

Alibaba Cloud Elasticsearch

Quick Start

Issue: 20190815

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 <code>cd / d C :/ windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid <i>Instance_ID</i></code>
[] or [a b]	It indicates that it is an optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>

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

Contents

Legal disclaimer.....	I
Generic conventions.....	I
1 Prerequisites.....	1
2 Calculate the amount of resources.....	5
3 Activate Alibaba Cloud Elasticsearch.....	8
4 Buy page parameters.....	11
5 Elasticsearch access test.....	26
6 Configure Elasticsearch features.....	32
7 Elasticsearch clients.....	34
8 Add Elasticsearch Head to Google Chrome.....	38

1 Prerequisites

Access Elasticsearch from the Internet

You can use the public network address of an Alibaba Cloud Elasticsearch instance to access the instance from the Internet. For more information, see [Basic information](#).



Note:

- To use public network addresses, you must configure the [Public network whitelist](#). By default, Elasticsearch forbids all public network addresses.
- You cannot use the transport client to access Alibaba Cloud Elasticsearch 6.7.0 with Commercial Feature through port 9300.
- You cannot use the transport client to access Alibaba Cloud Elasticsearch 6.3.2 with Commercial Feature through port 9300.

Access Elasticsearch from a VPC network

In a VPC network, you can use an ECS instance to access the internal network address of an Elasticsearch instance . For more information, see [Basic information](#).



Note:

- You cannot use the transport client to access Alibaba Cloud Elasticsearch 6.7.0 with Commercial Feature through port 9300.
- You cannot use the transport client to access Alibaba Cloud Elasticsearch 6.3.2 with Commercial Feature through port 9300.
- Data security is not guaranteed when you access an Alibaba Cloud Elasticsearch instance from the Internet. For data security, we recommend that you purchase an ECS instance that meets the requirements described in the following [Purchase instructions](#). You can then use the ECS instance to access the internal network address of the Elasticsearch instance from a VPC network.

Purchase instructions

- Before you purchase an Alibaba Cloud Elasticsearch instance, make sure that you have created a VPC network and VSwitch in the same region as the Elasticsearch instance.

- If you choose to use an Alibaba Cloud ECS instance to access an Alibaba Cloud Elasticsearch instance, make sure that the Elasticsearch instance and ECS instance are in the same region.
 - If your Alibaba Cloud ECS instance is connected to a VPC network, make sure that the ECS instance and your Alibaba Cloud Elasticsearch instance are connected to the same VPC network.
 - If your Alibaba Cloud ECS instance is connected to the classic network, you must make sure that the ECS instance meets the requirements described in the Supported VPC CIDR blocks topic in [Classic network errors](#).
- You can only select VSwitches in the same zone, region, and VPC network as the Elasticsearch instance.
 - If no VSwitch is available in the supported zones, you must manually create a VSwitch in one of the supported zones.
 - Make sure that the number of VSwitch IP addresses is no less than 200. Otherwise, the system displays a message indicating insufficient private IP addresses.

**Note:**

- If your Alibaba Cloud ECS instance and Elasticsearch instance are deployed in the same VPC network, same region, but different zones, you must create a VSwitch in the zone where the ECS instance is deployed to ensure that the ECS instance can access your Elasticsearch instance.
- New Elasticsearch users can choose to purchase free trial Elasticsearch instances for testing.

Guidelines for purchasing Subscription-based Alibaba Cloud Elasticsearch instances

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

Subscription-based Alibaba Cloud Elasticsearch instances purchased on the Alibaba Cloud International site do not support conditional refund or unconditional refund within five days. To cancel the subscription for an Elasticsearch instance, you need to sign in with your Alibaba Cloud account, choose Console > Billing > Billing Management, and click the switch to manually disable auto renewal. Before the end of the current billing cycle, you can still use the Elasticsearch instance. However, the

subscription fee is not refunded. Alibaba Cloud will stop renewing your instance in the next billing cycle.

Guidelines for Elasticsearch instance specifications

Alibaba Cloud has provided multiple types of Elasticsearch instances. The 1-Core 2 GB instances are only for testing purposes. These instances are not suitable for production. A 1-Core 2 GB instance is configured with the minimum specification, which consumes a large amount of resources and may cause service instability in production. In addition, the 1-Core 2 GB instance is not covered by the Elasticsearch Service Level Agreement (SLA). For more information, see [Pricing](#).

- If you are using a 1-Core 2 GB Elasticsearch instance in production, we recommend that you upgrade the instance to ensure the stability and availability of your service.
- The minimum specification for Elasticsearch instances used in production is 2-Core 4 GB.

Guidelines for purchasing Elasticsearch instance disks

Before determining the disk size of an Alibaba Cloud Elasticsearch node, learn about what type of data is saved in the disk. Determine the disk size based on your business requirements.

