipod ,  i - pod ,  i   pod
#  If   expand == false , " ipod ,  i - pod ,  i   pod "  is
 equivalent
#  to   the   explicit   mapping :
 ipod ,  i - pod ,  i   pod  =>  ipod
#  Multiple   synonym   mapping   entries   are   merged .
 foo  =>  foo   bar
 foo  =>  baz
#  is   equivalent   to
 foo  =>  foo   bar ,  baz
```

You can also directly define synonyms for the token filter in the configuration file. You must use `synonyms` instead of `synonyms_p ath` . Example:

```
 PUT  / test_index
{
    " settings ": {
        " index " : {
            " analysis " : {
                " filter " : {
                    " synonym " : {
                        " type " : " synonym ",
                        " synonyms " : [
                            " i - pod ,  i   pod  =>  ipod ",
                            " begin ,  start "
                        ]
                    }
                }
            }
        }
    }
}
```

We recommend that you use `synonyms_p ath` to define large synonym sets in the file. Using `synonyms` to define large synonym sets will increase the size of your cluster.

- WordNet synonyms

  Synonyms based on the WordNet format can be declared by using the following format:

```
 PUT  / test_index
{
    " settings ": {
        " index " : {
            " analysis " : {
                " filter " : {
                    " synonym " : {
                        " type " : " synonym ",
                        " format " : " wordnet ",
                        " synonyms " : [
                            " s ( 100000001 , 1 ,' abstain ', v , 1
 , 0 ).",
                            " s ( 100000001 , 2 ,' refrain ', v , 1
 , 0 ).",
                            " s ( 100000001 , 3 ,' desist ', v , 1 ,
 0 )."
                        ]
                    }
                }
            }
        }
    }
}
```

  You can also use `synonyms_p ath` to define WordNet synonyms in a file.

Example 1:

Upload a synonym dictionary file

1. Log on to the Alibaba Cloud Elasticsearch console.
2. Click Create in the upper-left corner to create an Alibaba Cloud Elasticsearch instance.
3. Click the instance to go to the configuration page.
4. In the left-side navigation pane, select Cluster Configuration, and then click Synonym Dictionary Configuration.

5. Click Upload, select the synonym dictionary file that you want to upload, and click Save . In this example, the TXT file that is generated in the format described in the preceding sections is uploaded.

   After the Alibaba Cloud Elasticsearch instance is activated and its status changes to Active, you can then use the synonym dictionary. In this example, file `aliyun_syn onyms . txt` is uploaded for testing. The file contains: `begin , start`

Configure and test the synonym dictionary

1. Click Kinana Console in the upper-right corner to go to the Kibana console.
2. In the left-side navigation pane, click Dev Tool.
3. Run the following command in the Console to create indexes:

```
 PUT   aliyun - index - test
{
" index ": {
 " analysis ": {
   " analyzer ": {
     " by_smart ": {
       " type ": " custom ",
       " tokenizer ": " ik_smart ",
       " filter ": [" by_tfr ," by_sfr "],
       " char_filte  r ": [" by_cfr "]
     },
     " by_max_wor  d ": {
       " type ": " custom ",
       " tokenizer ": " ik_max_wor  d ",
       " filter ": [" by_tfr ," by_sfr "],
       " char_filte  r ": [" by_cfr "]
     }
   },
   " filter ": {
     " by_tfr ": {
       " type ": " stop ",
       " stopwords ": [" "]
     },
     " by_sfr ": {
       " type ": " synonym ",
       " synonyms_p  ath ": " analysis / aliyun_syn  onyms . txt "
     }
   },
   " char_filte  r ": {
     " by_cfr ": {
       " type ": " mapping ",
       " mappings ": ["| => |"]
     }
   }
 }
}
```

```
}
```

4. Run the following command to configure the title field:

```
 PUT   aliyun - index - test / _mapping / doc
{
" properties ": {
 " title ": {
   " type ": " text ",
   " index ": " analyzed ",
   " analyzer ": " by_max_wor  d ",
   " search_ana  lyzer ": " by_smart "
 }
}
}
```

5. Run the following command to verify the synonyms:

```
 GET   aliyun - index - test / _analyze
{
" analyzer ": " by_smart ",
" text ":" begin "
}
```

The following results are returned if the configuration takes effect:

```
{
" tokens ": [
 {
   " token ": " begin ",
   " start_offs  et ":  0 ,
   " end_offset ":  5 ,
   " type ": " ENGLISH ",
   " position ":  0
 },
 {
   " token ": " start ",
   " start_offs  et ":  0 ,
   " end_offset ":  5 ,
   " type ": " SYNONYM ",
   " position ":  0
 }
]
}
```

6. Run the following command to add data for further testing:

```
 PUT   aliyun - index - test / doc / 1
{
" title ": " Shall   I   begin ?"
}
```

```
 PUT   aliyun - index - test / doc / 2
{
" title ": " I   start   work   at   nine ."
}
```

7. Run the following command to perform a query test:

```
 GET   aliyun - index - test / _search
```

```
{
 " query " : { " match " : { " title " : " begin " }},
 " highlight " : {
     " pre_tags " : ["< red >", "< bule >"],
     " post_tags " : ["</ red >", "</ bule >"],
     " fields " : {
         " title " : {}
     }
 }
}
```

If the query is successful, the following results are returned:

```
{
" took ":  11 ,
" timed_out ":  false ,
" _shards ": {
 " total ":  5 ,
 " successful ":  5 ,
 " failed ":  0 ,
},
" hits ": {
 " total ":  2 ,
 " max_score ":  0 . 41048482 ,
 " hits ": [
   {
     " _index ": " aliyun - index - test ",
     " _type ": " doc ",
     " _id ": " 2 ",
     " _score ":  0 . 41048482 ,
     " _source ": {
       " title ": " I   start   work   at   nine ."
     },
     " highlight ": {
       " title ": [
         " I  < red > start </ red >  work   at   nine ."
       ]
     }
   },
   {
     " _index ": " aliyun - index - test ",
     " _type ": " doc ",
     " _id ": " 1 ",
     " _score ":  0 . 39556286 ,
     " _source ": {
       " title ": " Shall   I   begin ?"
     },
     " highlight ": {
       " title ": [
         " Shall   I  < red > begin </ red >?"
       ]
     }
   }
 ]
}
}
```

Example 2

Follow these steps to directly import the synonyms and use the IK analyzer to filter the synonyms:

1. Configure synonym filter `my_synonym _filter` and a synonym dictionary.
2. Configure analyzer `my_synonym s`, and use IK analyzer `ik_smart` to split words.

   The IK analyzer `ik_smart` splits the words and then changes all letters to lowercase.

```
 PUT  / my_index
{
 " settings ": {
     " analysis ": {
         " analyzer ": {
             " my_synonym  s ": {
                 " filter ": [
                     " lowercase ",
                     " my_synonym  _filter "
                 ],
                 " tokenizer ": " ik_smart "
             }
         },
         " filter ": {
             " my_synonym  _filter ": {
                 " synonyms ": [
                     " begin , start "
                 ],
                 " type ": " synonym "
             }
         }
     }
 }
}
```

3. Run the following command to configure the title field:

```
 PUT  / my_index / _mapping / doc
{
" properties ": {
 " title ": {
   " type ": " text ",
   " index ": " analyzed ",
   " analyzer ": " my_synonym  s "
 }
}
}
```

4. Run the following command to verify the synonyms:

```
 GET  / my_index / _analyze
{
 " analyzer ":" my_synonym  s ",
 " text ":" Shall   I   begin ?"
}
```

   If the synonyms are verified, the following results are returned:

```
{
" tokens ": [
```

```
 {
   " token ": " shall ",
   " start_offs  et ":  0 ,
   " end_offset ":  5 ,
   " type ": " ENGLISH ",
   " position ":  0
 },
 {
   " token ": " i ",
   " start_offs  et ":  6 ,
   " end_offset ":  7 ,
   " type ": " ENGLISH ",
   " position ":  1
 },
 {
   " token ": " begin ",
   " start_offs  et ":  8 ,
   " end_offset ":  13 ,
   " type ": " ENGLISH ",
   " position ":  2
 },
 {
   " token ": " start ",
   " start_offs  et ":  8 ,
   " end_offset ":  13 ,
   " type ": " SYNONYM ",
   " position ":  2
 }
]
}
```

5. Run the following command to add data for further testing:

```
 PUT  / my_index / doc / 1
{
" title ": " Shall   I   begin ?"
}
```

```
 PUT  / my_index / doc / 2
{
" title ": " I   start   work   at   nine ."
}
```

6. Run the following command to perform a query test:

```
 GET  / my_index / _search
{
" query " : { " match " : { " title " : " begin " }},
" highlight " : {
  " pre_tags " : ["< red >", "< bule >"],
  " post_tags " : ["</ red >", "</ bule >"],
  " fields " : {
     " title " : {}
  }
}
}
```

7. If the query is successful, the following results are returned:

```
{
" took ":  11 ,
```

```
" timed_out ":  false ,
" _shards ": {
 " total ":  5 ,
 " successful ":  5 ,
 " failed ":  0 ,
},
" hits ": {
 " total ":  2 ,
 " max_score ":  0 . 41913947 ,
 " hits ": [
   {
     " _index ": " my_index ",
     " _type ": " doc ",
     " _id ": " 2 ",
     " _score ":  0 . 41913947 ,
     " _source ": {
       " title ": " I   start   work   at   nine ."
     },
     " highlight ": {
       " title ": [
         " I  < red > start </ red >  work   at   nine ."
       ]
     }
   },
   {
     " _index ": " my_index ",
     " _type ": " doc ",
     " _id ": " 1 ",
     " _score ":  0 . 39556286 ,
     " _source ": {
       " title ": " Shall   I   begin ?"
     },
     " highlight ": {
       " title ": [
         " Shall   I  < red > begin </ red >?"
       ]
     }
   }
 ]
}
}
```

3.3 YML configuration

Customize CORS requests

For more configurations, visit the Elasticsearch official website and view the HTTP information.

Configuration information

- Configurations in the table below are custom HTTP-based configurations provided by Alibaba Cloud Elasticsearch.
- For the following configurations, only static configuration is supported. Dynamic configuration is not supported. Note that for the following configurations to take effect, you must add the configurations to the `elasticsea rch . yml` file.

- Cluster network settings are used for the following configurations. (Network settings)

Configuration item	Description
`http . cors . enabled`	A CORS (Cross-Origin Resource Sharing) configuration item, which can be used to enable or disable CORS resource accesses. In other words, this setting is used to determine whether to allow Elasticsearch to receive requests sent by browsers to access resources in different domains. If the parameter is set to `true` , Elasticsearch can process `OPTIONS` CORS requests. If the domain information in the sent request is already declared in `http . cors . allow - origin` , Elasticsearch adds `Access - Control - Allow - Origin` in the header to respond to the CORS request. If the parameter is set to `false` (which is the default value), Elasticsearch ignores the domain information in the request header, not adding the `Access - Control - Allow - Origin` to the header, disabling CORS access. If the client neither supports `pre - flight` requests that add the domain information header, nor checks `Access - Control - Allow - Origin` in the header of the packet returned from the server, then the secured CORS access will be affected. If Elasticsearch disables CORS access, then the client can only check whether a response is returned by sending the `OPTIONS` request.

Configuration item	Description
`http.cors.allow-origin`	A CORS resource configuration item, which can be used to specify requests from which domains are accepted. The parameter is left blank, by default, with no domain is allowed. If `/` is added before the parameter value, then the configuration is identified as a regular expression, which means that `HTTP` and `HTTPS` domain requests that follow the regular expression are supported. For example`/Https?://Localhost(:[0-9]+)?/` means requests follow the regular expression can be responded to. `*` means that a configuration is valid and can be identified as enabling the cluster to support CORS requests from any domain, resulting in security risks to the Elasticsearch cluster.
`http.cors.max-age`	The browser can send an `OPTIONS` request to get the CORS configuration. `max-age` can be used to set how long the browser can retain the output result cache. The default value is `1728000` seconds (20 days).
`http.cors.allow-methods`	A request method configuration item. The optional values are `OPTIONS`, `HEAD`, `GET`, `POST`, `PUT`, and `DELETE`.
`http.cors.allow-headers`	A request header configuration item. The default value is `X-Requested-With, Content-Type, Content-Length`.
`http.cors.allow-credentials`	A credential configuration item, which is used to specify whether to return `Access-Control-Allow-Credentials` in the response header. If the parameter is set to `true`, Access-Control-Allow-Credentials is returned. The default value is `false`.

An example of custom cross-origin access configuration is as follows:

```
http . cors . enabled :  true
http . cors . allow - origin : "*"
http . cors . allow - headers : " X - Requested - With ,  Content -
Type ,  Content - Length ,  Authorizat  ion "
```

Customize remote re-indexing (whitelist)

The re-indexing component allows you to reconstruct the data index on the target remote Elasticsearch cluster. This function can work for all of the remote Elasticsea rch versions available, allowing you to index the data of earlier versions to the current version.

```
 POST   _   reindex
{
  " source ": {
    " remote ": {
      " Host ": " http : //  otherhost :  9200  ",
      " username ": " username ",
      " password ": " password ",
    },
    " index ": " source ",
    " query ": {
      " match ": {
        " test ": " data "
      }
    }
  },
  " dest ": {
    " index ": " test - 1 ",
  }
}
```

- `host` must contain the protocol supported, domain name, port, for example, `Https : //  otherhost :  9200` .
- `username` and `password` are optional. If the remote Elasticsearch server requires Basic Authorization, enter the username and password in the request. When use `Basic  Authorizat  ion` , also use the `https` protocol, otherwise the password will be transmitted as a text.
- The remote host address must be declared in `elasticsea  rch . yml` by using the `reindex . remote . whitelist` attribute for the API to be called remotely. The combination of host and port is allowed. The combination of `host` and `port` is allowed. However, note that multiple host configurations must be separated by commas (,), for example,

  ```
  otherhost :  9200 ,  another :  9200 , 127 . 0  . 10 . **:  9200 ,
  ```

```
            localhost :**
```

). The whitelist does not identify the protocol and only uses the host and port information for the security policy configuration.

- If the host address is already listed in the whitelist, the `query` request will not be verified or modified. Rather, the request will be directly sent to the remote server.

 Note:

- Indexing data from a remote cluster is not supportedManual SlicingOrAutomatic Slicing. For more information, see Manual slicing or Automatic slicing.

Multiple indexes settings

The remote service uses a stack to cache indexed data. The default maximum size is `100 MB`. If the remote index contains a large document, set the size of batch settings to a small value.

In the example below, the size of multiple index settings is 10, which is the minimum value:

```
 POST  _  reindex
{
  " source ": {
    " remote ": {
      " host ": " http :// otherhost : 9200 "
    },
    " index ": " source ",
    " size ":  10 ,
    " query ": {
      " match ": {
        " test ": " data "
      }
    }
  },
  " dest ": {
    " index ": " test - 1 ",
  }
}
```

Timeout period

- Use `socket_tim eout` to set the read timeout period of `socket`. The default value is `30s`.
- Use `connect_ti meout` to set the connection timeout period. The default value is `1s`.

In the example below, the read timeout period of `socket` is one minute, and the connection timeout period is 10 seconds.

```
 POST   _   reindex
{
  " source ": {
    " remote ": {
      " host ": " http :// otherhost : 9200 ",
      " socket_tim  eout ": " 1m ",
      " connect_ti  meout ": " 10s "
    },
    " index ": " source ",
    " query ": {
      " match ": {
        " test ": " data "
      }
    }
  },
  " dest ": {
    " index ": " test - 1 ",
  }
}
```

Customize the access log

Enable auditing

The index auditing configuration is as follows.

