

Alibaba Cloud

Elasticsearch
Quick Start

Document Version: 20201029

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 because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should 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 this document based on the "status quo", "being defective", and "existing functions" of its products and services. 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 take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, 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 contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, 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 this document 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 directly contact Alibaba Cloud for any errors of this document.

Document conventions








Style	Description	Example
 Danger	A danger notice 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.
 Warning	A warning notice 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 restart an instance.
 Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: If the weight is set to 0, the server no longer receives new requests.
 Note	A note indicates supplemental instructions, best practices, tips, and other content.	 Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings > Network > Set network type .
Bold	Bold formatting is used for buttons, menus, page names, and other UI elements.	Click OK .
<code>Courier font</code>	Courier font is used for commands	Run the <code>cd /d C:/window</code> command to enter the Windows system folder.
<i>Italic</i>	Italic formatting is used for parameters and variables.	<code>bae log list --instanceid</code> <i>Instance_ID</i>
[] or [a b]	This format is used for an optional value, where only one item can be selected.	<code>ipconfig [-all -t]</code>
{ } or {a b}	This format is used for a required value, where only one item can be selected.	<code>switch {active stand}</code>

Table of Contents

1.Overview	05
2.Considerations	07
3.Preparations	09
3.1. Create a VPC and a VSwitch	09
3.2. Evaluate specifications and storage capacity	09
4.Create an Elasticsearch instance	13
4.1. Create an Elasticsearch cluster	13
4.2. Parameters on the buy page	13
5.Step 2 (optional): Configure a cluster	22
6.Step 3: Access a cluster	24
7.Search services	27
7.1. Create an index	27
7.2. Create a document and insert data	28
7.3. Search for data	29
7.4. Delete documents	34
8.Step 5 (optional): Release a cluster	35

1. Overview

This topic describes how to quickly create an Alibaba Cloud Elasticsearch cluster, access the cluster, call the Elasticsearch RESTful API, and search for business data.

Background information

Make sure that you have understood the following information:

- [What is Alibaba Cloud Elasticsearch?](#)
- [Elasticsearch features](#)
- [Terms](#)
- [Compatibility matrixes](#)

Scenario

A finance service enterprise uses an online platform to manage its wealth management products. The enterprise uses conventional databases to provide search functionality for customers. The wealth management products provided by the enterprise offer satisfactory returns, and the customer base of the enterprise grows rapidly. The expansion of its business systems and increase of customer information cause the inherent issues of conventional databases to become noticeable. These issues include slow search responses, low accuracy, and the low performance of data service devices. To resolve these issues and improve customer satisfaction, the enterprise purchases the Alibaba Cloud Elasticsearch service. This topic uses this scenario to describe how to use Elasticsearch to create a cluster and search for data.

For example, the wealth management products of the enterprise are as follows:

```
{
  "products": [
    {"productName": "Daily Wealth Management for Comprehensive Health", "annual_rate": "3.2200%", "describe": "180-day wealth management product. Minimum investment of CNY 20,000. Low-risk investment. Select whether to receive push messages for returns"}
    {"productName": "Western Tongbao", "annual_rate": "3.1100%", "describe": "90-day wealth management product. Minimum investment of CNY 10,000. Daily push messages when returns are credited to your account."}
    {"productName": "Anxiang Livestock Industry", "annual_rate": "3.3500%", "describe": "270-day wealth management product. Minimum investment of CNY 40,000. Daily push messages when returns are immediately credited to your account."}
    {"productName": "Monthly 5G Device Purchase Profit", "annual_rate": "3.1200%", "describe": "90-day wealth management product. Minimum investment of CNY 12,000. Daily push messages when returns are credited to your account."}
    {"productName": "New Energy Power Wealth Management", "annual_rate": "3.0100%", "describe": "30-day wealth management product. Minimum investment of CNY 8,000. Daily push messages for returns."}
    {"productName": "Microcredit Profit", "annual_rate": "2.7500%", "describe": "3-day popular wealth management product. No service fees. Minimum investment of CNY 500. Push messages for returns."}
  ]
}
```

Procedure

Before you perform the following operations, make sure that you have read the instructions provided in [Considerations](#). The procedure is as follows:

1. Make preparations.

Create a Virtual Private Cloud (VPC) and a VSwitch, and evaluate the specifications and storage capacity for your Elasticsearch cluster. For more information, see [Create a VPC and a VSwitch](#) and [Evaluate specifications and storage capacity](#).

2. [Create an Elasticsearch cluster](#).

For more information about the parameters that are required for creating an Elasticsearch cluster, see [Parameters on the buy page](#). In this topic, Alibaba Cloud Elasticsearch V6.7 is used as an example.

3. (Optional) [Configure the cluster](#).

You can enable auto indexing, and configure plug-ins, security settings, and monitoring and alerting settings.

4. [Access the cluster](#).

You can use the Kibana console, a curl command, or a client to access the cluster.

5. [Query business data](#).

Alibaba Cloud Elasticsearch V6.7 allows you to use the Kibana console to perform the following operations: [creating an index](#), [creating a document and inserting data into the document](#), [searching for data](#), and [deleting an index](#).

6. (Optional) [Release the cluster](#).

You can release an Elasticsearch cluster that you no longer use to avoid resource waste. After the cluster is released, you are no longer charged for the cluster. In addition, the data in the cluster is deleted and can no longer be recovered.

2. Considerations

This topic describes the issues to consider before you purchase an Alibaba Cloud Elasticsearch instance.

Considerations of purchasing a subscription-based Elasticsearch instance

Discounts are offered on subscription-based Alibaba Cloud Elasticsearch instances based on the subscription duration.

You can request a refund within five days after you have purchased a subscription-based Elasticsearch instance. Refunds are not supported after five days from the date of purchase.

Considerations of purchasing an Elasticsearch instance

Currently, Alibaba Cloud Elasticsearch provides multiple **instance specifications**. The **1-core 2 GB** specification is only for testing purposes. The specification is not covered under our Service Level Agreement (SLA). Do not use Elasticsearch instances of this specification in a production environment. A production environment always requires high resource consumption. This specification provides a low configuration, which may cause service instability in a production environment.


- If you have purchased a **1-core 2 GB** instance and the instance is being used in a production environment, we recommend that you perform a [Upgrade the configuration of a cluster](#) to improve the stability and availability of your services. When you upgrade a cluster, additional fees may be charged. For more information, see [Pricing](#).
- The minimum specification for an Elasticsearch instance used in a production environment is **2-core 4 GB**.

Considerations of purchasing disks for Elasticsearch instances

Before you determine the size of the disks purchased for your Elasticsearch instance, make sure that you understand the types of data that will be stored in Elasticsearch.

Alibaba Cloud Elasticsearch is used to store the following types of data:

- User data that has been pushed to Elasticsearch.
- Index replicas. The number of index replicas is user-configurable. However, each index must have a minimum of one replica.
- Alibaba Cloud Elasticsearch cluster logs, including the operation log, access log, and slow log. The amount of storage space that Elasticsearch consumes increases with the number of queries and pushes that Elasticsearch has received. By default, Elasticsearch only retains log data generated in the last seven days. Cluster logs currently are not accessible to users.
- X-Pack monitoring indexes for troubleshooting. X-Pack is a plug-in of Alibaba Cloud Elasticsearch. Monitoring indexes include the following:
 - `.monitoring-es-6-2018.01.08`: The indexes consume a large amount of storage space. By default, Elasticsearch only keeps indexes created in the last seven days.
 - `.monitoring-kibana-6-2018.01.08`: The amount of storage space that the indexes consume increases with the number of the indexes. By default, Elasticsearch only keeps indexes created in the last seven days.
 - `.watcher-history-3-2018.01.08`: The indexes consume only a small amount of storage space. You need to manually delete these indexes if you no longer need them.

 **Notice** The smaller the disk space, the higher the disk utilization of the Elasticsearch cluster logs and X-Pack monitoring indexes.

3. Preparations

3.1. Create a VPC and a VSwitch

Alibaba Cloud Elasticsearch only supports Virtual Private Clouds (VPCs). Before you purchase an Alibaba Cloud Elasticsearch cluster, you must create a VPC and a VSwitch.

For more information about how to create a VPC and a VSwitch, see [Create an IPv4 VPC network](#).

During the creation, take note of the following items:

- If you want to use an Elastic Compute Service (ECS) instance to access your Elasticsearch cluster, the ECS instance and Elasticsearch cluster must reside in the same region. Additionally, the following requirements must be met:
 - If the ECS instance is connected to a VPC, the Elasticsearch cluster must also be connected to the VPC.
 - If the ECS instance is connected to a classic network, follow the instructions provided in [Access to an Alibaba Cloud Elasticsearch cluster from the classic network](#) to access the Elasticsearch cluster.
 - If the ECS instance and Elasticsearch cluster are connected to the same VPC but reside in different zones, create a VSwitch in each of the zones. This ensures that the network is accessible between the zones.
- The **VSwitch** drop-down list only displays the available VSwitches in a specific VPC that reside in the same zone as your Elasticsearch cluster.
 - If no VSwitches are available in the zone, create a VSwitch in the zone. For more information, see [Create a VSwitch](#).
 - The number of available VSwitch IP addresses is at least 50. Otherwise, the system displays the **Don't have enough private IPs in this switch** message.

3.2. Evaluate specifications and storage capacity

Before you use Alibaba Cloud Elasticsearch, you must evaluate the total amount of required resources, such as the disk space, node specifications, number of shards, and size of each shard. Based on test results and user feedback, Alibaba Cloud offers some common methods for the evaluation. You can purchase a cluster or upgrade the configuration of a cluster based on the evaluation results.

Precautions

- Different users may have different requirements on data schemas, query complexity, data sizes, performance, and data changes. This topic is used for reference only. We recommend that you measure the specifications and storage capacity for your Elasticsearch cluster based on actual data and business scenarios.
- If you purchase a cluster before evaluating its specifications and storage capacity, we recommend that you enable the auto scaling feature for the cluster. This feature allows you to resize disks, add nodes, and upgrade node specifications at all times based on evaluation results.

Supported disk types

This topic is suitable for Elasticsearch clusters that use standard SSDs.

Disk space evaluation

The disk space of an Elasticsearch cluster is determined by the following factors:

- Number of replica shards: Each primary shard must have at least one replica shard.
- Indexing overheads: In most cases, indexing overheads are 10% greater than those of source data. The overheads of the `_all` parameter are not included.
- Space reserved by the operating system: By default, the operating system reserves 5% of disk space for critical processes, system recovery, and disk fragments.
- Elasticsearch overheads: Elasticsearch reserves 20% of disk space for internal operations such as logging and segment merging.
- Security threshold overheads: Elasticsearch reserves at least 15% of disk space as the security threshold.

Based on these factors, the minimum required disk space is calculated by using the following formula:

$$\begin{aligned} \text{Minimum required disk space} &= \text{Volume of source data} \times (1 + \text{Number of replica shards}) \times \text{Indexing overheads} / (1 - \text{Linux reserved space}) / (1 - \text{Elasticsearch overheads}) / (1 - \text{Security threshold overheads}) \\ &= \text{Volume of source data} \times (1 + \text{Number of replica shards}) \times 1.7 \\ &= \text{Volume of source data} \times 3.4 \end{aligned}$$

Note

- We recommend that you disable the `_all` parameter unless it is required by your business.
- Indexes that have the `_all` parameter enabled incur larger overheads on disk usage. Based on test results and user feedback, we recommend that you evaluate the disk space by 1.5 times the original value. This indicates that the minimum required disk space is calculated by using the following formula:

$$\begin{aligned} \text{Minimum required disk space} &= \text{Volume of source data} \times (1 + \text{Number of replica shards}) \times 1.7 \\ &\times (1 + 0.5) = \text{Volume of source data} \times 5.1 \end{aligned}$$

- For an Elasticsearch V6.7 or V7.4 cluster of the Standard Edition, an ultra disk can offer a maximum storage space of 20 TiB for a single node. When you purchase an Elasticsearch cluster, you can specify the storage space as needed. Before you resize the disk for an Elasticsearch V6.7 cluster, make sure that the kernel of the cluster is updated. For more information about how to update a kernel, see [Update the kernel of a cluster](#). For information about how to resize a disk, see [Upgrade the configuration of a cluster](#).

Node specification evaluation

The performance of an Elasticsearch cluster is determined by the specifications of each node in the cluster. Based on test results and user feedback, we recommend that you determine node specifications by following these guidelines:

- Maximum number of nodes per cluster:

$$\text{Maximum number of nodes per cluster} = \text{Number of vCPUs per node} \times 5$$

- Maximum volume of data per node:

The maximum volume of data that a node in an Elasticsearch cluster can store depends on the scenario. Examples:

- Acceleration or aggregation on data queries: Maximum volume of data per node = Memory size per node (GiB) × 10
- Log data importing or offline analytics: Maximum volume of data per node = Memory size per node (GiB) × 50
- In general scenarios: Maximum volume of data per node = Memory size per node (GiB) × 30

The following table lists some node specifications.