```
xpack . security . audit . index . bulk_size :  5000
xpack . security . audit . index . events . emit_reque  st_body :
false
xpack . security . audit . index . events . exclude :  run_as_den
ied , anonymous_  access_den  ied , realm_auth  entication  _failed ,
access_den  ied , connection  _denied
xpack . security . audit . index . events . include :  authentica
tion_faile  d , access_gra  nted , tampered_r  equest , connection
_granted , run_as_gra  nted
xpack . security . audit . index . flush_inte  rval :  180s
xpack . security . audit . index . rollover :  hourly
xpack . security . audit . index . settings . index . number_of_
replicas :  1
xpack . security . audit . index . settings . index . number_of_
shards :  10
```

Index auditing output

Alibaba Cloud Elasticsearch instances do not support displaying request-related log files. Therefore, to view information about the Elasticsearch instance requests, such as the access_log, you must log in to the Elasticsearch console and enable the access log index function.

After this function is enabled, the access log is output to indexes on the Elasticsearch instance. The name of indexes starts with `. security_a  udit_log -*`.

Audit log indexing configuration

Note:

- Filtering is not supported during audits because sensitive data may be audited in plain text when the `request body` is included in audit events.
- Audit log indexing occupies Alibaba Cloud Elasticsearch instance storage space. You must manually clear old audit log indexes because no policy is available for clearing expired indexes.

Feature	Default value	[DO NOT TRANSLATE]
`xpack . security . audit . index . bulk_size`	`1 , 000`	Indicates how many audit events are batched into a single write file.
`xpack . security . audit . index . flush_inte rval`	`1 s`	Indicates how often buffered events are flushed to the index.
`xpack . security . audit . index . rollover`	`daily`	Indicates how often to roll over to a new index. Options include `hourly` , `daily` , `weekly` , or `monthly` .
`Xpack . security . audit . index . events . include`	`anonymous_ access_den ied` , `authentica tion_faile d` , `realm_auth entication _failed` , `access_gra nted` , `access_den ied` , `tampered_r equest` , `connection _granted` , `connection _denied` , `run_as_gra nted` , `run_as_den ied`	Specifies the audit events to be indexed. For more information about audit event types, see Audit event types.

Feature	Default value	[DO NOT TRANSLATE]
`xpack . security . audit . index . events . exclude`		Excludes the specified auditing events from indexing.
`xpack . security . audit . index . events . emit_request_body`	`false`	Indicates whether to include the request body in REST requests in certain event types, such as `authentication_failed` .

Audit indexing settings

The configuration item `xpack . security . audit . index . settings` in the `elasticsearch . yml` file specifies the settings for the indexes in which the events are stored.

The following example sets both the number of shards and the number of replicas to `1` for the audit indexes.

```
xpack . security . audit . index . settings :
  index :
    number_of_shards : 1
    number_of_replicas : 1
```

Note:

You can pass custom settings to xpack.security.audit.index.settings when enabling audit indexing. Once you apply the change to the Elasticsearch instance, audit indexes will be available on the Elasticsearch instance. Otherwise, the elasticsearch instance audit log is set to the default `Number_of_shards : 5` , and `Number_of_replicas : 1` .

Remote audit log indexing settings

Indexing settings for remote audit logs are currently unavailable.

Customize thread pool queue size

You can set `Thread_pool . bulk . queue_size` , `Thread_pool . write . queue_size` , and `Thread_pool . search . queue_size` to customize the queue size of the write and search thread pools, respectively..

In the following example, both the write and search queue size are set to `500` .

Note:

The following parameters are not specifically identified for an ES version and by default are compatible with ES version 5.5.3 and 6.3.2.

```
thread_poo l . bulk . queue_size : 500 ( Only applicable to
the Elasticsea rch 5 . 5 . 3 with X - Pack version )
thread_poo l . write . queue_size : 500 ( Only applicable to
the Elasticsea rch 6 . 3 . 2 with X - Pack version )
thread_poo l . search . queue_size : 500
```

Parameter optimization

Configuration Item	Description
Index. codec	The ES data compression algorithm defaults to LZ4. Usually, by setting LZ4 to best_compression in a warm or cold cluster using a high-speed cloud disk, a higher compression ratio DEFLATE algorithm can be used. After the algorithm is changed, segment merges will use the newest version of the algorithm. Note that using best_compr ession will result in reduced write performance.

REST API settings

You can set the `index . codec` parameter by using REST API.

Note:

- `close` the corresponding index before running the command.
- $index_name: Replace with the index name you need to set.