Specification	Maximum number of nodes	Maximum disk space per node in query scenarios	Maximum disk space per node in logging scenarios	Maximum disk space per node in general scenarios
2 vCPUs and 4 GiB of memory	10	40 GiB	200 GiB	100 GiB
2 vCPUs and 8 GiB of memory	10	80 GiB	400 GiB	200 GiB
4 vCPUs and 16 GiB of memory	20	160 GiB	800 GiB	512 GiB
8 vCPUs and 32 GiB of memory	40	320 GiB	1.5 TiB	1 TiB
16 vCPUs and 64 GiB of memory	50	640 GiB	2 TiB	2 TiB

Shard evaluation

The number of shards and the size of each shard determine the stability and performance of an Elasticsearch cluster. You must properly plan shards for all indexes in an Elasticsearch cluster. This prevents numerous shards from affecting cluster performance when it is difficult to define business scenarios.

Note In versions earlier than Elasticsearch V7.X, one index has five primary shards and each shard has one replica shard by default. In Elasticsearch V7.X and later, one index has one primary shard and the shard has one replica shard by default.


Before you plan shards, take note of the following items:

- Volume of data stored on each index
- Whether the volume will increase
- Node specifications
- Whether you will delete or merge temporary indexes on a regular basis

Based on the preceding items, Alibaba Cloud provides the following guidelines for you to plan shards. These guidelines are for reference only.

- Before you allocate shards, evaluate the volume of data that you want to store. If the total data volume is large, write a small amount of data to reduce the workloads of your Elasticsearch cluster. In this case, configure multiple primary shards for each index and one replica shard for each primary shard. If both the total data volume and the volume of data that you want to write are small, configure one primary shard for each index and one or more replica shards for each primary shard.

- Ensure that the size of each shard is no more than 30 GB. In special cases, the size can be no more than 50 GB. If the evaluation result exceeds the limit, properly allocate shards before creating indexes and perform a reindex operation in the future. This operation ensures the normal running of your Elasticsearch cluster but is time-consuming.

 **Note** If the evaluated data volume is less than 30 GB, you can configure one primary shard and multiple replica shards for each index to implement load balancing. For example, the size of each index is 20 GB and your Elasticsearch cluster has five data nodes. In this case, you can configure one primary shard and four replica shards for each index.

- For log analytics or extremely large indexes, ensure that the size of each shard is no more than 100 GB.
- The total number of primary shards and replica shards is the same as or a multiple of the number of data nodes.

 **Note** The more shards, the more performance overheads of your Elasticsearch cluster.


- Configure a maximum of five shards for an index on a node.
- Calculate the number of shards for all indexes on a single node by using one of the following formulas:
 - For clusters with small specifications: Number of shards on a single data node = Memory size of the data node × 30
 - For clusters with large specifications: Number of shards on a single data node = Memory size of the data node × 50

Notice

When you calculate the number of shards, you must also take data volume into account. If the data volume is less than 1 TB, we recommend that you calculate the number of shards by using the formula for clusters with small specifications.

By default, the maximum number of shards on a single node in an Elasticsearch V7.X cluster is 1,000. We recommend that you do not change the maximum number. If you want to change the number of shards on a single node, you can expand node capacity and change the number before you use the cluster.

- Add at least two independent client nodes. The ratio of client nodes to data nodes must be 1:5, and the vCPU-to-memory ratio of each client node must be 1:4 or 1:8. For example, your Elasticsearch cluster contains 10 data nodes and each data node offers 8 vCPUs and 32 GiB of memory. In this case, you can configure two independent client nodes, each of which offers 8 vCPUs and 32 GiB of memory.

 **Note** After you use independent client nodes, you can perform a reduce operation on the evaluation result. In this case, if severe garbage collection (GC) occurs in the reduce stage, data nodes cannot be affected.

- If the Auto Indexing feature is enabled, enable index lifecycle management or call an Elasticsearch API operation to delete automatically created indexes.
- Delete small indexes in a timely manner to free up heap memory.

4. Create an Elasticsearch instance

4.1. Create an Elasticsearch cluster

This topic describes how to create an Alibaba Cloud Elasticsearch cluster.

Prerequisites

- An Alibaba Cloud account is created.
Click to [create an Alibaba Cloud account](#).
- A Virtual Private Cloud (VPC) and a VSwitch are created.
For more information, see [Create a VPC and a VSwitch](#).
- The specifications and storage capacity of your cluster are evaluated.
For more information, see [Evaluate specifications and storage capacity](#).

Procedure

1. Go to the [Elasticsearch cluster buy page](#).
2. On the buy page, complete cluster launch configurations. This tutorial uses an example of version 6.7. For other configurations, please customize the selection. For more information of configuration parameters, see [Parameters on the buy page](#).
3. Click **Buy Now**.
4. After the cluster is created, click **Console** to go to the **Overview** page of the Elasticsearch console.
5. In the left-side navigation pane, click **Elasticsearch Clusters** to view the created Elasticsearch cluster.

What's next

After the state of the cluster changes from **Initializing** to **Active**, you can configure or access the cluster. For more information, see [Step 2 \(optional\): Configure a cluster](#) or [Step 3: Access a cluster](#).

4.2. Parameters on the buy page

When you purchase an Alibaba Cloud Elasticsearch cluster, follow the instructions in this topic to configure the parameters on the buy page.

Billing methods

Elasticsearch provides two billing methods: **subscription** and **pay-as-you-go**. You can select a billing method as required.

- **Pay-as-you-go**: We recommend that you purchase **pay-as-you-go** Elasticsearch clusters for program development or functional tests.

You can log on to the Elasticsearch console, click **More** in the Actions column that corresponds to an Elasticsearch cluster, and select **Release** to manually release the cluster.

- **Subscription**: Discounts are offered for subscription Elasticsearch clusters based on subscription duration. However, refunds are not provided after the date of purchase.

Manual renewal and auto-renewal are supported. For more information, see [Enable auto renewal](#) and [Manually renew an Elasticsearch cluster](#). Subscription Elasticsearch clusters cannot be manually released in the console.

Basic settings

Parameter	Description
Instance Type	The value of this parameter can only be X-Pack Version .
Elasticsearch Version	<p>Valid values: 7.4, 6.8, 6.7, 6.3, 5.6, and 5.5.</p> <p>Note We recommend that you select the latest version to reduce version upgrades in the future. You may encounter differences in performance optimization and bug fixes between different versions. For more information about these differences, see the open-source Elasticsearch release notes of each version.</p>
Region	For more information, see Regions and zones .
Zone	For more information, see Regions and zones .
Number of Zones	<ul style="list-style-type: none"> • 1-AZ: This is the default deployment method. It is used to handle non-critical workloads. • 2-AZ: This deployment method implements cross-zone disaster recovery. It is used to handle production workloads. • 3-AZ: This deployment method implements high availability. It is used to handle production workloads that require high service availability. <p>Note</p> <ul style="list-style-type: none"> • You can deploy an Elasticsearch cluster across three zones only in the China (Hangzhou), China (Beijing), China (Shanghai), or China (Shenzhen) region. • When you deploy an Elasticsearch cluster across zones, you do not need to specify each zone. The system automatically selects the zones. • For more information about the precautions for deploying and using cross-zone Elasticsearch clusters, see Precautions.
Network Type	The value of this parameter can only be VPC .
VPC	<p>Select a Virtual Private Cloud (VPC) in the current region.</p> <p>Note If you want to use an Elastic Compute Service (ECS) instance to access your Elasticsearch cluster in a VPC, make sure that the ECS instance and Elasticsearch cluster reside in the same VPC.</p>

Parameter	Description
VSwitch	After you specify a VPC, all the available VSwitches in the selected zone are displayed.



Regions and zones



The following table lists the [regions and zones](#) where Elasticsearch clusters are available.



Country/District	Region	Zone
China	China (Hangzhou)	Zone I, Zone H, Zone G, Zone F, Zone E, and Zone B
	China (Beijing)	Zone H, Zone G, Zone F, Zone E, Zone D, Zone C, and Zone A
	China (Shanghai)	Zone G, Zone F, Zone E, Zone D, and Zone B
	China (Shenzhen)	Zone E, Zone D, Zone C, Zone B, and Zone A
	China (Qingdao)	Zone B and Zone C
	China (Zhangjiakou-Beijing Winter Olympics)	Zone B and Zone A
	China (Hong Kong)	Zone C and Zone B
Asia Pacific	Singapore	Zone A and Zone B
	Malaysia (Kuala Lumpur)	Zone A and Zone B
	Japan (Tokyo)	Zone A
	Australia (Sydney)	Zone A
	Indonesia (Jakarta)	Zone A
Europe & Americas	US (Silicon Valley)	Zone A and Zone B
	Germany (Frankfurt)	Zone A and Zone B
Middle East & India	India (Mumbai)	Zone A



Node settings


Parameter	Description
-----------	-------------

Parameter	Description
<p>Data Node Type</p>	<p>Data nodes store index data. You can use data nodes to add, delete, search for, modify, and aggregate data in documents. Data nodes have high CPU, memory, and I/O requirements. When you optimize the performance of an Elasticsearch cluster, you must monitor the status of the data nodes in the cluster. If the resources of the cluster are insufficient, we recommend that you add data nodes to the cluster.</p> <p>Data nodes support specifications such as 1C 2GB (for testing purposes), 2C 4GB, and 2C 8GB.</p> <div data-bbox="512 580 1385 1249" style="background-color: #e6f2ff; padding: 10px;"> <p> Notice</p> <p>The 1C 2GB specifications are designed for testing purposes. Do not use clusters with such specifications for production purposes. The service-level agreement (SLA) does not apply to these clusters. Therefore, we recommend that you do not select the specifications.</p> <p>If you have any problem when using clusters with the 1C 2GB specifications, you can use one of the following methods to resolve the problem:</p> <ul style="list-style-type: none"> • Increase the specifications. For more information, see Upgrade the configuration of a cluster. We recommend that you increase the specifications based on the vCPU-to-memory ratio of 1:4. • On the Monitoring page of the Kibana console, view the monitoring information or logs of your cluster to obtain the specific information about the problem. Then, resolve the problem based on the information. For example, if the information indicates that the indexes of your cluster occupy too much memory, you can delete some indexes. For more information about how to log on to the Kibana console of your cluster, see Log on to the Kibana console. </div>
<p>Data Nodes</p>	<p>The number of data nodes that you want to purchase. The default value of this parameter is 3. Valid values: 2 to 50.</p> <div data-bbox="512 1379 1385 1529" style="background-color: #e6f2ff; padding: 10px;"> <p> Notice You must purchase a minimum of two data nodes. However, a cluster that contains only two data nodes has a greater risk of split-brain. Therefore, exercise caution when you set this parameter.</p> </div>
	<p>You can use dedicated master nodes to perform operations on clusters. You can create or delete indexes, track nodes, and allocate shards. The stability of dedicated master nodes is important to the health of clusters. By default, every node in a cluster may be selected as a dedicated master node. Operations, such as data indexing, search, and queries, require a large number of CPU, memory, and I/O resources. To ensure the stability of a cluster, we recommend that you purchase dedicated master nodes to separate the dedicated master nodes from data nodes.</p>



Parameter	Description
<p>Dedicated Master Node</p>	<p>The default value of this parameter is No for a cluster deployed in only one zone and is Yes for a cluster deployed across zones. On the buy page or configuration upgrade page, click Yes next to Dedicated Master Node to purchase dedicated master nodes. You can also upgrade purchased dedicated master nodes on the configuration upgrade page. Your cluster is then billed based on the new specifications. For more information about the prices of the specifications, see Pricing.</p> <div data-bbox="512 450 1385 1039" style="background-color: #e1f5fe; padding: 10px; border: 1px solid #cfe2f3;"> <p> Notice</p> <ul style="list-style-type: none"> To improve the stability of your services, we recommend that you purchase dedicated master nodes. You cannot release the dedicated master nodes that you have purchased. If you have purchased 10 or more data nodes, the default value of this parameter is No. You must manually purchase dedicated master nodes. When you upgrade the configuration of an Elasticsearch cluster, if the dedicated master nodes of the cluster are free of charge, these nodes will start to incur fees after the upgrade. If dedicated master nodes are purchased and the value of the Dedicated Master Node parameter on the configuration upgrade page is Yes, the specifications of the dedicated master nodes are 1C 2GB. </div> <p>After you set this parameter to Yes, you can configure the following parameters:</p> <ul style="list-style-type: none"> Dedicated Master Nodes The value of this parameter can only be 3. Dedicated Master Node Type The default value of this parameter is 2C 8GB, which is the minimum specifications for a dedicated master node. You can set the parameter as required. <div data-bbox="539 1391 1385 1473" style="background-color: #e1f5fe; padding: 10px; border: 1px solid #cfe2f3;"> <p> Notice You cannot downgrade dedicated master nodes.</p> </div> <ul style="list-style-type: none"> Dedicated Master Node Disk Type The value of this parameter can only be Cloud SSD. Dedicated Master Node Storage Space The value of this parameter can only be 20G.