```
PUT $ index_name / _ settings
{
  " index ": {
    " codec ": " best_compr ession "
  }
}
```

4 Plug-ins

4.1 Overview

Based on open-source community plug-ins, Alibaba Cloud Elasticsearch provides a variety of plug-ins and other extensions. The topic describes the plug-ins feature of Alibaba Cloud Elasticsearch. This feature allows you to use plug-ins provided by Alibaba Cloud Elasticsearch to meet business demands.

Use plug-ins

Log on to the Alibaba Cloud Elasticsearch console, and select Instance ID > Plug-ins.

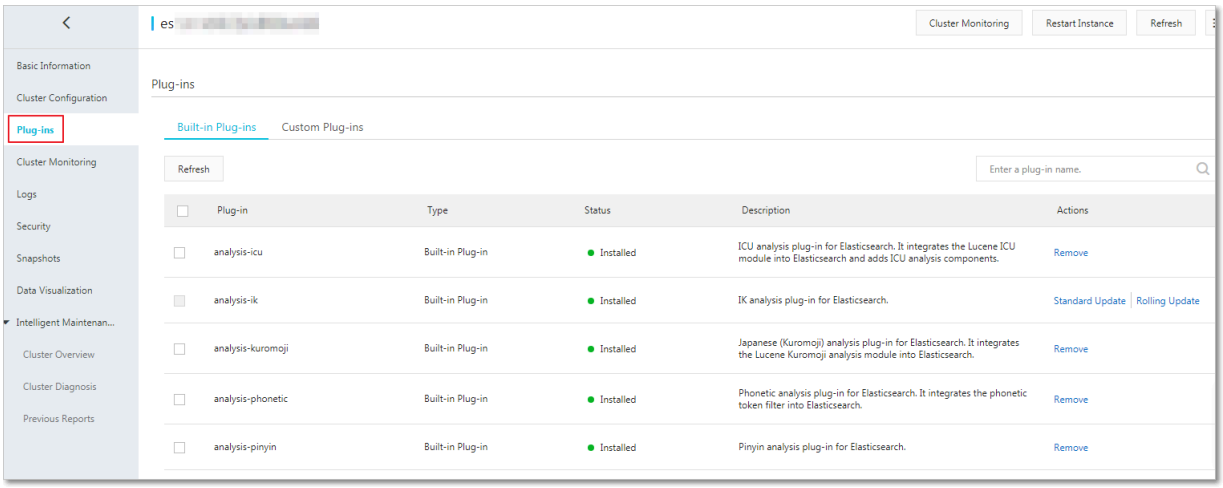

On the Plug-ins page, you can check Built-in Plug-ins and Custom Plug-ins.

- Built-in plug-ins

  You cannot remove the analysis-ik and elasticsearch-repository-oss plug-ins in the Built-in Plug-ins list. With the analysis-ik plug-in, you can use the standard update or rolling update method to upload and update IK dictionaries. This allows you to update customized dictionaries. For more information, see #unique_29.

- Custom plug-ins

  You can upload, install, and remove custom plug-ins to meet your business demands. For more information, see #unique_30.

4.2 Custom plug-ins

This topic describes how to upload, install, and remove custom plug-ins for Alibaba Cloud Elasticsearch.

Upload and install a custom plug-in

Notice:

- After you upload a custom plug-in, Elasticsearch needs to restart the cluster to install the plug-in. The custom plug-in may adversely affect the stability of the cluster. Make sure that the uploaded custom plug-in is secure and can run normally on the cluster.
- When the Elasticsearch cluster is upgraded, it will not upgrade the custom plug-in at the same time. To upgrade the plug-in, you have to upload the new version of the plug-in to the cluster.
- If your plug-in is not included in any privacy policies, we hope that you can make it open-source to help us develop our open-source community plug-ins.

1. On the Plug-ins page, select Custom Plug-ins > Upload.

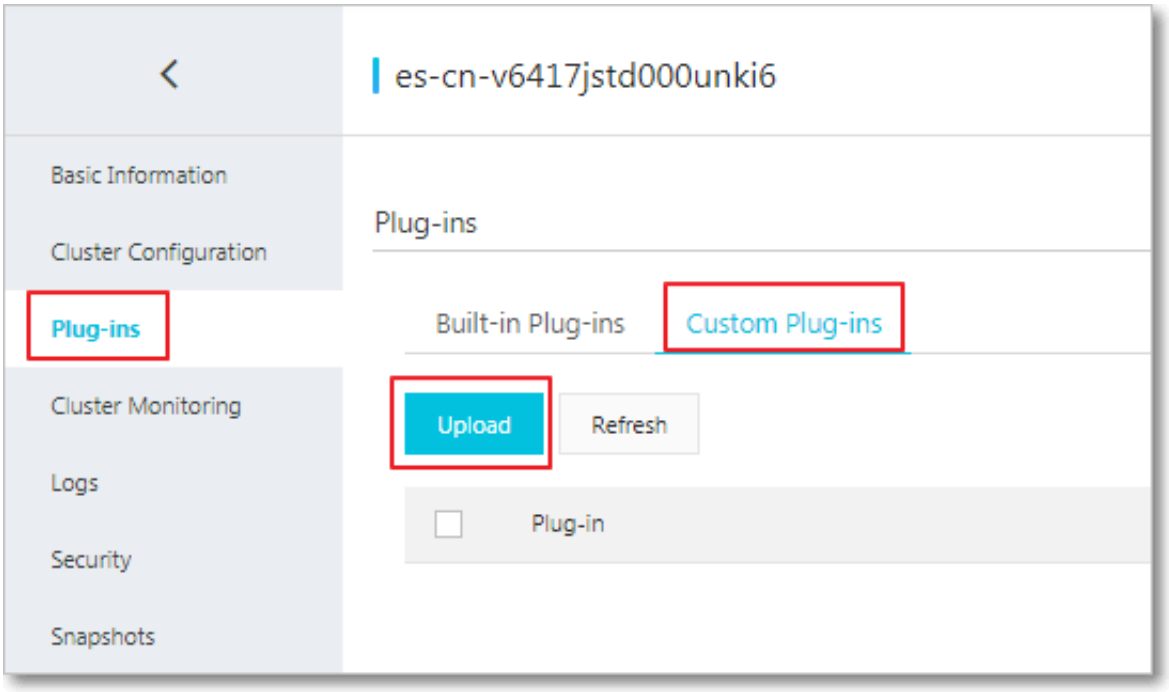

2. In the Upload Plug-in dialog box, click Select files, or drag and drop files to this area, and select the custom plug-in that you want to upload.

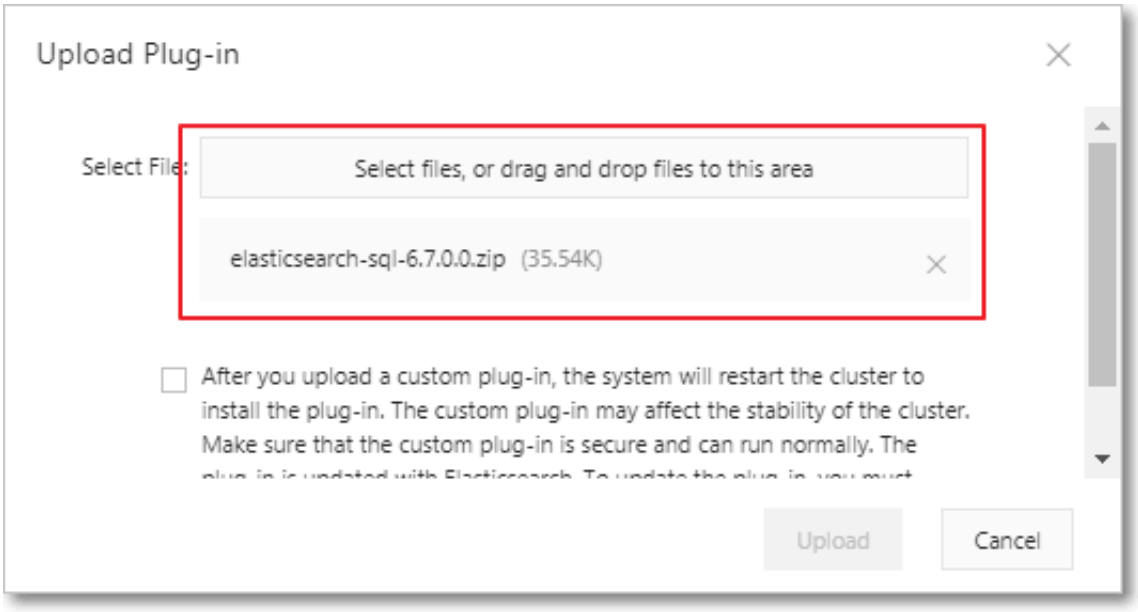

You can also drag and drop a custom plug-in file to this area to upload the plug-in. As shown in the preceding figure, the plug-in file Elasticsearch-sql-6.7.0.0 has been added.

Note:
You can add multiple custom plug-ins in the same way.

3. Read the agreement carefully, select the check box, and click Upload.

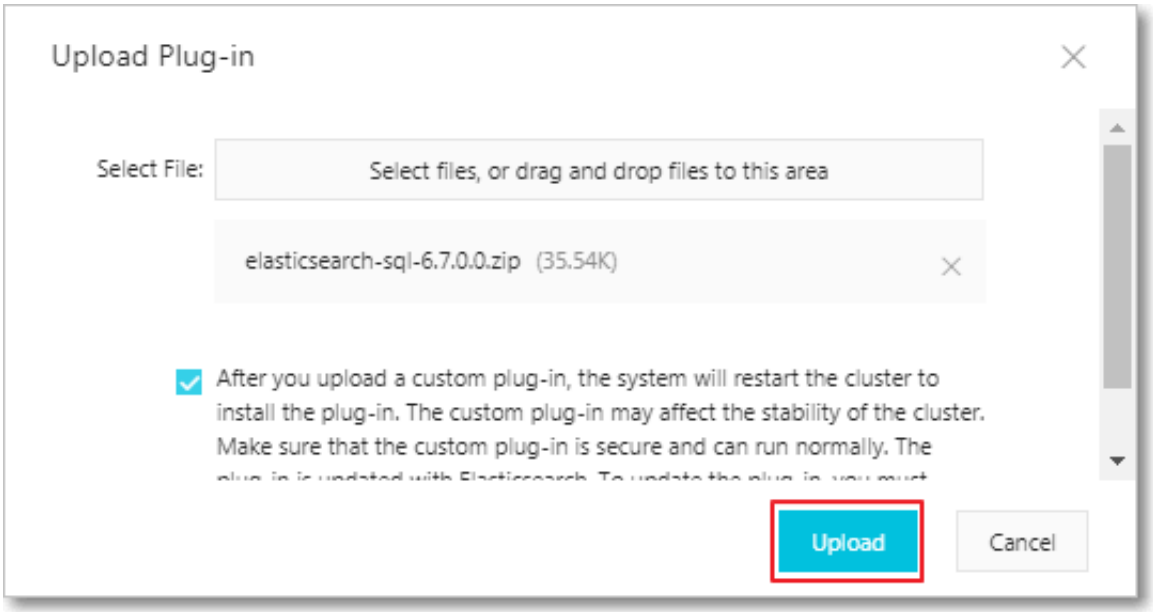

After you upload the plug-in, Elasticsearch will restart the cluster to install the plug-in. After the cluster is restarted, you can check the plug-in in the Custom

Plug-ins list. The status of the plug-in that you upload will display Installed. This indicates that the plug-in has been uploaded and installed successfully.

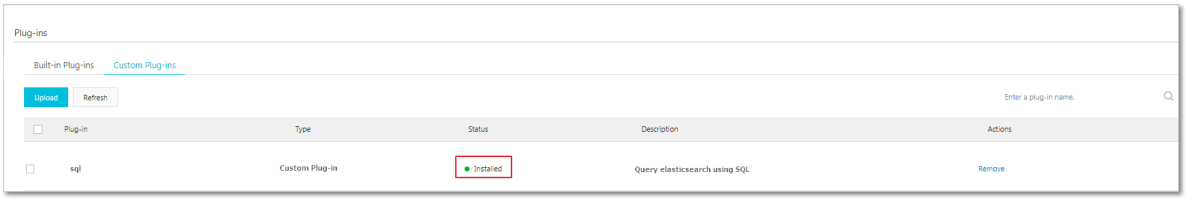

If you no longer need the plug-in, click Remove on the right side to remove the plug-in. For more information, see #unique_29/unique_29_Connect_42_section_d0y_kyx_fu0.

5 Cluster monitoring

5.1 Cluster monitoring

Cluster alarm

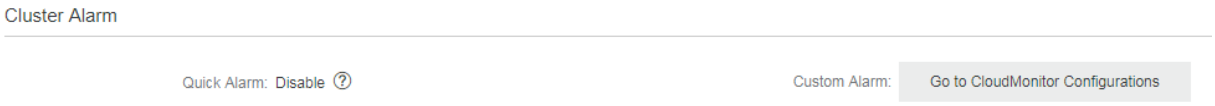

Quick alarm

1. Elasticsearch supports quick alarm. This feature is disabled by default. You can go to the clusters list page and click Quick Alarm to enable or disable this feature.

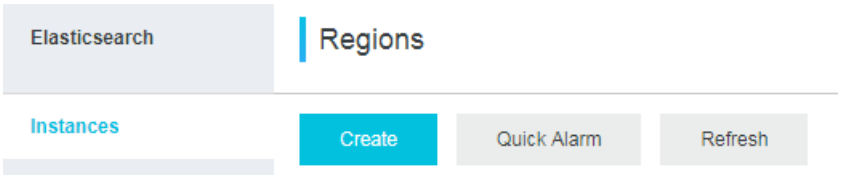

2. If this feature is disabled, click Quick Alarm, and then click Enable Now in the dialog box to manually enable it.

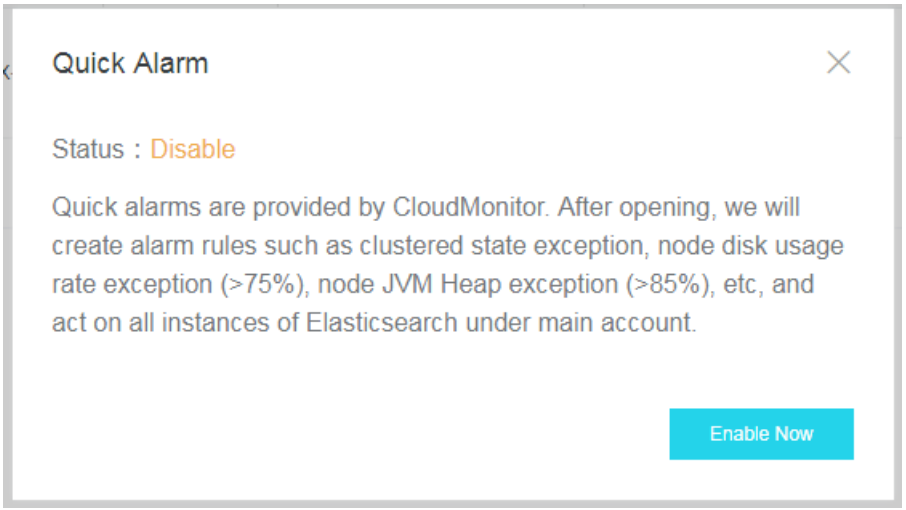

Custom alarms

You can click Cluster Monitor to create custom alarm rules. For more information about creating alarm rules, see #unique_34.

Cluster monitor

You can view Elasticsearch instance parameters and workloads.

Preset time

You can click a time option to view cluster metrics that are collected in the specified time period.

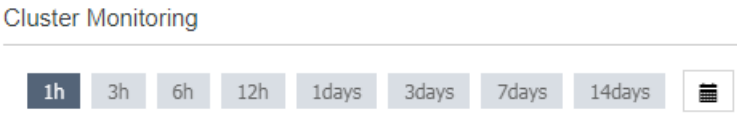

Custom cluster monitoring time

You can click Custom to specify the start time and end time to define a time window and view cluster monitoring data collected within the time window.

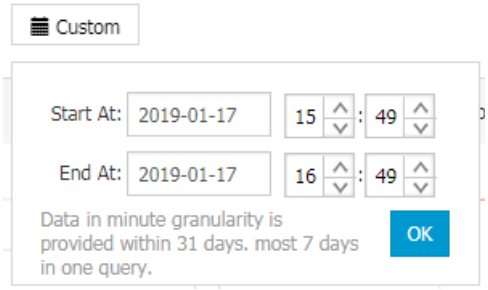

Note:

You can query up to 7 continuous days of data in the last 31 days by the minute.

Cluster monitoring metrics

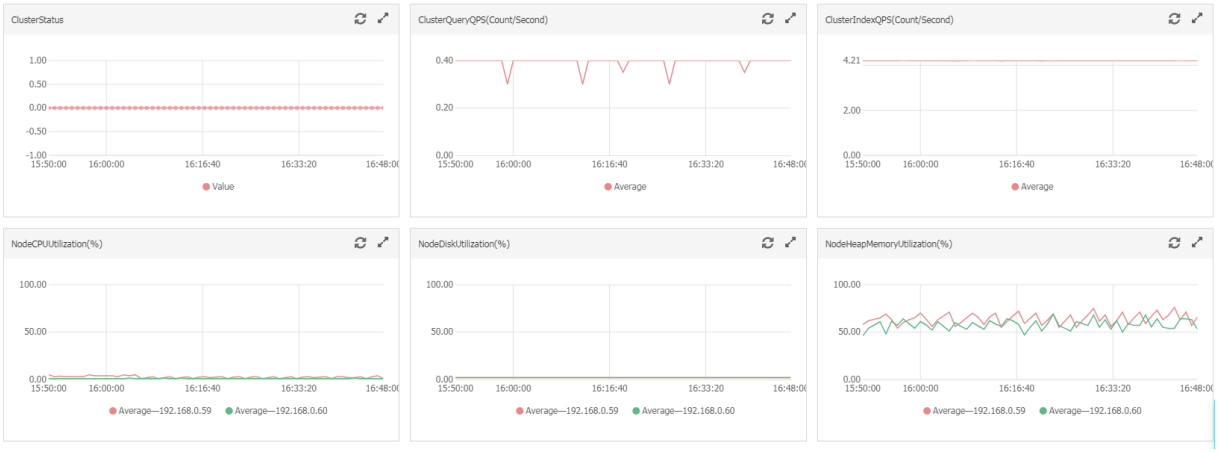

6 Query logs

Alibaba Cloud Elasticsearch allows you to search and view multiple types of logs, including the Elasticsearch instance log, search slow log, indexing slow log, and GC log.

You can search for specific log entries by entering keywords and setting a time range . All Alibaba Cloud Elasticsearch log entries are sorted in time descending order. You can search for log entries that are stored within the last seven days.

Alibaba Cloud Elasticsearch allows you to use Lucene to query logs. For more information, see Query String Query.

 Note:
Due to the restrictions Elasticsearch puts on query conditions, a maximum of 10,000 log entries can be returned. If the log entries that you have queried are not contained in the returned 10,000 log entries, set a more specific time range to narrow down the search results.

Example

The following example shows how to search for Elasticsearch instance logs whose content contains the keyword `health` , level is set to `info` , and host is set to `192 . 168 . 1 . 123 .`

1. Log on to the Alibaba Cloud Elasticsearch console, select the target instance, and click Manage in the Actions column to go to the Basic Information page. On the Basic Information page, click Logs in the left-side navigation pane and then click the Instance Log tab.
2. Enter `host : 192 . 168 . 1 . 123  AND  content : health  AND  level : info` in the search box.

3. Specify a time range and click Search.

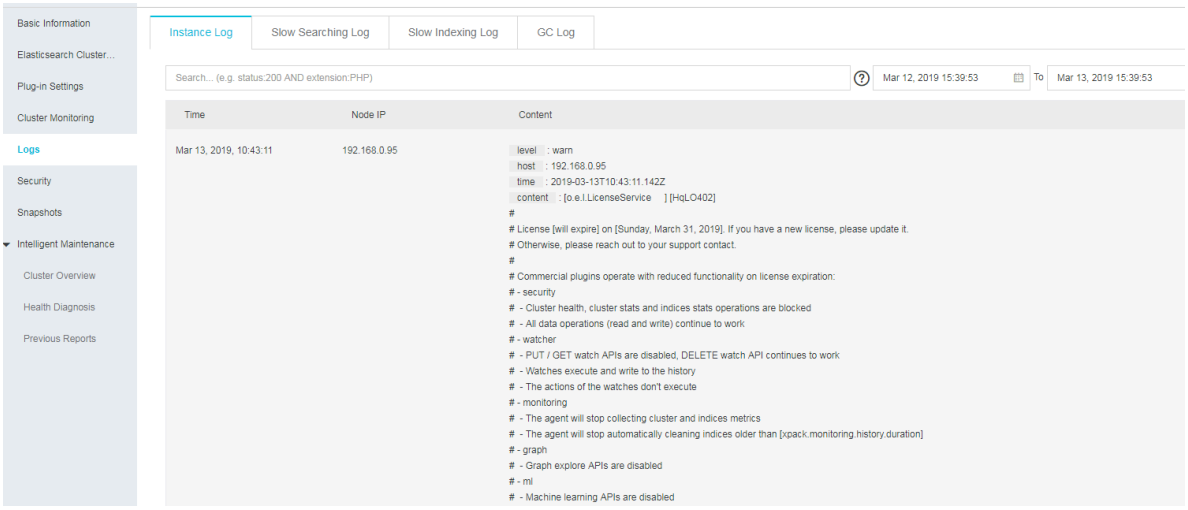

Note:

- If you do not specify the end time, it defaults to the current system time.
- If you do not specify the start time, it defaults to one hour later than the end time.
- The word `AND` connecting search conditions that you enter in the search box must be capitalized.

Log description