Parameter	Description
Client Node	<p>You can purchase client nodes to share the CPU overheads of data nodes. This further improves the computing performance and service stability of your Elasticsearch cluster. For CPU-intensive services, we recommend that you purchase client nodes. For example, if a number of aggregation or query operations are performed, you can use client nodes to share overheads. For more information, see Open-source Elasticsearch node types.</p> <p>The default value of this parameter is No. On the buy page or configuration upgrade page, click Yes next to Client Node to purchase client nodes. You can also upgrade purchased client nodes on the configuration upgrade page. Your cluster is then billed based on the new specifications. For more information about the prices of the specifications, see Pricing.</p> <div data-bbox="512 678 1385 763"> Notice You cannot release the client nodes that you have purchased.</div> <p>After you set this parameter to Yes, you can configure the following parameters:</p> <ul style="list-style-type: none">• Client Nodes The default value of this parameter is 2. Valid values: 2 to 25.• Client Node Type The default value of this parameter is 2C 8GB. You can set the parameter as required. <div data-bbox="539 1084 1385 1169"> Notice You cannot downgrade client nodes.</div> <ul style="list-style-type: none">• Client Node Disk Type The value of this parameter can only be Efficient cloud disk.• Client Node Storage Space The value of this parameter can only be 20G.

Parameter	Description
<p>Warm Node</p>	<p>If your business includes both of the following index types, we recommend that you purchase warm nodes to implement the hot-warm architecture. This architecture improves the computing performance and service stability of Elasticsearch. For more information, see "Hot-Warm" Architecture in Elasticsearch 5.x.</p> <ul style="list-style-type: none"> • Frequently queried or written indexes • Infrequently queried or written indexes, typically indexes of records <p>The default value of this parameter is No. On the buy page or configuration upgrade page, click Yes next to Warm Node to purchase warm nodes. You can also upgrade purchased warm nodes on the configuration upgrade page. Your cluster is then billed based on the new specifications. For more information about the prices of the specifications, see Pricing.</p> <div style="background-color: #e1f5fe; padding: 5px; border: 1px solid #cfe2f3;"> <p> Notice You cannot release the warm nodes that you have purchased.</p> </div> <p>After you purchase nodes, the system adds <code>-Enode.attr.box_type</code> to their startup parameters as follows:</p> <ul style="list-style-type: none"> • Data nodes: <code>-Enode.attr.box_type=hot</code> • Warm nodes: <code>-Enode.attr.box_type=warm</code> <p>After you set this parameter to Yes, you can configure the following parameters:</p> <ul style="list-style-type: none"> • Warm Nodes The default value of this parameter is 2. Valid values: 2 to 25. • Warm Node Type The default value of this parameter is 2C 8GB. You can set the parameter as required. <div style="background-color: #e1f5fe; padding: 5px; border: 1px solid #cfe2f3;"> <p> Notice You cannot downgrade warm nodes.</p> </div> <ul style="list-style-type: none"> • Warm Node Disk Type The value of this parameter can only be Efficient cloud disk. • Warm Node Disk Encryption Disk encryption offers the maximum data security without the need to make additional changes to your business and applications. However, disk encryption may have a small impact on the performance of your Elasticsearch cluster. Disk encryption is free of charge. Reading data from or writing data to encrypted disks does not incur any additional fees. • Warm Node Storage Space The minimum value of this parameter is 500. Unit: GiB. You can set the parameter as required.
<p>Kibana Node</p>	<p>The value of this parameter can only be Yes.</p>


Parameter	Description
Kibana Node Type	Alibaba Cloud offers you a free Kibana node with the specifications of 1C 2GB. You can choose to purchase a Kibana node with higher specifications.
Username	<p>The username of the account that is used to access an Elasticsearch cluster and log on to the Kibana console. The default value of this parameter is elastic.</p> <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p> Notice If you use the elastic account to access your Elasticsearch cluster and then reset the password of the account, it may require some time for the new password to take effect. During this period, you cannot use the elastic account to access the cluster. Therefore, we recommend that you do not use the elastic account to access an Elasticsearch cluster. You can log on to the Kibana console and create a user with the required role to access an Elasticsearch cluster.</p> </div>
Password	The password of the elastic account. You must specify this parameter.

Storage settings

Parameter	Description
Disk Type	<p>The disk type of an Elasticsearch cluster. Valid values:</p> <ul style="list-style-type: none"> • Cloud SSD: This is the default value. A standard SSD provides a maximum of 2,048 GiB of storage space. Standard SSDs are ideal for online data analysis and search that require high IOPS and fast responses. • Efficient cloud disk: Ultra disks are cost-effective and are ideal for logging and analyzing large amounts of data. <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p> Notice Ultra disks with the storage space larger than 2,560 GiB cannot be resized because these disks are designed to run in disk arrays or RAID 0.</p> </div>
Disk Encryption	<p>Disk encryption offers the maximum data security without the need to make additional changes to your business and applications. However, disk encryption may have a small impact on the performance of your Elasticsearch cluster. Disk encryption is free of charge. Reading data from or writing data to encrypted disks does not incur any additional fees.</p> <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p> Notice</p> <ul style="list-style-type: none"> • Only cloud disks can be encrypted. • You cannot enable disk encryption for purchased disks. • You cannot disable disk encryption for encrypted disks. • During a cluster configuration upgrade, you cannot change the disk encryption attribute for the disks that you have purchased. However, you can enable disk encryption when you purchase warm nodes and cloud disks. </div>

Parameter	Description
Node Storage	<p>The storage space of each node. It depends on the disk type. Unit: GiB.</p> <ul style="list-style-type: none"> • If the disk type is Cloud SSD, the maximum value of this parameter is 2048. • If the disk type is Efficient cloud disk, the maximum value of this parameter is 5120. <ul style="list-style-type: none"> ◦ If the volume of the data that you want to store exceeds 2,048 GiB, you can set the storage space to 2560, 3072, 3584, 4096, 4608, or 5120. ◦ If the storage space of the disk for a purchased Elasticsearch cluster is less than 2,048 GiB, you can resize the disk to a maximum of 2,048 GiB. If the storage space of the disk is greater than 2,048 GiB, you cannot resize the disk.