You can view log entries that are retrieved based on specified search conditions on the log search page. Each log entry contains the following parts: Time, Node IP, and Content.

Time

The time when the log entry was created.

Node IP

The IP address of the Alibaba Cloud Elasticsearch node.

Content

The information about the level, host, time, and content.

- level: the level of the log entry. Log levels include trace, debug, info, warn, and error. GC log entries do not have levels.
- host: indicates the IP address of the Elasticsearch node. You can view the IP address on the Nodes tab in the Kibana console.
- time: indicates the time when the log entry was created.

- **content: displays major information about the log entry.**

7 Security configuration

This topic describes the security configuration of Alibaba Cloud Elasticsearch, including the Elasticsearch instance password, public network whitelist, VPC whitelist, and HTTPS protocol.

Network settings

You can reset the Elasticsearch instance password, configure the VPC whitelist, enable Public network access, and configure the Public network whitelist and Enable HTTPS in network settings.

Elasticsearch instance password

To reset the Elasticsearch instance password, click Reset, and enter a new password for the administrator account elastic. After you reset the password, it takes up to 5 minutes for the new password to take effect.

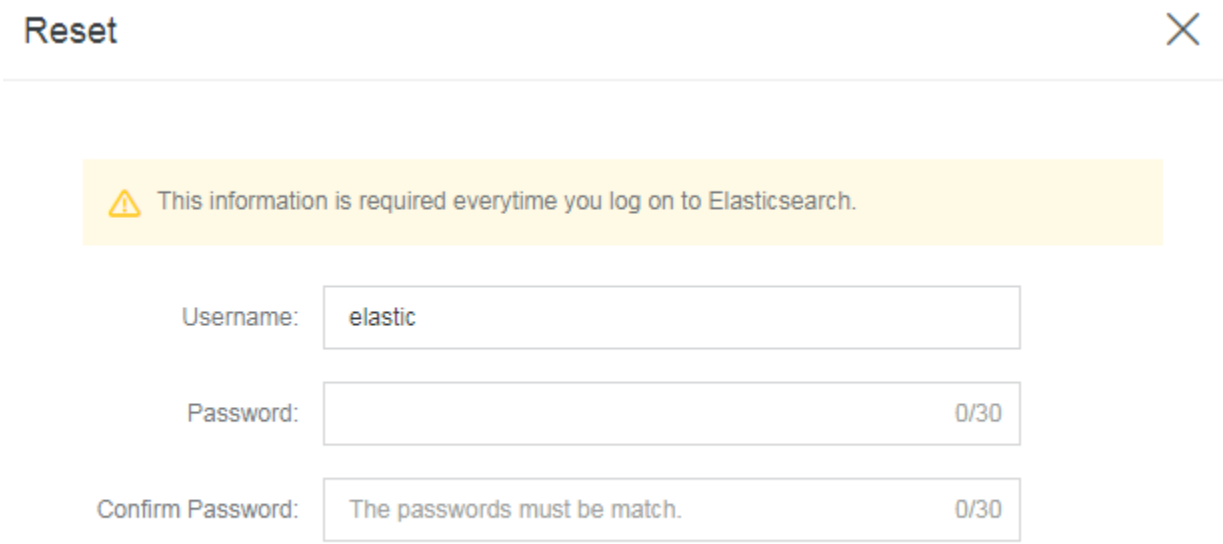

If you use the elastic account to log on to the Alibaba Cloud Elasticsearch instance or Kibana console, then you must use the new password.

Note:

- The reset operation only resets the password of the elastic account. The operation does not reset the password of other accounts that are used to log on to the instance. We recommend that you do not use the elastic account to log on to your Alibaba Cloud Elasticsearch instance.
- The Reset operation does not restart the Alibaba Cloud Elasticsearch instance.

VPC whitelist

When you need to access an Alibaba Cloud Elasticsearch instance from an ECS instance in a VPC network, you must add the IP address of the ECS instance to the VPC whitelist.

Click Update, enter the IP address in the VPC whitelist dialog box, and click OK.

You can add IP addresses and CIDR blocks to the whitelist in the format of `192 . 168 . 0 . 1` and `192 . 168 . 0 . 0 / 24` , respectively. Separate these IP addresses and CIDR blocks with commas (,). Enter `127 . 0 . 0 . 1` to forbid all IPv4 addresses or enter `0 . 0 . 0 . 0 / 0` to allow all IPv4 addresses.

Note:

- By default, all private IPv4 addresses are allowed to access Elasticsearch.
- The VPC whitelist is used to control access from internal network addresses in VPC networks.

Public network access

Click the Public Network Access switch to enable public network access. After this feature is enabled, the switch is in green. By default, the switch is in gray, which means that public network access is disabled. To access your Alibaba Cloud Elasticsearch instance through the Internet, you must enable public network access.

Public network whitelist

Before you configure the public network whitelist, you must toggle on the Public Network Access switch. By default, the public network access feature forbids all public network addresses.

To access your Alibaba Cloud Elasticsearch instance through the Internet, you must add the IP address of your client to the public network whitelist.

You can add IP addresses and CIDR blocks in the format of `192 . 168 . 0 . 1` and `192 . 168 . 0 . 0 / 24`, respectively. Separate these IP addresses and CIDR blocks with commas (,). Enter `127 . 0 . 0 . 1` to forbid all IPv4 addresses or enter `0 . 0 . 0 . 0 / 0` to allow all IPv4 addresses.

If your Elasticsearch instance is deployed in the China (Hangzhou) region, then you can add IPv6 addresses and CIRD blocks to the whitelist in the format of `2401 : b180 : 1000 : 24 :: 5` and `2401 : b180 : 1000 ::/ 48`, respectively. Enter `:: 1` to forbid all IPv6 addresses or enter `::/ 0` to allow all IPv6 addresses.

Enable HTTPS

Hypertext Transfer Protocol Secure (HTTPS) is a secure version of HTTP. HTTPS uses Secure Socket Layer (SSL) for secure data transmission. This means that HTTPS still uses HTTP for communications. SSL is used to encrypt the data.

`Procedure`

Notice:

- Alibaba Cloud Elasticsearch allows you to enable and disable HTTPS. To protect your data, we recommend that you enable HTTPS.

- Before you enable HTTPS, you must purchase client nodes.

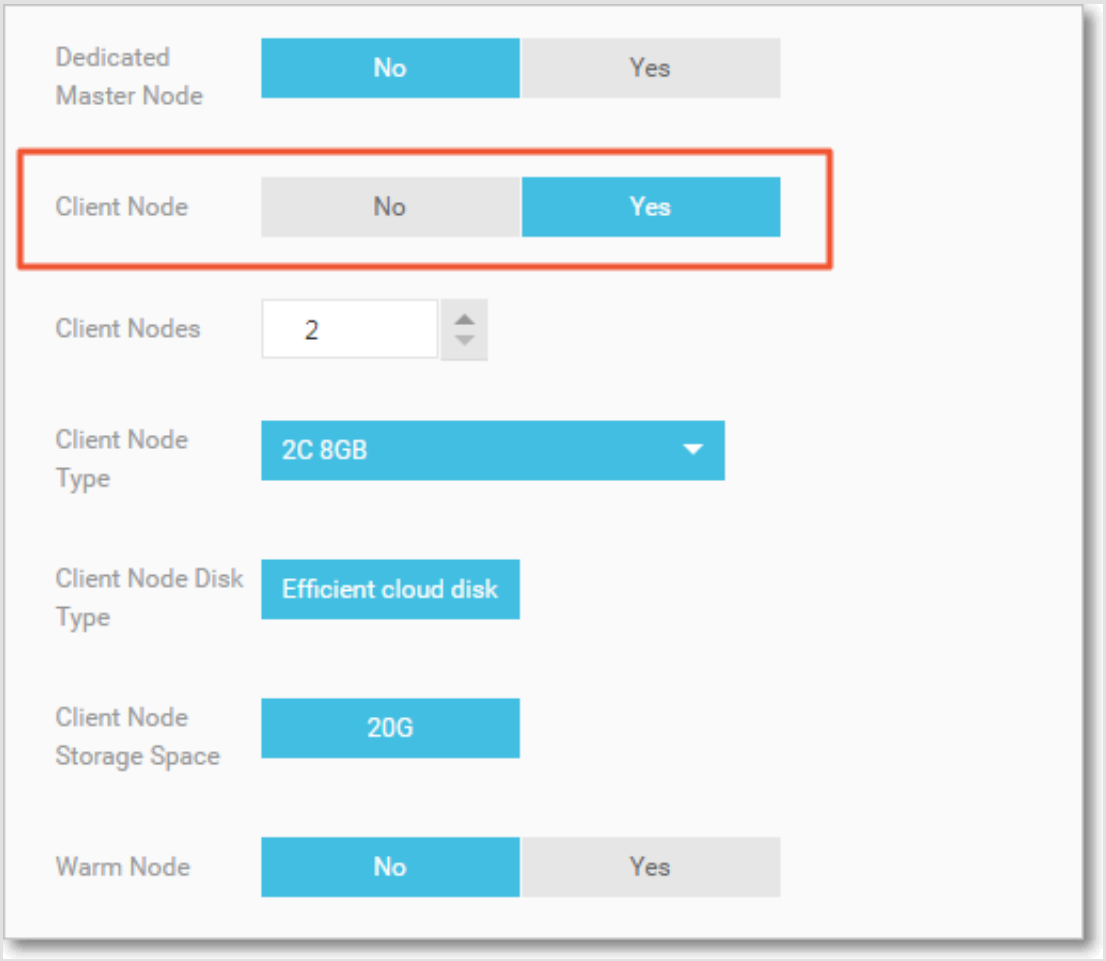

1. Log on to the Alibaba Cloud Elasticsearch console, click Instance ID/Name > Security, and click the HTTPS switch to enable HTTPS.

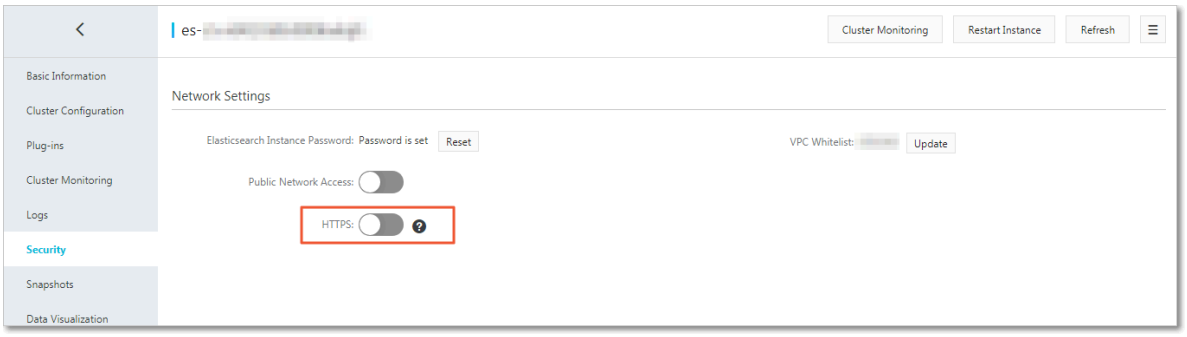

Notice:

- Before you enable HTTPS, you must update the code of the client that is used to access the Elasticsearch instance. Otherwise, you may fail to access the instance. For more information, see Sample client code for enabling or disabling HTTPS.

- During the process of enabling or disabling HTTPS, the services running on the instance will be interrupted and the instance will be restarted. Before you enable or disable HTTPS, make sure that your businesses will not be adversely affected.

2. In the Confirm Operation dialog box, select I have edited the code of the Elasticsearch instance, and then click OK.

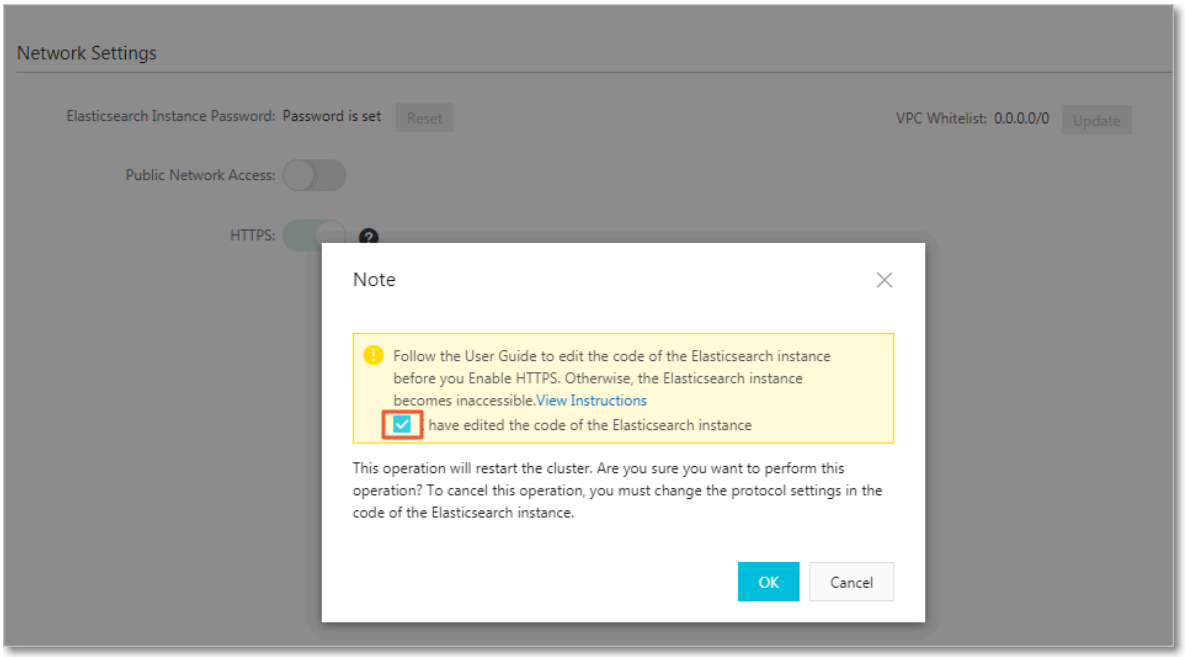

Note:

If you have not purchased client nodes, after you enable HTTPS, the system prompts a notification requiring you to purchase client nodes. You can follow the instructions to purchase client nodes.

After you confirm to enable or disable HTTPS, the instance will restart. You can click the Tasks icon in the upper-right corner to check the progress. After the instance is restarted, you can then access the instance through HTTPS.

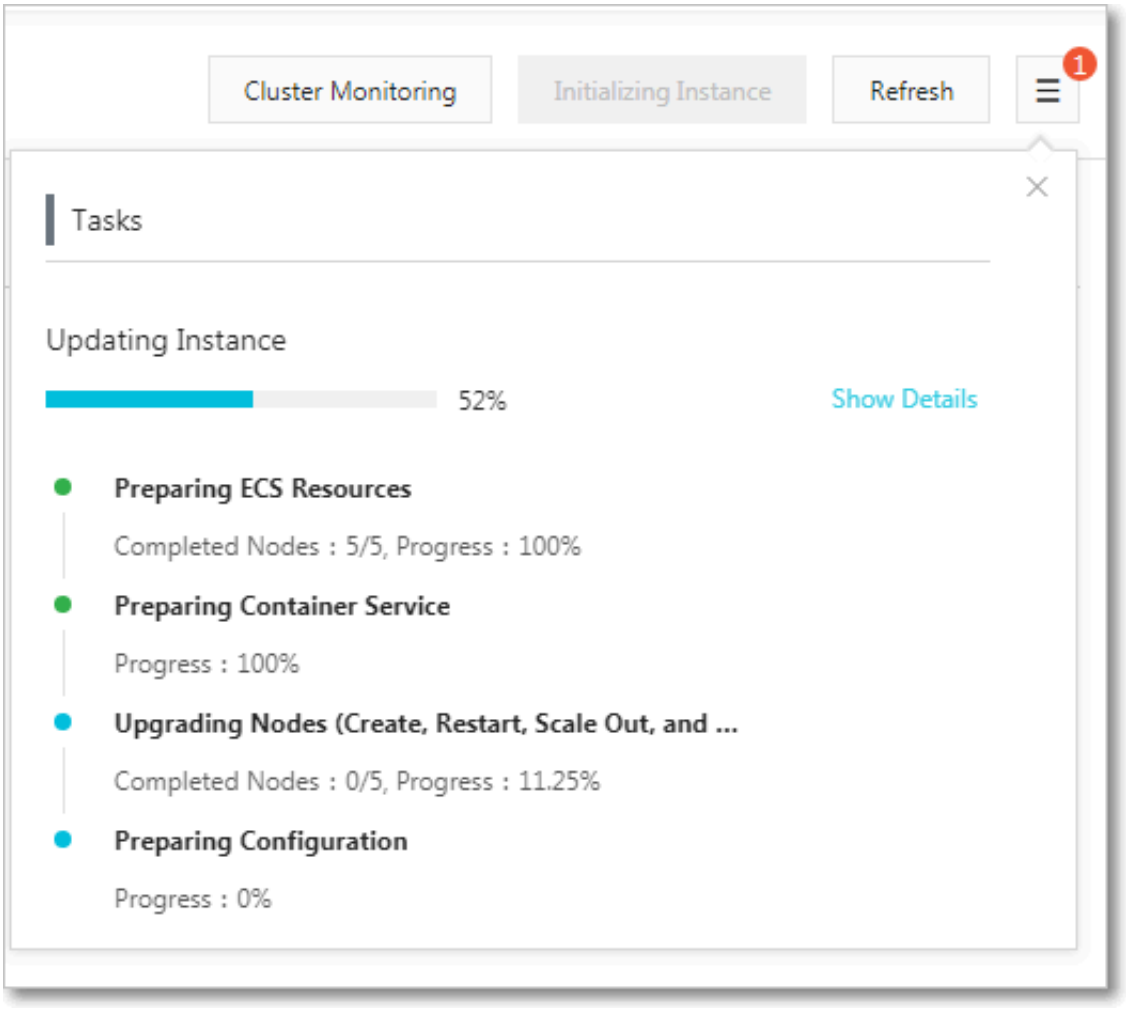

Sample client code for enabling or disabling HTTPS

The following example shows the changes that need to be made to the code of the Elasticsearch REST client after you enable HTTPS.

- The code of the REST client before HTTPS is enabled:

```
final Credential sProvider credential sProvider = new
BasicCrede ntialsProv ider ();
        credential sProvider . setCredent ials ( AuthScope . ANY
,
            new UsernamePa sswordCred entials (" elastic ", "
Your password "));
RestClient Builder restClient Builder = RestClient . builder
(
```

```
            new   HttpHost (" es - cn - xxxxx . elasticsea  rch .
aliyuncs . com ",  9200 ));
        RestClient   restClient  =  restClient  Builder .
setHttpCli  entConfigC  allback (
            new   RestClient  Builder . HttpClient  ConfigCall
back () {
               @ Override
                public   HttpAsyncC  lientBuild  er   customizeH
ttpClient ( HttpAsyncC  lientBuild  er   httpClient  Builder ) {
                    return   httpClient  Builder . setDefault
Credential  sProvider ( credential  sProvider );
                }
           }). build ();
```

- The code of the REST client after HTTPS is enabled:

```
final   Credential  sProvider   credential  sProvider  =  new
BasicCrede  ntialsProv  ider ();
        credential  sProvider . setCredent  ials ( AuthScope . ANY
,
            new   UsernamePa  sswordCred  entials (" elastic ", "
Your   password "));
RestClient  Builder   restClient  Builder  =  RestClient . builder
(
            new   HttpHost (" es - cn - xxxxx . elasticsea  rch .
aliyuncs . com ",  9200 , " https "));
        RestClient   restClient  =  restClient  Builder .
setHttpCli  entConfigC  allback (
            new   RestClient  Builder . HttpClient  ConfigCall
back () {
               @ Override
                public   HttpAsyncC  lientBuild  er   customizeH
ttpClient ( HttpAsyncC  lientBuild  er   httpClient  Builder ) {
                    return   httpClient  Builder . setDefault
Credential  sProvider ( credential  sProvider );
                }
           }). build ();
```

As shown in the preceding example, after you enable HTTPS, you must include the `https` parameter in `HttpHost : new   HttpHost (" es - cn - xxxxx . elasticsea  rch . aliyuncs . com ",  9200 , " https "));`

8 Data backup

8.1 Snapshots

This topic describes the snapshot feature of Alibaba Cloud Elasticsearch.

Log on to the Alibaba Cloud Elasticsearch console, click Instance Name > Snapshots to navigate to the Snapshots (Free Trial) page.

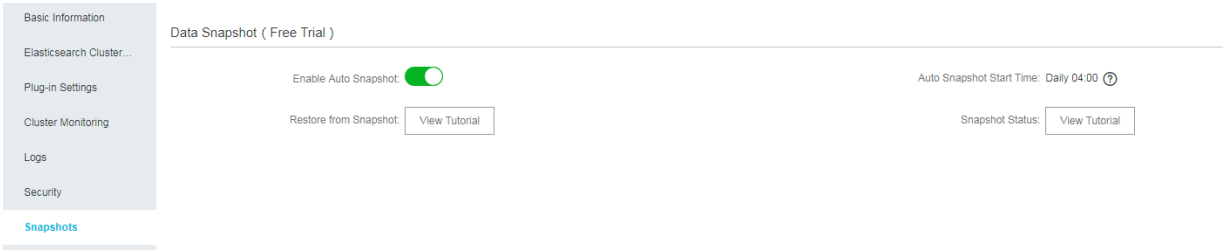

Parameter	Description
Auto Snapshot	When the Auto Snapshot switch is in the green color, auto snapshot is enabled. By default, auto snapshot is disabled.
Auto Snapshot Period	If auto snapshot is disabled, the You must enable auto snapshot first message is displayed. Notice: Auto snapshot uses the system time of the region where the Elasticsearch instance is created. Do not perform any snapshot operations when the system is creating snapshots.

Parameter	Description
Edit Configuration	If auto snapshot is enabled, you can click Edit Configuration in the upper-right corner to open the Auto Snapshot Configuration dialog box and then set the time for creating snapshots. 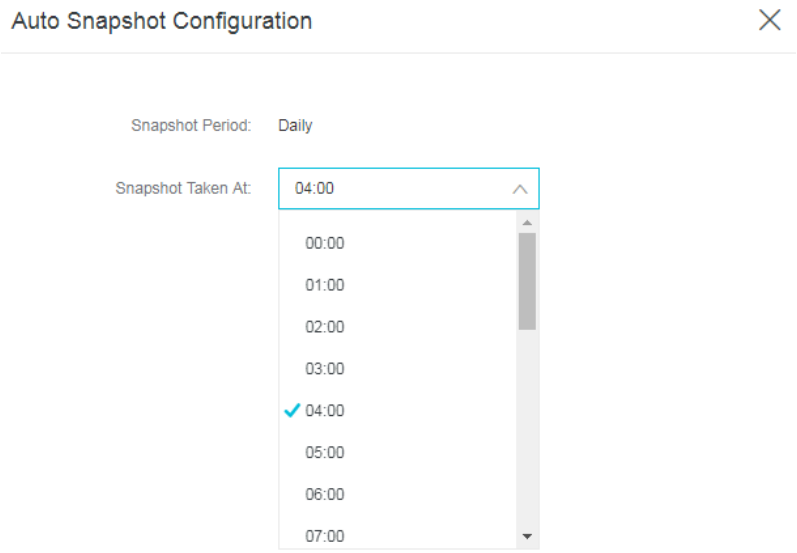 **Notice:** · The Frequency parameter is set to Daily. · The Create Snapshot At parameter specifies the specific time for creating a snapshot daily. Valid values are from 0 to 23 hours. · Alibaba Cloud Elasticsearch only stores snapshots that are created in the last three days.
Restore from Snapshot	Click View Tutorial to learn how to restore data from a snapshot.
Snapshot Status	Click View Tutorial to learn more information about snapshot status.

8.2 View backup information

View automatic backup information

After enabling automatic backup, you can log on to the Kibana console that has been integrated into Alibaba Cloud Elasticsearch and run the Elasticsearch `snapshot` command in Dev Tools to view snapshots.

View all snapshots

Run the following command to view all the snapshots that are located in the aliyun_auto_snapshot repository.

```
GET   _   snapshot / aliyun_aut  o_snapshot / _   all
```

Response:

```
{
  " snapshots ": [
    {
      " snapshot ": " es - cn - abcdefghij  klmn_20180  628092236 ",
      " uuid ": " n7YIayyZTm  2hwg8BeWby  dA ",
      " version_id ":  5050399 ,
      " version ": " 2 . 0 . 0 ",
      " indices ": [
        ". kibana "
      ],
      " state ": " SUCCESS ",
      " start_time ": " 2018 - 06 - 28T01 : 22 : 39 . 609Z ",
      " start_time  _in_millis ":  1530148959  609 ,
      " end_time ": " 2018 - 06 - 28T01 : 22 : 39 . 923Z ",
      " end_time_i  n_millis ":  1530148959  923 ,
      " duration_i  n_millis ":  314 ,
      " failures ": [],
      " _shards " : {
        " total ": 1
        " failed " :  0
        " successful ":  1 ,
      }
    },
    {
      " snapshot ": " es - cn - abcdefghij  klmn_20180  628092500 ",
      " uuid ": " frdl1YFzQ5  Cn5xN9ZWuK  LA ",
      " version_id ":  5050399 ,
      " version ": " 2 . 0 . 0 ",
      " indices ": [
        ". kibana "
      ],
      " state ": " SUCCESS ",
      " start_time ": " 2018 - 06 - 28T01 : 25 : 00 . 764Z ",
      " start_time  _in_millis ":  1530149100  764 ,
      " end_time ": " 2018 - 06 - 28T01 : 25 : 01 . 482Z ",
      " end_time_i  n_millis ":  1530149101  482 ,
      " duration_i  n_millis ":  718 ,
      " failures ": [],
      " _shards " : {
        " total ": 1
        " failed " :  0
        " successful ":  1 ,
      }
    }
  ]
```

```
}
```

- state: Specifies the status of a snapshot. The snapshot status includes the following:
  - `IN_PROGRES S` : The snapshot is being restored.
  - `SUCCESS` : The snapshot has been restored and all shards have been successfully stored.
  - `FAILED` : The snapshot has been restored with an error. Some data cannot be stored.
  - `PARTIAL` : The snapshot has been successfully restored to an instance. However, one or more shards cannot be stored.
  - `INCOMPATIB LE` : The snapshot version is incompatible with the current instance version.

View specified snapshot

Run the following command to view detailed information about the specified snapshot in the aliyun_auto_snapshot repository.