Purchase plan

Parameter	Description
Duration	<p>This parameter is available only for subscription clusters. The default value of this parameter is 1 month. Valid values: 1 month, 2 month, 3 month, 4 month, 5 month, 6 month, 7 month, 8 month, 9 month, 1 yr, 2 yr, and 3 yr.</p>
Auto Renew	<p>Only subscription clusters support the auto-renewal feature. This feature is disabled by default.</p> <ul style="list-style-type: none"> • You can select Auto Renew to enable this feature. • For purchased subscription Elasticsearch clusters, you can enable this feature in the Billing Management console. For more information, see Enable auto renewal. <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p> Notice</p> <ul style="list-style-type: none"> ◦ Monthly subscription: The auto-renewal cycle is one month. ◦ Yearly subscription: The auto-renewal cycle is one year. </div>

Node types

The following table lists the node types supported by Alibaba Cloud Elasticsearch.

Node type	Description
Data node	<p>If dedicated master nodes are purchased, data nodes are used only as data nodes. If no dedicated master nodes are purchased, data nodes are also used as dedicated master nodes.</p>
Dedicated master node	<p>Dedicated master nodes are used only as dedicated master nodes.</p>
Client node	<p>Client nodes are used only as client nodes.</p>
Warm node	<p>If no dedicated master nodes are purchased, warm nodes are used as both data nodes and dedicated master nodes. If dedicated master nodes are purchased, warm nodes are used only as data nodes.</p>

5.Step 2 (optional): Configure a cluster

After you create an Alibaba Cloud Elasticsearch cluster, you can configure the cluster to improve search efficiency and service security.

Enable auto indexing

Notice

- The **auto indexing** feature is only available for Elasticsearch clusters of versions earlier than V7.0.
- The **auto indexing** feature is designed for testing purposes. We recommend that you do not enable this feature in production environments.
- Before you upload a document to an Elasticsearch cluster, you must manually create indexes and mappings for the document. If you do not create both indexes and mappings, errors may occur. For example, if you delete the mappings for a document and then upload another document that does not have any mappings, the auto indexing feature creates mappings for both documents. The created mappings may not be suitable for your application. To avoid mapping errors, the **auto indexing** feature is disabled by default.
- After you enable the **auto indexing** feature, the system restarts your Elasticsearch cluster. Therefore, before you enable this feature, make sure that the restart does not affect your services.

If you must use the **auto indexing** feature, follow these steps to enable this feature:

1. Log on to the [Alibaba Cloud Elasticsearch console](#).
2. In the top navigation bar, select the region where your cluster resides.
3. In the left-side navigation pane, click **Elasticsearch Clusters**. On the page that appears, find the target cluster and click its ID in the **Cluster ID/Name** column.
4. In the left-side navigation pane of the page that appears, click **Cluster Configuration**.
5. On the **Cluster Configuration** page, click **Modify Configuration** on the right side of **YML Configuration**.
6. In the **YML Configuration** pane, set **Auto Indexing** to **Enable**.
7. Select **This operation will restart the cluster. Continue?** and click **OK**.
The system then restarts the Elasticsearch cluster. After the cluster is restarted, the **auto indexing** feature is enabled.

Configure plug-ins

On the Plug-ins page, you can install or remove built-in plug-ins, use the standard or rolling update method to update the IK analyzer plug-in, or upload custom plug-ins. For more information, see [Configure built-in plug-ins](#) and [Configure custom plug-ins](#).

Configure security settings

On the Security page, you can reset the password that is used to access your Elasticsearch cluster, configure a Virtual Private Cloud (VPC) whitelist, and enable the Public Network Access feature. You can also configure a public IP address whitelist, enable HTTPS, or configure cluster interconnection. For more information, see [Security](#).

Configure monitoring and alerting settings

Alibaba Cloud Elasticsearch can monitor the following metrics of Elasticsearch clusters and send SMS messages to alert users. You can customize alert thresholds for the metrics. For more information, see [Configure Elasticsearch alerts in CloudMonitor](#).

- ClusterStatus
- ClusterQueryQPS(Count / Second)
- ClusterIndexQPS(Count / Second)
- NodeCPUUtilization(%)
- NodeDiskUtilization(%)
- NodeHeapMemoryUtilization(%)
- NodeLoad_1m

6.Step 3: Access a cluster

Access an Elasticsearch cluster

Before you use an Alibaba Cloud Elasticsearch cluster to query your business data, make sure that you can access the cluster. You can access your Elasticsearch cluster by using the Kibana console, a curl command, or a client.

Use the Kibana console to access a cluster

The Kibana console is a part of the Elastic ecosystem and is seamlessly integrated into Elasticsearch. The Kibana console allows you to monitor the status of your Elasticsearch clusters and manage these clusters. This section describes how to use the Kibana console to manage your Elasticsearch cluster.

1. Log on to the [Alibaba Cloud Elasticsearch console](#).
2. In the top navigation bar, select the region where your cluster resides.
3. In the left-side navigation pane, click **Elasticsearch Clusters**. On the page that appears, find the target cluster and click its ID in the **Cluster ID/Name** column.
4. In the left-side navigation pane of the page that appears, click **Data Visualization**.
5. In the **Kibana** section, click **Console**.
6. Enter the username and password and click **Log in**.

Note

- The username is `elastic`. The password is the one that is specified when you create your Elasticsearch cluster.
- Before you log on to the Kibana console, make sure that the **Public Network Access** feature is enabled and the IP address of the host used to log on to the Kibana console is in the Kibana whitelist. By default, the **Public Network Access** feature is enabled. You can configure these settings in the [Network Access Configuration](#) section of the Kibana Configuration page.