```
GET _ snapshot / aliyun_aut o_snapshot /< snapshot >/ _ status
```

- `< Snapshot >`: Specifies the name of the snapshot, for example, `Es - cn - abcdefghij klmn_20180 628092236` .

Response:

```
{
  " Snapshots ": {
    {
      " snapshot ": " es - cn - abcdefghij klmn_20180 628092236 ",
      " repository ": " aliyun_aut o_snapshot ",
      " uuid ": " n7YIayyZTm 2hwg8BeWby dA ",
      " state ": " SUCCESS ",
      " shards_sta ts ": {
        " initializi ng ": 0 ,
        " started ": 0 ,
        " finalizing ": 0 ,
        " done ": 1 ,
        " failed " : 0
        " total ": 2
      },
      " stats ": {
        " number_of_ files ": 4 ,
        " processed_ files ": 4 ,
        " total_size _in_bytes ": 3296 ,
        " processed_ size_in_by tes ": 3296 ,
        " start_time _in_millis ": 1530148959 688 ,
        " time_in_mi llis ": 77
      },
      " indices ": {
        ". kibana ": {
```

```
          " shards_sta  ts ": {
            " initializi  ng ":  0 ,
            " started ":  0 ,
            " finalizing ":  0 ,
            " done ":  1 ,
            " failed " :  0
            " total ":  2
          },
          " stats ": {
            " number_of_  files ":  4 ,
            " processed_  files ":  4 ,
            " total_size  _in_bytes ":  3296 ,
            " processed_  size_in_by  tes ":  3296 ,
            " start_time  _in_millis ":  1530148959  688 ,
            " time_in_mi  llis ":  77
          },
          " shards ": {
            " 0 ": {
              " stage ": " DONE ",
              " stats ": {
                " number_of_  files ":  4 ,
                " processed_  files ":  4 ,
                " total_size  _in_bytes ":  3296 ,
                " processed_  size_in_by  tes ":  3296 ,
                " start_time  _in_millis ":  1530148959  688 ,
                " time_in_mi  llis ":  77
              }
            }
          }
        }
      }
    }
  ]
}
```


8.3 Auto snapshot guide

Enable auto snapshot

1. Log on to the Alibaba Cloud Elasticsearch console.
2. On the Instances page, click the target instance ID. You will be directed to the Basic Information page.
3. In the left-side navigation pane, click Snapshots.

4. On the Snapshots page, switch on Enable Auto Snapshot.

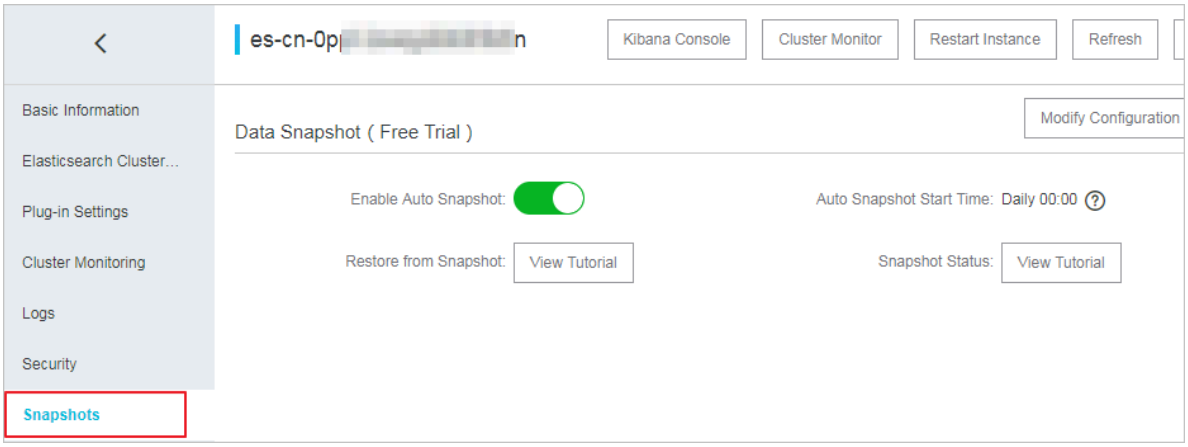

5. Click Modify Configuration in the upper-right corner to set the time when the daily snapshot is created.

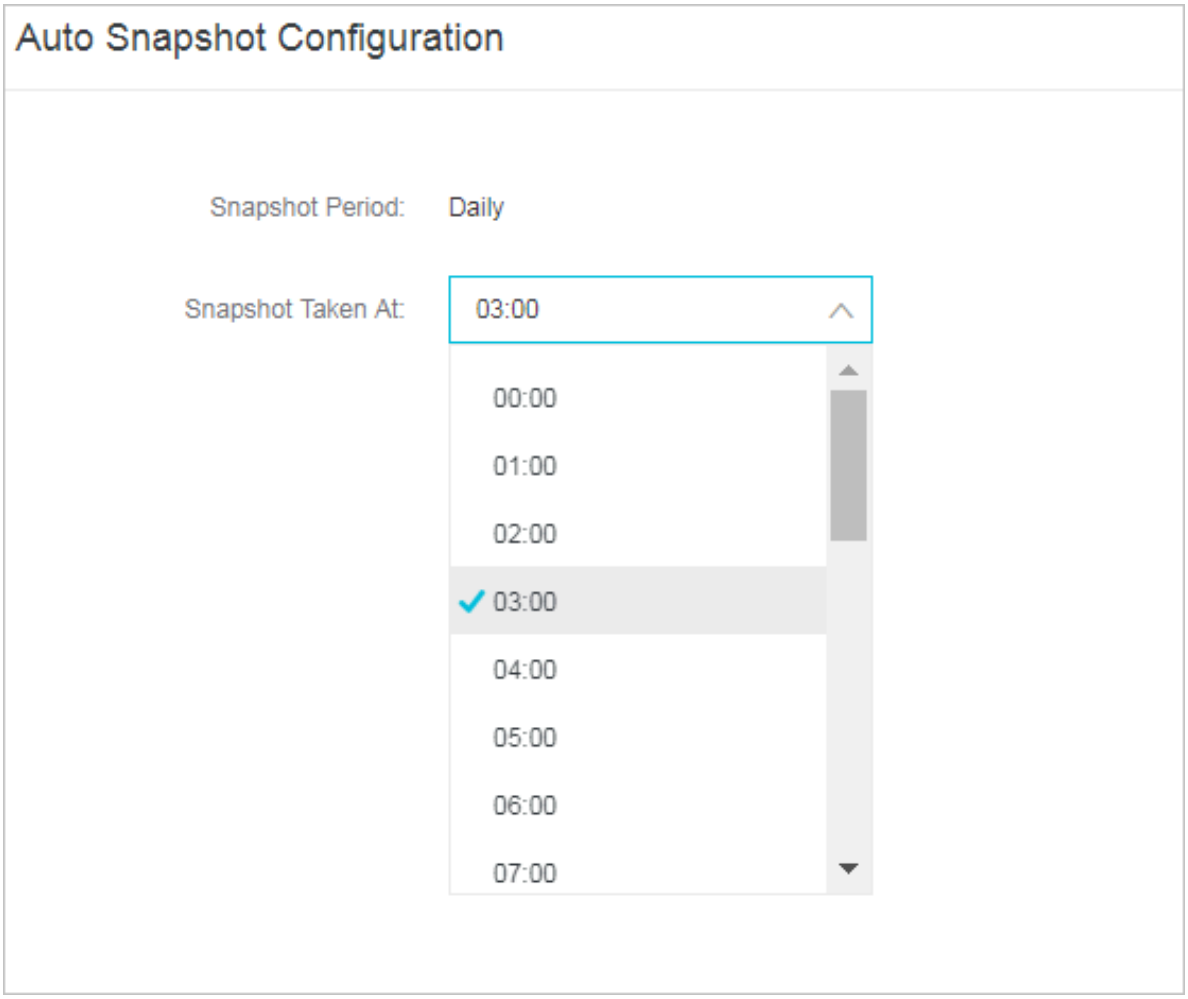

Restore snapshots into instances

If you have enabled auto snapshot for a specified Alibaba Cloud Elasticsearch instance, snapshots will be automatically created on a daily basis. You can call the

corresponding `snapshot` operation to restore a snapshot into the Alibaba Cloud Elasticsearch instance where the snapshot is created.

Note:

- The first snapshot is a complete backup created on a running Alibaba Cloud Elasticsearch instance. The following snapshots are created based on the incremental data of the Elasticsearch instance. Therefore, it takes a longer time to create the first snapshot, but a shorter time to create subsequent snapshots.
- A snapshot does not store monitoring data generated by an Alibaba Cloud Elasticsearch instance, such as the `. monitoring` and `. security_a udit` files.
- An auto snapshot can only be restored into the Alibaba Cloud Elasticsearch instance where the snapshot is created.
- An auto snapshot repository is created when the first snapshot is created.

View all snapshot repositories

You can run the `GET _snapshot` command to view all snapshot repositories.

The following response is returned:

```
{
  " aliyun_aut  o_snapshot ": {
    " type ": " oss ",
    " settings ": {
      " compress ": " true ",
      " base_path ": " xxxx ",
      " endpoint ": " xxxx "
    }
  }
}
```

- `aliyun_aut o_snapshot` : the name of the repository.
- `type` : the storage medium where snapshots are stored. This example uses Alibaba Cloud Object Storage Service (OSS).
- `compress : true` : enables compression of an index's metadata files.
- `base_path` : the location of the snapshots.
- `endpoint` : the region of the OSS instance.

View all snapshots

You can run the `GET _snapshot / aliyun_aut o_snapshot / _all` command to view all snapshots stored in the repository `aliyun_aut o_snapshot`

The following response is returned:

```
{
  " snapshots ": [
    {
      " snapshot ": " es - cn - abcdefghij  klmn_20180  627091600 ",
      " uuid ": " MMRniVLPRA  iawSCm8D8D  ug ",
      " version_id ":  5050399 ,
      " version ": " 5 . 5 . 3 ",
      " indices ": [
        " index_1 ",
        ". security ",
        ". kibana "
      ],
      " state ": " SUCCESS ",
      " start_time ": " 2018 - 06 - 27T01 : 16 : 01 . 009Z ",
      " start_time  _in_millis ":  1530062161  009 ,
      " end_time ": " 2018 - 06 - 27T01 : 16 : 05 . 632Z ",
      " end_time_i  n_millis ":  1530062165  632 ,
      " duration_i  n_millis ":  4623 ,
      " failures ": [],
      " shards ": {
        " total ":  12 ,
        " failed ":  0 ,
        " successful ":  12
      }
    }
  ]
}
```

Default parameters

Auto snapshots also support the following parameters that are not displayed:

- `max_snapsh ot_bytes_p er_sec : 40mb` : throttles per node snapshot rate. The default snapshot rate is 40 MB per second.
- `max_restor e_bytes_pe r_sec : 40mb` : throttles per node restore rate. The default restore rate is 40 MB per second.
- `chunk_size : Max 1Gb` : large files can be broken into smaller chunks during the snapshot process if needed. The maximum size of a chunk is 1 GB.

Restore a snapshot into an instance

You can run the `_restore` command to restore a snapshot into an instance:

- Restore all indexes in a specified snapshot that is stored in the `aliyun_aut o_snapshot` repository. The restore tasks are executed in the background.

```
POST _snapshot / aliyun_aut o_snapshot /< snapshot >/ _restore
```

`< snapshot >`: replace it with the name of the specified snapshot. Example: `es - cn - abcdefghij klmn_20180 627091600`

- Restore all indexes in the specified snapshot that is stored in the `aliyun_aut o_snapshot` repository, and receive a response after all restore tasks are completed:

  The `_restore` command runs restore tasks asynchronously. The Alibaba Cloud Elasticsearch instance will return a response immediately if the restore command is executable. Restore tasks are executed in the background. You can add the `wait_for_c ompletion` parameter to the command. This parameter requires the Alibaba Cloud Elasticsearch instance to return the response only after the restore tasks are completed.

```
POST _snapshot / aliyun_aut o_snapshot /< snapshot >/ _restore
? wait_for_c ompletion = true
```

`< snapshot >`: replace it with the name of the specified snapshot. Example: `es - cn - abcdefghij klmn_20180 627091600` .

- Restore indexes in the specified snapshot that is stored in the `aliyun_aut o_snapshot` repository, and rename the restored indexes. The restore tasks are executed in the background.