7. In the left-side navigation pane of the Kibana console, click **Dev Tools**. On the **Console** tab, run the following command to access your Elasticsearch cluster:


```
GET /
```

If the connection is established, the following result is returned:



```
{
  "name" : "YnO5kEc",
  "cluster_name" : "es-cn-v641cjgnr000b****",
  "cluster_uuid" : "Xff3sz-GQQaLAqdOzj****",
  "version" : {
    "number" : "6.7.0",
    "build_flavor" : "default",
    "build_type" : "tar",
    "build_hash" : "8453f77",
    "build_date" : "2019-03-21T15:32:29.844721Z",
    "build_snapshot" : false,
    "lucece_version" : "7.7.0",
    "minimum_wire_compatibility_version" : "5.6.0",
    "minimum_index_compatibility_version" : "5.0.0"
  },
  "tagline" : "You Know, for Search"
}
```

Use a curl command to access a cluster

Before you use a curl command to access your Elasticsearch cluster, you must create an ECS instance that resides in the same region, zone, and Virtual Private Cloud (VPC) as the cluster. For more information, see [Create an instance by using the provided wizard](#).

 **Notice** You can also use an existing ECS instance. Make sure that this instance resides in the same region, zone, and VPC as your Elasticsearch cluster. For more information about how to use an ECS instance deployed in a classic network to access an Elasticsearch cluster deployed in a VPC, see [Access to an Alibaba Cloud Elasticsearch cluster from the classic network](#).

1. Connect to an ECS instance. For more information, see [Connect to an ECS instance](#).
2. Run the following curl command to access your Elasticsearch cluster:

 **Note** If the system displays "curl command not found", run the yum install curl command to install curl on your ECS instance.

```
curl -u <username>:<password> http://<host>:<port>
```

Parameter	Description
-----------	-------------

Parameter	Description
<username>	<p>The account that is used to access your Elasticsearch cluster. We recommend that you use an account other than the elastic account.</p> <div style="border: 1px solid #add8e6; padding: 10px; background-color: #e6f2ff;"> <p>Notice</p> <ul style="list-style-type: none"> ◦ If you use the elastic account to access your Elasticsearch cluster and then reset the password of the account, it may require some time for the new password to take effect. During this period, you cannot use the elastic account to access your Elasticsearch cluster. Therefore, we recommend that you do not use the elastic account to access your Elasticsearch cluster. ◦ If the version of your Elasticsearch cluster contains "with_X-Pack", you must specify both the username and password to access the cluster. </div>
<password>	The password that is used to access your Elasticsearch cluster. The password is the one specified when you create the cluster or initialize Kibana.
<host>	The internal endpoint of your Elasticsearch cluster. You can obtain the internal endpoint from the Basic Information page of your cluster. For more information, see View basic information of a cluster .
<port>	The port of your Elasticsearch cluster. The default port is 9200 . You can obtain the port number from the Basic Information page of the cluster. For more information, see View basic information of a cluster .

Example command:

```
curl -u elastic:es_password http://es-cn-vxxxxxxxxxxxxmedp.elasticsearch.aliyuncs.com:9200
```

If the connection is established, the result shown in the following figure is returned.



Use a client to access a cluster

You can use a PHP, Python, Java, or Go client to access your Elasticsearch cluster. For more information, see [Access Alibaba Cloud Elasticsearch by using a client](#).

7. Search services

7.1. Create an index

Before you perform operations on a document, you must create an index for the document. This topic describes how to use Alibaba Cloud Elasticsearch to create an index.


Context

The sample code provided in this topic is only suitable for Elasticsearch V6.7. For more information about the code for other versions, see [Index API](#).

Procedure

1. Log on to the Kibana console of your Elasticsearch cluster. For more information, see [Log on to the Kibana console](#).
2. In the left-side navigation pane, click **Dev Tools**.
3. On the **Console** tab, run the following command to create an index named `product_info`.

```
PUT /product_info
{
  "settings": {
    "number_of_shards": 5,
    "number_of_replicas": 1
  },
  "mappings": {
    "products": {
      "properties": {
        "productName": {"type": "text", "analyzer": "ik_smart"},
        "annual_rate": {"type": "keyword"},
        "describe": {"type": "text", "analyzer": "ik_smart"}
      }
    }
  }
}
```

 **Notice** In open-source Elasticsearch 7.0.0 and later, mapping types are removed. However, these mapping types are still supported in earlier versions. For more information, see [Removal of mapping types](#). If mapping types are used in Elasticsearch 7.0.0 and later, the system displays `"type": "mapper_parsing_exception"`.

The preceding example creates an index named `product_info`. The index is of the `products` type and contains the `productName`, `annual_rate`, and `describe` fields.

If the index is created, the following result is returned:

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

7.2. Create a document and insert data

This topic describes how to use Alibaba Cloud Elasticsearch to create a document and insert data into the document.

Procedure

1. Log on to the Kibana console of your Elasticsearch cluster. For more information, see [Log on to the Kibana console](#).
2. In the left-side navigation pane, click **Dev Tools**.
3. On the **Console** tab, run the following command to create a document and insert data. In this example, all data is inserted at a time.

```
POST /product_info/products/_bulk
{"index":{}}
{"productName":"Daily Wealth Management for Comprehensive Health","annual_rate":"3.2200%","describe":"180-day wealth management product. Minimum investment of CNY 20,000. Low-risk investment. Select whether to receive push messages for returns."}
{"index":{}}
{"productName":"Western Tongbao","annual_rate":"3.1100%","describe":"90-day wealth management product. Minimum investment of CNY 10,000. Daily push messages when returns are credited to your account."}
{"index":{}}
{"productName":"Anxiang Livestock Industry","annual_rate":"3.3500%","describe":"270-day wealth management product. Minimum investment of CNY 40,000. Daily push messages when returns are immediately credited to your account."}
{"index":{}}
{"productName":"Monthly 5G Device Purchase Profit","annual_rate":"3.1200%","describe":"90-day wealth management product. Minimum investment of CNY 12,000. Daily push messages when returns are credited to your account."}
{"index":{}}
{"productName":"New Energy Power Wealth Management","annual_rate":"3.0100%","describe":"30-day wealth management product. Minimum investment of CNY 8,000. Daily push messages for returns."}
{"index":{}}
{"productName":"Microcredit Profit","annual_rate":"2.7500%","describe":"3-day popular wealth management product. No service fees. Minimum investment of CNY 500. Push messages for returns."}
```

If `"errors" : false` is returned, data is inserted into the created document.

7.3. Search for data

This topic describes how to use Alibaba Cloud Elasticsearch to search for data. Elasticsearch supports two search methods: full-text search and search by condition.

Full-text search

1. Log on to the Kibana console of your Elasticsearch cluster. For more information, see [Log on to the Kibana console](#).
2. In the left-side navigation pane, click **Dev Tools**.
3. On the **Console** tab, run the following command to search for products with descriptions that contain `Daily push messages when returns are credited to your account`.

```
GET /product_info/products/_search
{
  "query": {
    "match": {
      "describe": "Daily push messages when returns are credited to your account"
    }
  }
}
```

If the command is executed successfully, the following result is returned:

```
{
  "took" : 21,
  "timed_out" : false,
  "_shards" : {
    "total" : 5,
    "successful" : 5,
    "skipped" : 0,
    "failed" : 0
  },
  "hits" : {
    "total" : 6,
    "max_score" : 1.3968885,
    "hits" : [
      {
        "_index" : "product_info",
        "_type" : "products",
        "_id" : "WLvWYXAB8RqL5AUxLqUU",
        "_score" : 1.3968885,
        "source" : {

```

```
  "_source" : {
    "productName" : "Western Tongbao",
    "annual_rate" : "3.1100%",
    "describe" : "90-day wealth management product. Minimum investment of CNY 10,000. Daily push
messages when returns are credited to your account."
  }
},
{
  "_index" : "product_info",
  "_type" : "products",
  "_id" : "WrvWYXAB8RqI5AUxLqUU",
  "_score" : 1.3968885,
  "_source" : {
    "productName" : "Monthly 5G Device Purchase Profit",
    "annual_rate" : "3.1200%",
    "describe" : "90-day wealth management product. Minimum investment of CNY 12,000. Daily push
messages when returns are credited to your account."
  }
},
{
  "_index" : "product_info",
  "_type" : "products",
  "_id" : "WbvWYXAB8RqI5AUxLqUU",
  "_score" : 1.3547361,
  "_source" : {
    "productName" : "Anxiang Livestock Industry",
    "annual_rate" : "3.3500%",
    "describe" : "270-day wealth management product. Minimum investment of CNY 40,000. Daily push
messages when returns are immediately credited to your account."
  }
},
{
  "_index" : "product_info",
  "_type" : "products",
  "_id" : "W7vWYXAB8RqI5AUxLqUU",
  "_score" : 1.1507283,
  "_source" : {
    "productName" : "New Energy Power Wealth Management",
    "annual rate" : "3.0100%",
    "describe" : "30-day wealth management product. Minimum investment of CNY 8,000. Daily push m
essages for returns."
  }
}
```

```
}
},
{
  "_index": "product_info",
  "_type": "products",
  "_id": "XLvWYXAB8RqL5AUxLqUU",
  "_score": 0.5753642,
  "_source": {
    "productName": "Microcredit Profit",
    "annual_rate": "2.7500%",
    "describe": "3-day popular wealth management product. No service fees. Minimum investment of CNY 500. Push messages for returns."
  }
},
{
  "_index": "product_info",
  "_type": "products",
  "_id": "V7vWYXAB8RqL5AUxLqUU",
  "_score": 0.31854028,
  "_source": {
    "productName": "Daily Wealth Management for Comprehensive Health",
    "annual_rate": "3.2200%",
    "describe": "180-day wealth management product. Minimum investment of CNY 20,000. Low-risk investment. Select whether to receive push messages for returns."
  }
}
]
}
```

Alibaba Cloud Elasticsearch allows you to create a tokenizer based on indexes to search for data. You can also sort search results by score. In the preceding result, the descriptions of the first two products contain `Daily push messages when returns are credited to your account`, while the descriptions of the last two products contain only `push messages`. The higher the ranking of a product in the search result, the higher the matching degree and score of the product.

Search by condition

Run the following command to search for products with an annualized rate of 3.0000% to 3.1300%.

```
GET /product_info/products/_search
```

```
{
  "query": {
    "range": {
      "annual_rate": {
        "gte": "3.0000%",
        "lte": "3.1300%"
      }
    }
  }
}
```

If the command is executed successfully, the following result is returned:


```
{
  "took" : 10,
  "timed_out" : false,
  "_shards" : {
    "total" : 5,
    "successful" : 5,
    "skipped" : 0,
    "failed" : 0
  },
  "hits" : {
    "total" : 2,
    "max_score" : 1.0,
    "hits" : [
      {
        "_index" : "product_info",
        "_type" : "products",
        "_id" : "WLVWYXAB8RqI5AUxLqUU",
        "_score" : 1.0,
        "_source" : {
          "productName" : "Western Tongbao",
          "annual_rate" : "3.1100%",
          "describe" : "90-day wealth management product. Minimum investment of CNY 10,000. Daily push messages when returns are credited to your account."
        }
      },
      {
        "_index" : "product_info",
        "_type" : "products",
        "_id" : "WrvWYXAB8RqI5AUxLqUU",
        "_score" : 1.0,
        "_source" : {
          "productName" : "Monthly 5G Device Purchase Profit",
          "annual_rate" : "3.1200%",
          "describe" : "90-day wealth management product. Minimum investment of CNY 12,000. Daily push messages when returns are credited to your account."
        }
      }
    ]
  }
}
```

Elasticsearch finds products that meet requirements based on the search condition and displays the products in descending order.

For more information, see [Query DSL](#).

7.4. Delete documents

This document describes how to use Alibaba Cloud Elasticsearch (ES) to delete documents.

Use the [Kibana console to access Alibaba Cloud Elasticsearch](#), and send the following request to delete documents.

- Delete a document with a specified ID:

```
DELETE /my_index/my_type/2
```

If the request is successful, the following result is returned:

```
{
  "_index" : "my_index",
  "_type" : "my_type",
  "_id" : "2",
  "_version" : 4,
  "result" : "deleted",
  "_shards" : {
    "total" : 2,
    "successful" : 2,
    "failed" : 0
  },
  "_seq_no" : 3,
  "_primary_term" : 1
}
```

- Delete all documents in a specified index:

```
DELETE /my_index
```

If the request is successful, the following result is returned:

```
{
  "acknowledged" : true
}
```

8.Step 5 (optional): Release a cluster

If you no longer require a cluster, you can release the cluster. After the cluster is released, you are no longer charged for the cluster. In addition, the data in the cluster is deleted and can no longer be recovered.

Prerequisites

Snapshots are created. For more information, see [Commands for creating snapshots and restoring data](#).

Context

You can follow the instructions provided in this topic to release only pay-as-you-go clusters.

Procedure

1. Log on to the [Alibaba Cloud Elasticsearch console](#).
2. In the top navigation bar, select the region where your cluster resides.
3. In the left-side navigation pane, click **Elasticsearch Clusters**. On the page that appears, find the target cluster. Then, in the **Actions** column, click **More** and select **Release**.



4. In the message that appears, click **OK**.