```
POST _snapshot / aliyun_aut o_snapshot /< snapshot >/ _restore
{
" indices ": " index_1 ",
" rename_pat tern ": " index_ (.+)",
" rename_rep lacement ": " restored_i ndex_ $ 1 "
}
```

  - `< snapshot >`: replace it with the name of the specified snapshot. Example: `es - cn - abcdefghij klmn_20180 627091600` .
  - `indices` : specifies names of the indexes that you need to restore.
  - `rename_pat tern` : uses a regular expression to match the restored indexes. This parameter is optional.
  - `rename_rep lacement` : renames the index that matches the regular expression. This parameter is optional.

9 Data visualization

9.1 Kibana

9.1.1 Log on to the Kibana console

This topic describes how to log on to the Kibana console. After you purchase an Alibaba Cloud Elasticsearch instance, Elasticsearch provides you a free Kibana node with one core and 2 GB of memory. The Kibana console supports data query, data visualization, and other features.

Prerequisites

To log on to the Kibana console, you must first purchase an Elasticsearch instance. Make sure that #unique_46/unique_46_Connect_42_section_bbj_euc_ly7 is Active.

Context

Alibaba Cloud Elasticsearch provides the Kibana console for you to scale your business. The Kibana console is a part of the Elasticsearch ecosystem, which has been seamlessly integrated into Elasticsearch. The Kibana console enables you to monitor the status of your Elasticsearch instances and manage these instances.

Procedure

1. Log on to the Alibaba Cloud Elasticsearch console, and clickInstance ID/Name > Data Visualization.

2. On the Data Visualization page, click Console under Kibana.

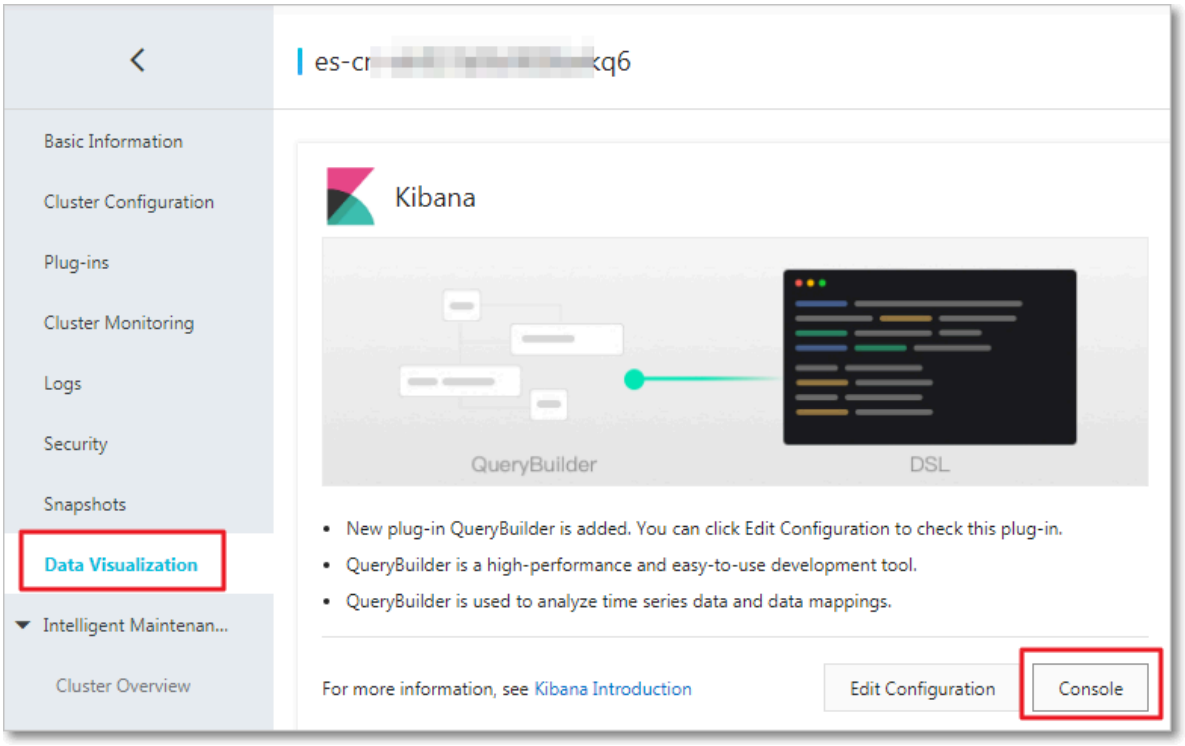

3. Enter the username and password on the logon page, and then click LOG IN.

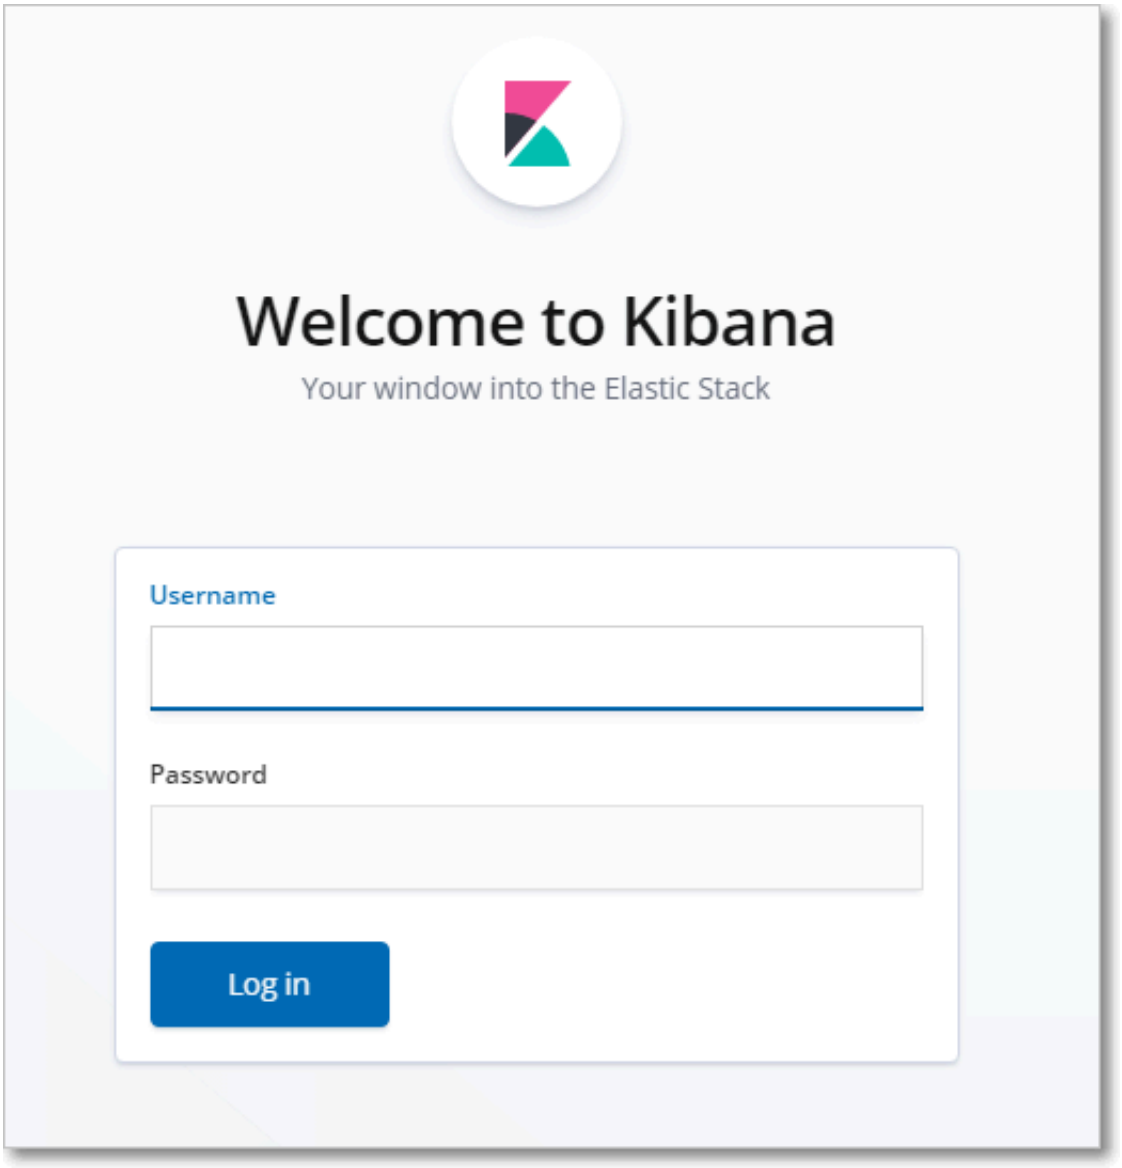

- Username: the default username is elastic.
- Password: enter the password that you have set when you purchase the Elasticsearch instance.

The following figure shows the Kibana console logged on from an Alibaba Cloud Elasticsearch instance 6.7. If you use other Elasticsearch versions, the actual console may look slightly different from the one in the figure.

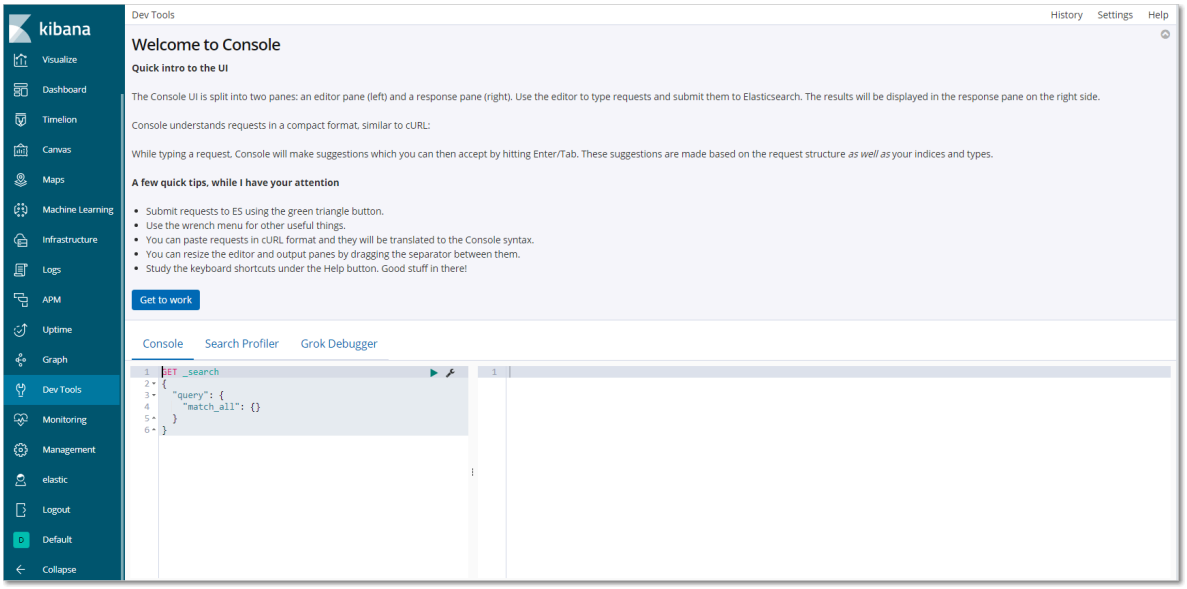

What's next

After you log on to the Kibana console, you can then perform operations such as query data or create dashboards. For more information, see Kibana User Guide.

9.1.2 Basic configuration (6.7.0)

This topic introduces the basic configuration of the Kibana node. You can switch the language of the Kibana console in the basic configuration.

> **Notice:**
> The basic configuration of the Kibana node is only available in Alibaba Cloud Elasticsearch 6.7.0 with Commercial Feature.

Switch the language of the Kibana console

1. Log on to the Alibaba Cloud Elasticsearch console, and then clickInstance ID/Name > Data Visualization.

2. Click Edit Configuration under Kibana to go to the Kibana Configuration page.

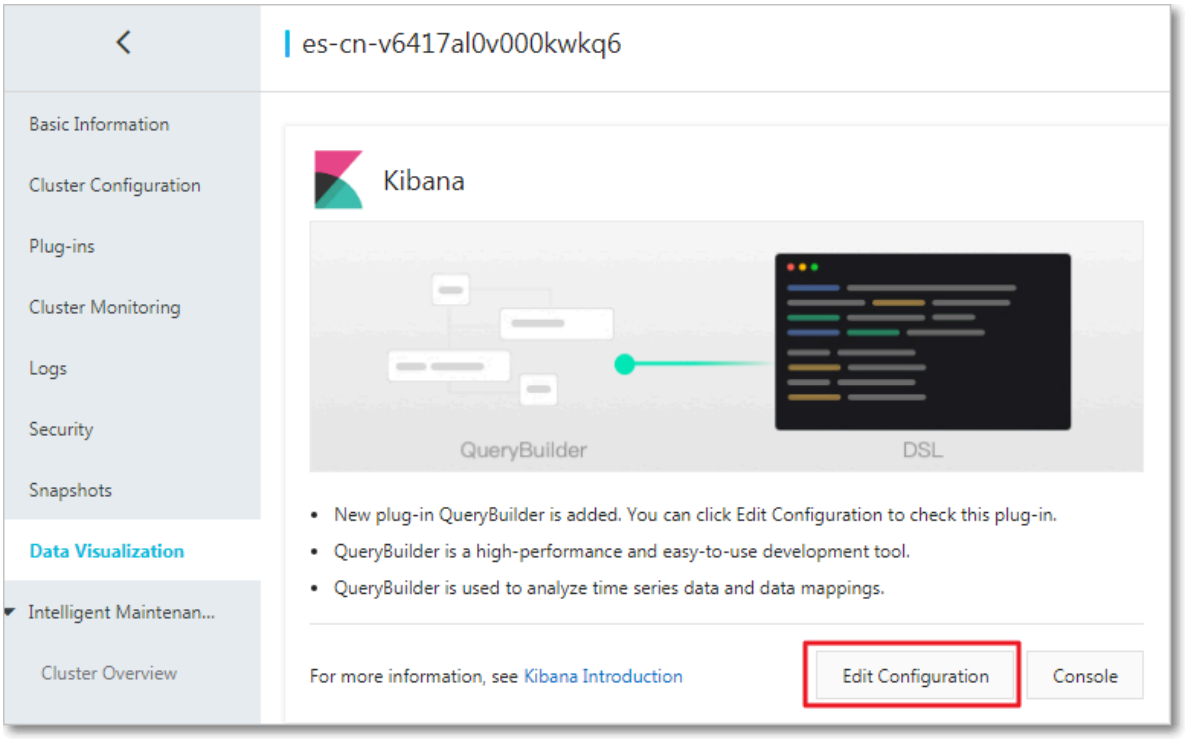

You can then view the Basic Configuration on the Kibana Configuration page. In the Basic Configuration area, follow these steps to switch the language of the Kibana console. By default, the language is set to English.

3. Click Edit Configuration on the right side of Basic Configuration.

Notice:

The system must restart the Kibana node for the changes to take effect. Make sure that the restart process does not affect your operations on the Kibana console before you perform the following steps:

4. On the Edit Basic Configuration page, select a language from the Select Language list, and click OK.

Note:

The Kibana console supports both English and Chinese. The default language is English.

After you click OK, the Kibana node will automatically restart. After the Kibana node is restarted, #unique_48 and verify that the console is switched to the selected language.

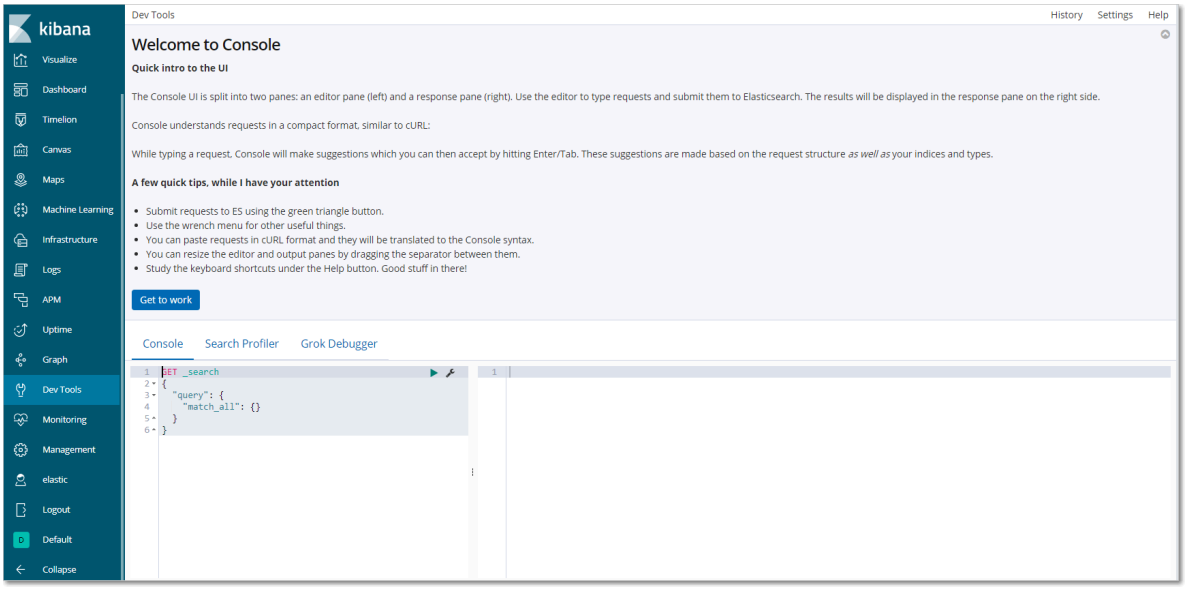

9.1.3 Network access configuration

This topic describes the network access configuration of Kibana clusters. The network access configuration includes the public network access configuration and Kibana whitelist.

Go to the network access configuration page

1. Log on to the Alibaba Cloud Elasticsearch console, and click Instance ID/Name > Data Visualization.
2. Click Edit Configuration under Kibana to go to the Kibana Configuration page.

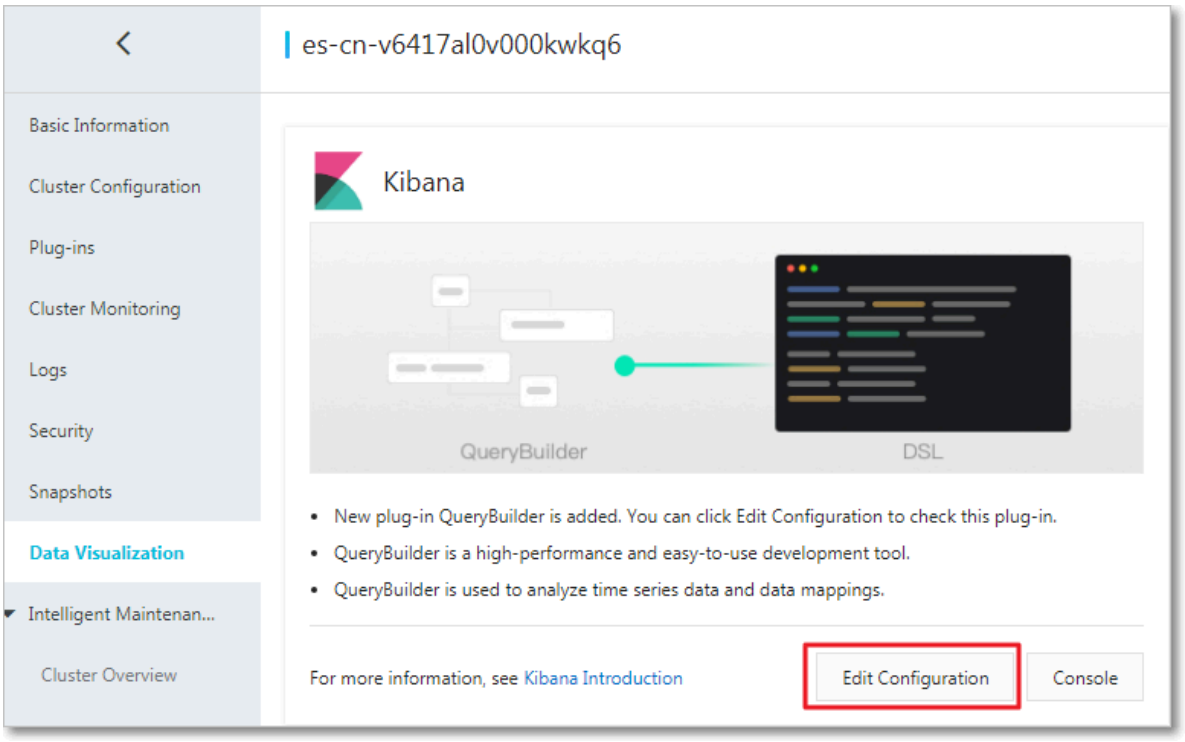

You can then view the Network Access Configuration on the Kibana Configuration page. In the Network Access Configuration area, you can enable or disable Public network access, and configure the Kibana whitelist. By default, the public network access feature is enabled.

Public network access

By default, the Public Network Access switch is toggled on (green). You can click the Public Network Access switch to disable this feature. When this feature is disabled,

the switch is gray. When the Public Network Access feature is disabled, you cannot log on to the Kibana console through the Internet.

Kibana whitelist

To configure the Kibana whitelist, click Update next to the Kibana whitelist, enter IP addresses into the dialog box, and click OK.

Note:
By default, all public network addresses are allowed to access the Kibana console.

The Kibana console supports both IP addresses and CIDR blocks. Enter IP addresses and CIDR blocks in the format of `192.168.0.1` and `192.168.0.0/24`, respectively. Separate these IP addresses and CIDR blocks with commas (,). You can enter `127.0.0.1` to forbid all IPv4 addresses or enter `0.0.0.0/0` to allow all IPv4 addresses.

If your Kibana node is deployed in the China (Hangzhou) region, then you can add IPv6 addresses to the Kibana whitelist. Enter IPv6 addresses and CIDR blocks in the format of `2401:b180:1000:24::5` and `2401:b180:1000::/48`, respectively. Enter `::1` to forbid all IPv6 addresses and enter `::/0` to allow all IPv6 addresses.

9.1.4 Plug-in configuration

Alibaba Cloud Kibana provides multiple plug-ins based on open-source community plug-ins. This topic introduces Alibaba Cloud Kibana plug-ins and describes how to install and remove these plug-ins.

Plug-ins

BSearch-QueryBuilder

BSearch-QueryBuilder is an advanced query plug-in, as well as a UI component.

- Easy to learn: the BSearch-QueryBuilder plug-in is a UI component, allowing you to create Elasticsearch DSL queries in a visualized manner. You can customize search conditions without coding. This saves the costs on learning complex DSL statements. It also helps developers write and verify DSL statements.
- Easy to use: all queries that you have defined are saved in Kibana, which are ready for use at anytime.

- Compact: BSearch-QueryBuilder only consumes about 14 MB of disk space. BSearch-QueryBuilder does not stay resident in the memory. This means that it will not adversely affect the performance of Kibana and Elasticsearch.
- Secure and reliable: BSearch-QueryBuilder does not rewrite, store, or forward any user data. The source code of BSearch-QueryBuilder has been verified by Alibaba Cloud security auditing.

 Note:
BSearch-QueryBuilder currently only supports Alibaba Cloud Elasticsearch instances V6.3 and V6.7. Version 5.5.3 is not supported.

Install a plug-in

 Notice:
After you purchase an Alibaba Cloud Elasticsearch instance, Elasticsearch offers you a free Kibana node with one core and 2 GB of memory. A plug-in consumes resources. Before you install a plug-in, you must upgrade the Kibana node to 2-core, 4 GB or higher. For more information, see #unique_52.

1. Log on to the Alibaba Cloud Elasticsearch console, and purchase an Elasticsearch instance.
2. ClickInstance ID/Name > Data Visualization.

3. Click Edit Configuration under Kibana.

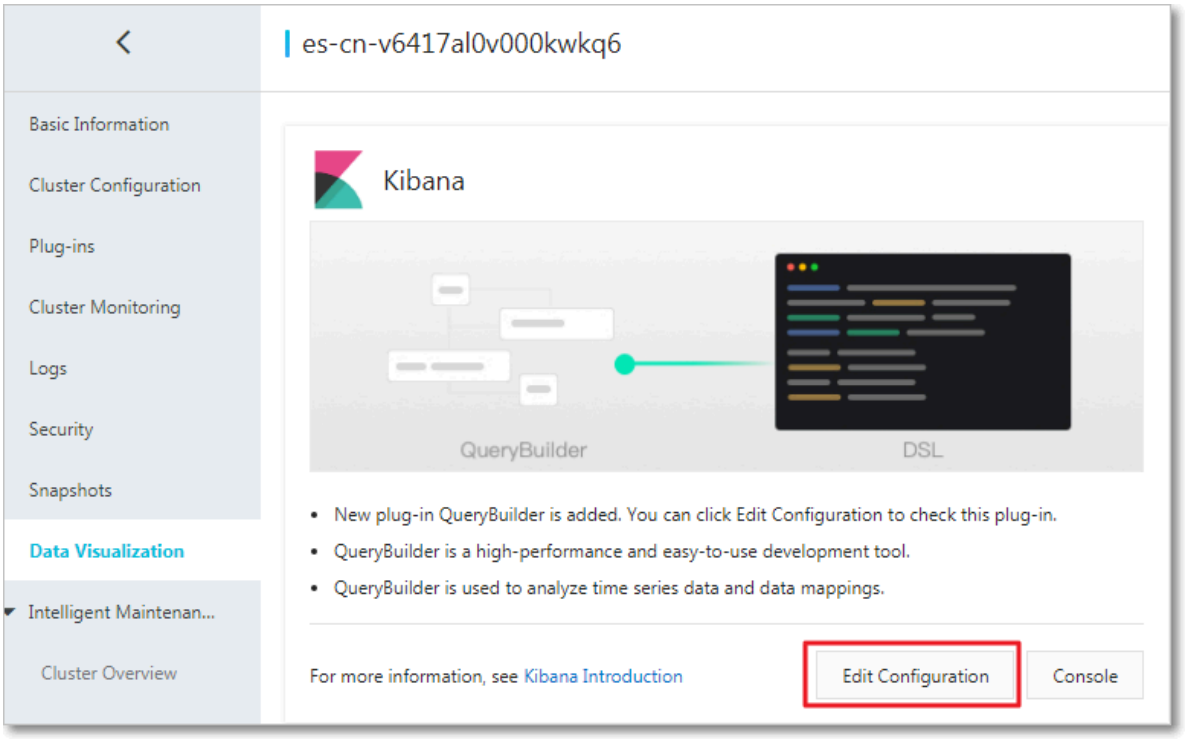

4. On the Kibana Configuration page, click Install in the Actions column in the Plug-in Configuration list.

Notice:

- After you confirm the install operation, the system will restart the Kibana node. During the restart process, Kibana cannot provide services normally. Therefore, before you confirm the operation, make sure that the restart process does not affect your operations on the Kibana console.
- If the specification of your Kibana node is lower than 2-core, 4 GB, the system prompts a notification requiring you to upgrade the instance. Follow the instructions to upgrade the Kibana node to 2-core, 4 GB or higher.

5. Confirm the operation and restart the Kibana node.

   After the Kibana node is restarted, the installation process is then completed. The plug-in will be in the Installed state.

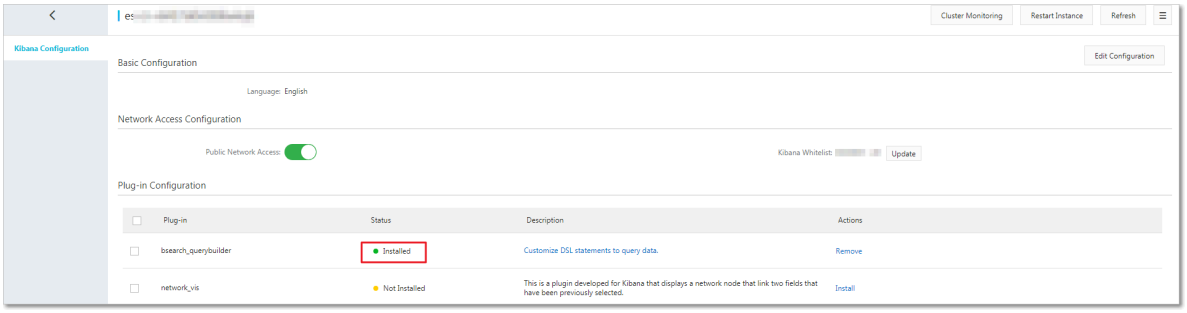

> **Note:**
> The installation process may be time-consuming.

Remove a plug-in

1. Follow the steps in Install a plug-in to go to the Kibana Configuration page, and then click Remove in the Actions column in the Plug-in Configuration list.

   > **Notice:**
   > After you confirm the remove operation, the system will restart the Kibana node. During the restart process, Kibana cannot provide services normally. Therefore, before you confirm the operation, make sure that the restart process does not affect your operations on the Kibana console.

2. Confirm the operation and restart the Kibana node.

   After the Kibana node is restarted, the remove process is then completed. The plug-in will be in the Not Installed state.

9.1.5 Use BSearch-QueryBuilder

BSearch-QueryBuilder is an advanced query plug-in, as well as a UI component. With the BSearch-QueryBuilder plug-in, you no longer need to write complex DSL statements for data query. It allows you to create complex queries in a visualized manner. This document describes how to use the BSearch-QueryBuilder plug-in to create a query.

Features

BSearch-QueryBuilder has the following features:

- Easy to learn: the BSearch-QueryBuilder plug-in is a UI component, allowing you to create Elasticsearch DSL queries in a visualized manner. You can customize search conditions without coding. This saves the costs of learning complex DSL statements. It also helps developers write and verify DSL statements.
- Easy to use: all queries that you have defined are saved in Kibana, which are ready for use at anytime.
- Compact: BSearch-QueryBuilder only consumes about 14 MB of disk space. BSearch-QueryBuilder does not stay resident in the memory. This means that it will not adversely affect the performance of Kibana and Elasticsearch.
- Secure and reliable: BSearch-QueryBuilder does not rewrite, store, or forward any user data. The source code of BSearch-QueryBuilder has been verified by Alibaba Cloud security auditing.

Background

QueryDSL is an open-source Java framework used to define SQL type-safe queries . It allows you to use API operations to send queries instead of writing statements . Currently, QueryDSL supports JPA, JDO, SQL, Java Collections, RDF, Lucene, and Hibernate Search.

Elasticsearch provides a complete JSON query DSL for you to define queries. QueryDSL provides various query expressions. Some queries can wrap other queries, such as the boolean queries. Some queries can wrap filters, such as the constant score queries. Some queries can wrap other queries and filters at the same time, such as the filtered queries. You can use any query expressions and filters supported by Elasticsearch to create complex queries and filter the returned result. DSL is only mastered by a few programmers. You may make mistakes when writing DSL statements. QueryBuilder can help users that do not have much knowledge in Elasticsearch DSL or those that want to create DSL queries efficiently.

Preparations

To use the BSearch-QueryBuilder plug-in, you must first purchase an Elasticsearch instance. The version of the instance must be 6.3 or 6.7. Version 5.5.3 is not supported.

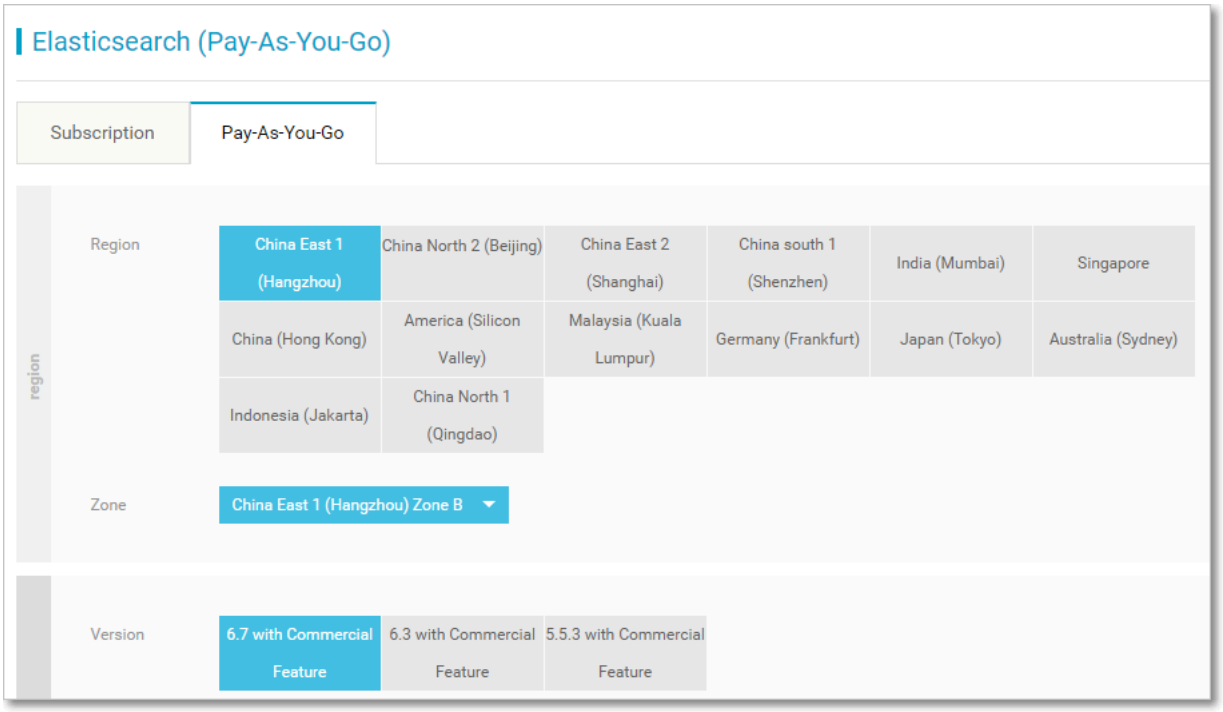

Note:

You can also use an existing instance. If the instance version does not meet the requirements, upgrade the instance.

Install the BSearch-QueryBuilder plug-in

Notice:

Before you install the BSearch-QueryBuilder plug-in, make sure that the specification of your Kibana node is 2-core, 4 GB or higher. Otherwise, #unique_52.

1. Log on to the Alibaba Cloud Elasticsearch console.
2. Click the name of the Elasticsearch instance, and then click Data Visualization in the left-side navigation pane.

3. On the Data Visualization page, click Edit Configuration under Kibana.

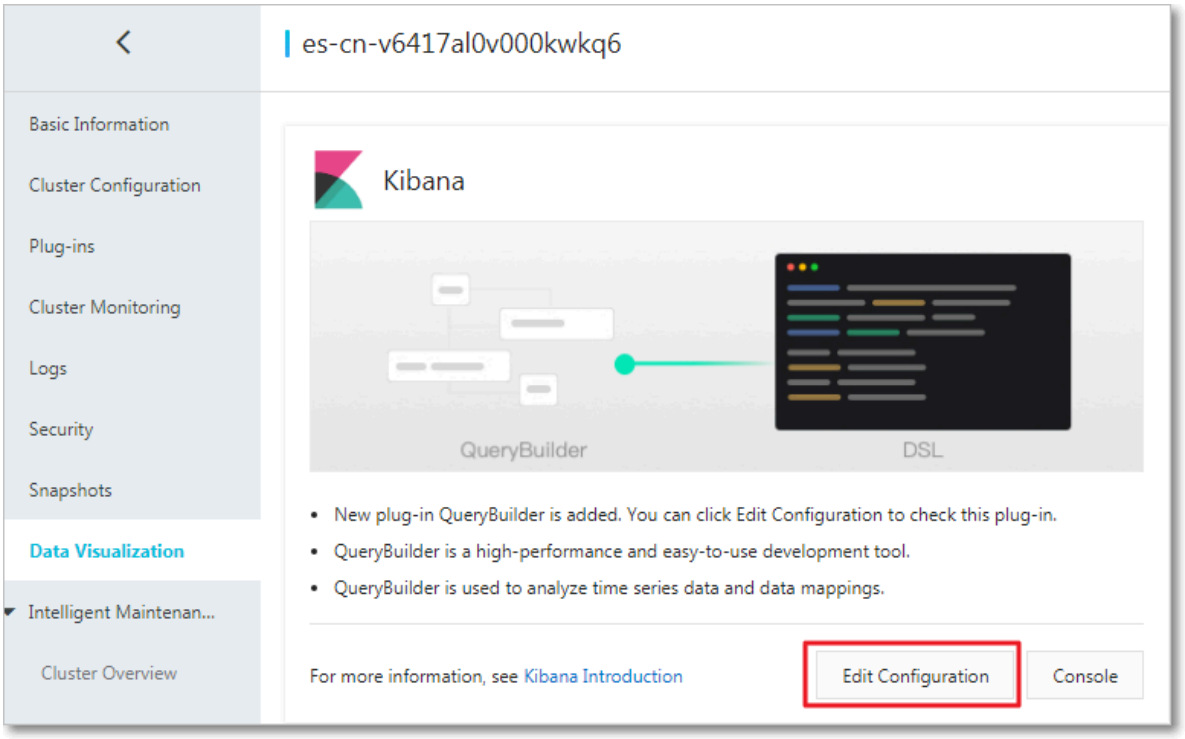

4. On the Kibana Configuration page, click Install on the right side of Bsearch_querybuilder in the Plug-in Configuration list.

> **Notice:**
> After you confirm the install operation, the system will restart the Kibana node. Therefore, before you confirm the operation, you must make sure that the restart process does not affect your operations on the Kibana console.

5. Confirm the operation and restart the Kibana node.

   After the Kibana node is restarted, the installation process is then completed. The plug-in will be in the Installed state.

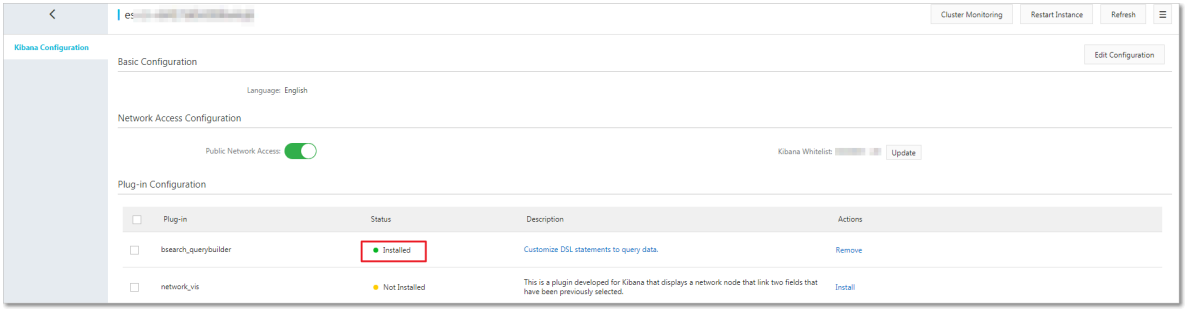

> **Note:**
> The installation process may be time-consuming.

Use the BSearch-QueryBuilder plug-in

1. Go back to the Data Visualization page, click Console under Kibana.
2. Enter the username and password, and then click LOG IN to log on to the Kibana console.

   The default username is elastic. Enter the password that you have set when purchasing the Elasticsearch instance.
3. In the Kibana console, select Discover > Query.

   > **Notice:**
   > Before querying, make sure that you have created an index pattern. To create an index pattern, in the Kibana console, click Management, find the Kibanaarea, and click Index Patterns > Create index pattern.
4. In the query area, select a search condition and filter, and click Submit.

   After you submit the query, the system shows the query result.

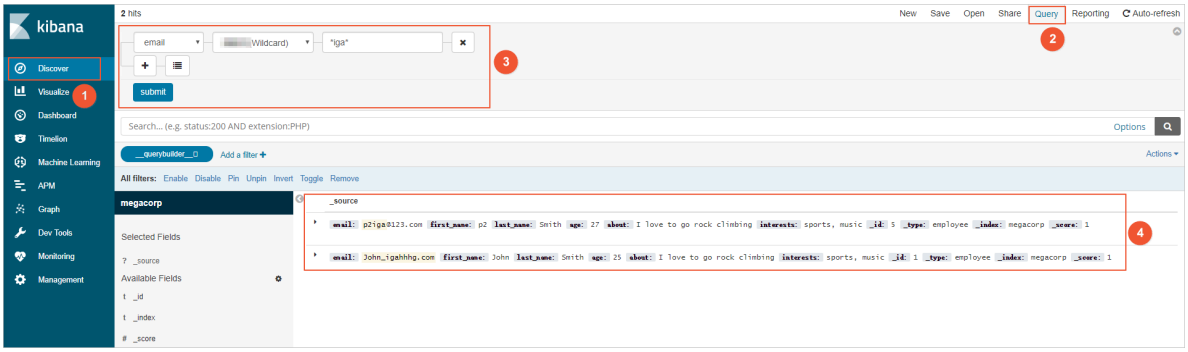

   In the query area, click the  icon to add a search condition, click the  icon to add a filter for the condition, or click the  icon to delete a search condition or filter.

   For more information, see Examples.

Examples

The BSearch-QueryBuilder plug-in allows you to create a variety of queries, such as regexp queries, boolean queries, and range queries.

- Regexp queries

  As shown in the following figure, the email condition is added for fuzzy match. The email condition matches all email addresses that contain the iga keyword.

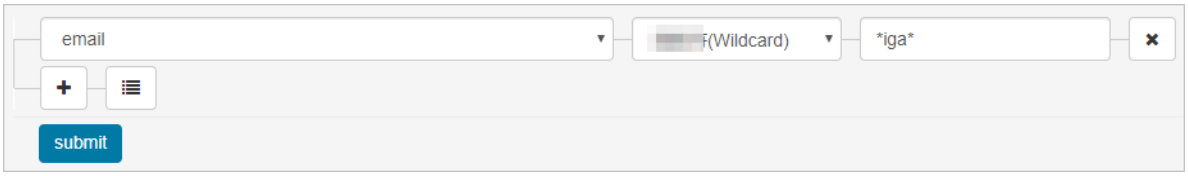

  The following figure shows the returned result:

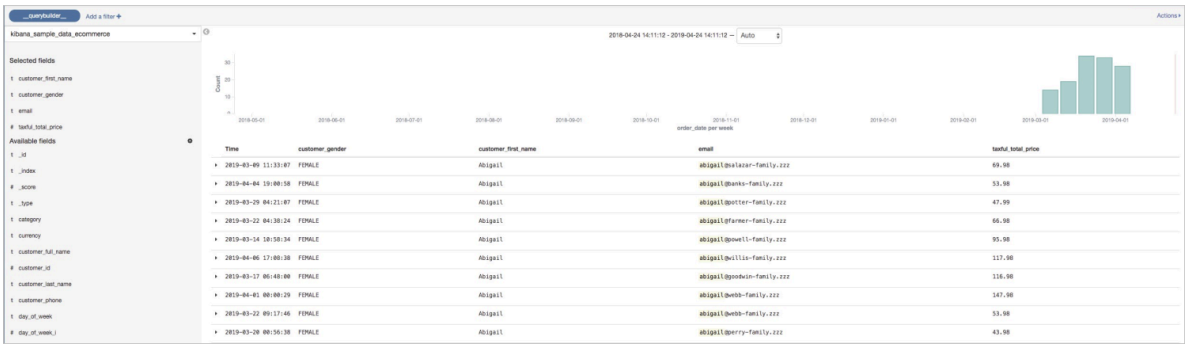

- Boolean queries

  As shown in the following figure, the index condition is set to tryme_book. An OR condition containing multiple filters is also added to filter data by type. These

type filters are set to Undergraduate teaching materials, Math, Foreign language teaching, and Undergraduate textbooks.

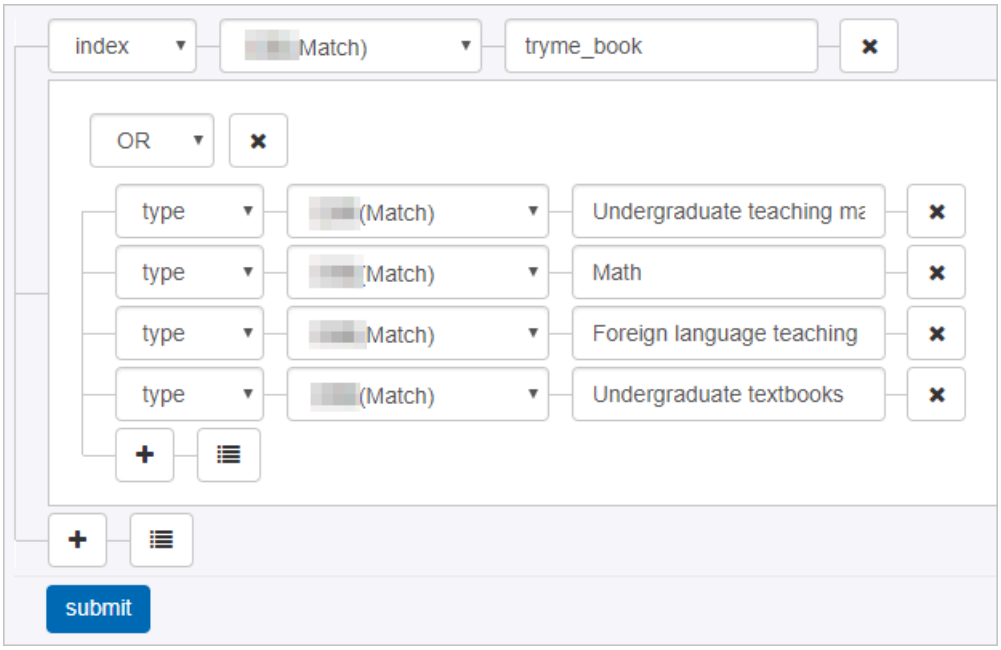

The following figure shows the returned result.

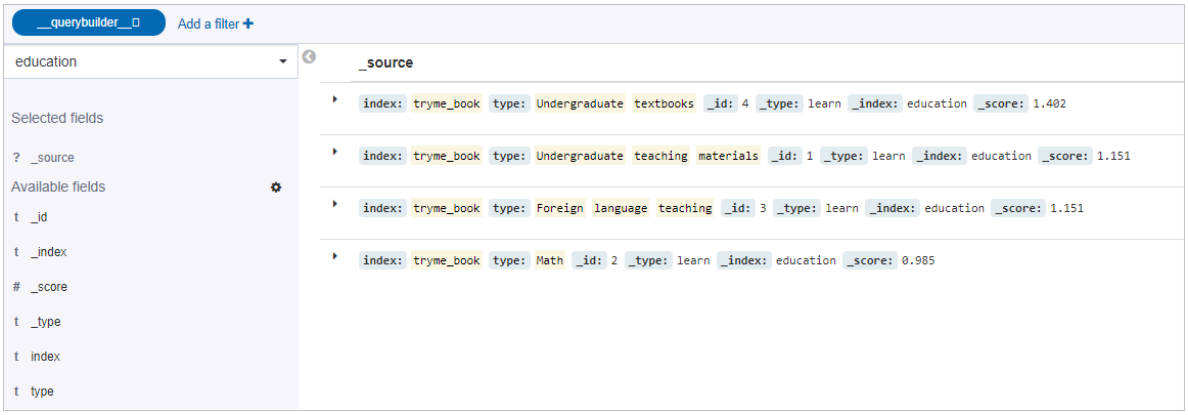

- Range queries

  Range queries allow you to search data by date. As shown in the following figure, the range condition is used to filter data based on the utc_time field. Only data entries created `in the last 240 days` are returned.

  The following figure shows the returned result.

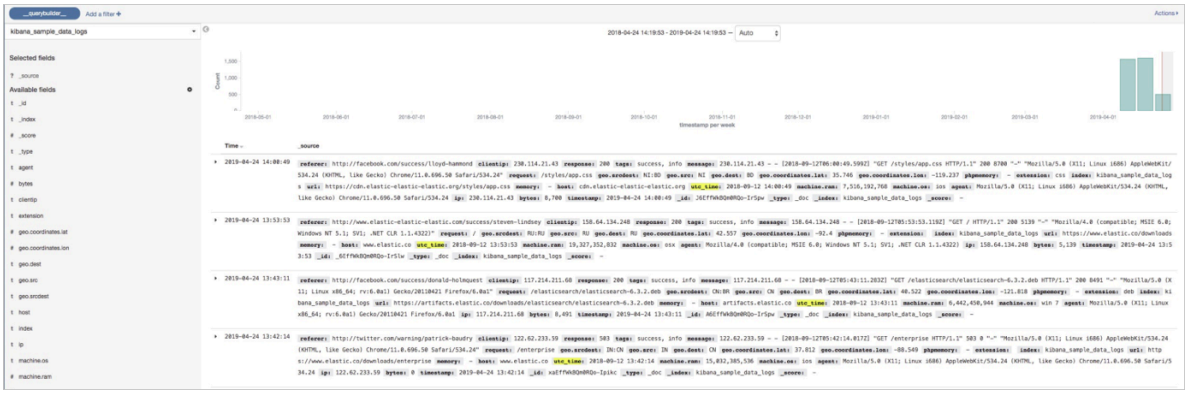

With all these search conditions and filters, you can define a complex query as follows:

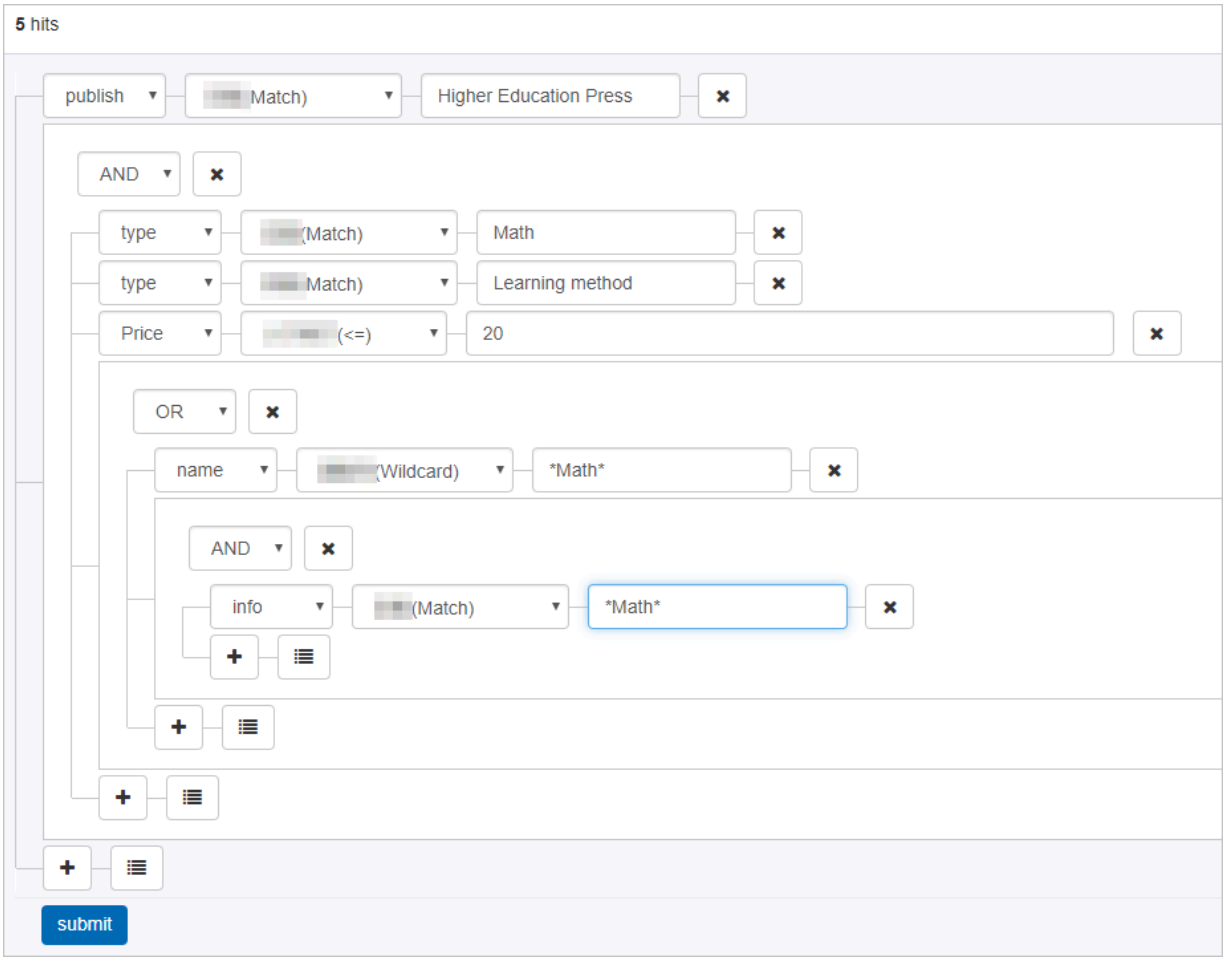

The actual DSL statement for the query is as follows:

```
"query": {
  "bool": {
    "must": [
      {
        "bool": {
          "must": [
            {
              "match_phrase": {
                "publish": "Higher Education Press"
              }
            },
            {
              "bool": {
                "must": [
                  {
                    "match_phrase": {
                      "type": "Math"
                    }
                  },
                  {
                    "match_phrase": {
                      "type": "Learning method"
                    }
                  },
                  {
                    "range": {
                      "Price": {
                        "lte": 20
                      }
                    }
                  },
                  {
                    "bool": {
                      "should": [
                        {
                          "wildcard": {
                            "name": "*Math*"
                          }
                        },
                        {
                          "bool": {
                            "must": [
                              {
                                "match_phrase": {
                                  "info": "*Math*"
                                }
                              }
                            ]
                          }
                        }
                      ]
                    }
                  }
                ]
              }
            }
          ]
        }
      }
```

As shown in the preceding examples, BSearch-QueryBuilder significantly simplifies the complexity of Elasticsearch queries.