We recommend that you preserve disk space for storing the system logs when selecting a disk. System logs include the operation log and monitoring log.

Storage space usage



Note:

The storage space of a disk determines the amount of Elasticsearch cluster logs and X-Pack monitoring indexes that you can store.

- Stores user data that has been pushed to Elasticsearch.
- Stores Elasticsearch replicas. The number of replicas is user-configurable. However, Elasticsearch will retain a minimum of one replica.
- Stores Elasticsearch cluster logs. The amount of storage space that Elasticsearch consumes increases with the number of queries and pushes that Elasticsearch

has received. By default, Elasticsearch only keeps cluster logs for seven days after creation. Elasticsearch cluster logging is not currently available.

- Stores the operation Log.
- Stores the access log.
- Stores slow logs.
- Stores X-Pack monitoring indexes for troubleshooting exceptions. X-Pack is an Elasticsearch component. Monitoring indexes include the following:
 - `.monitoring-es-6-2018.01.08`. The indexes consume a large amount of storage space. Elasticsearch creates only one index each day and keeps indexes for seven days after creation.
 - `.monitoring-kibana-6-2018.01.08`. The amount of storage space that the indexes consume increases with the number of the indexes. Elasticsearch creates only one index each day and keeps indexes for seven days after creation.
 - `.watcher-history-3-2018.01.08`. The indexes consume only a small amount of storage space. Elasticsearch creates only one index each day and keeps all indexes that have been created. You must manually delete the indexes that you no longer need.



Note:

When an Elasticsearch instance has a high disk usage, such as a disk usage higher than 80%, the Elasticsearch cluster health status changes to yellow or red, and the Elasticsearch instance cannot be restarted. Therefore, before you restart an Elasticsearch instance, make sure that the health status of the instance is green. For more information, see [Instance management](#).

2 Calculate the amount of resources

Before you purchase an Alibaba Cloud Elasticsearch instance, you must estimate and calculate the amount of Elasticsearch resources that you need. Based on testing results and user feedback, Alibaba Cloud offers some common methods to estimate and calculate the amount of Elasticsearch resources. These methods are for reference only.

Supported disk types

This topic applies to Alibaba Cloud Elasticsearch instances that use SSD disks.

Calculate disk capacity

The capacity of disks that are used by Alibaba Cloud Elasticsearch instances is determined based on the following factors:

- The number of replicas. You must store a minimum of one replica.
- The index overheads are generally 10% larger than those of the source data. The index overheads of the `_all` field are not calculated.
- Reserved space of the operating system. The operating system reserves 5% of the disk space for critical processes, system recovery, and disk fragments by default.
- Elasticsearch overheads. Elasticsearch reserves 20% of the disk space for internal operations, such as segment merging and logging.
- A minimum of 15% of the disk space must be reserved as the security threshold.

Minimum disk space = Size of source data x 3.4

```
Total disk space = Size of source data × ( 1 + Number
of replicas ) × ( 1 + Indexing overheads ) / ( 1 - Linux
reserved space ) / ( 1 - Elasticsearch overheads ) / ( 1 -
Threshold overheads )
= Size of source data × ( 1 + Number of replicas ) × 1
. 7
= Size of source data × 3 . 4
```

We recommend that you do not set the `_all` field unless it is required by your business.

Indexes that have this field enabled generate larger overheads. Based on testing results and user feedback, we recommend that you add extra 50% of the estimated space to the final amount of disk space.

```
Total disk space = Size of source data × ( 1 + Number
of replicas ) × 1 . 7 × ( 1 + 0 . 5 )
```

$$= \text{Size of source data} \times 5.1$$

Determine cluster specifications

The performance of an Alibaba Cloud Elasticsearch cluster is determined by the specifications of the individual Elasticsearch instances in the cluster. Based on testing results and user feedback, we recommend that you determine the specification of your cluster as follows:

$$\text{Maximum number of nodes in a cluster} = \frac{\text{Number of vCPUs per node} \times 5}{5.1}$$

The maximum workload of an Elasticsearch node varies depending on different scenarios. Examples:

- Data acceleration and query aggregation

$$\text{Maximum workload per node} = \frac{\text{Memory size per node (GB)}}{40} \times 10$$

- Logging and analysis

$$\text{Maximum workload per node} = \frac{\text{Memory size per node (GB)}}{40} \times 50$$

- Common scenarios

$$\text{Maximum workload per node} = \frac{\text{Memory size per node (GB)}}{40} \times 30$$

Cluster specifications

Specification	Maximum number of cluster nodes	Maximum memory size per node (query)	Maximum memory size per node (log)	Maximum memory size per node (common)
2C 4 GB	10	40 GB	200 GB	100 GB
2C 8 GB	10	80 GB	400 GB	200 GB
4C 16GB	20	160 GB	800 GB	512 GB
8C 32 GB	40	320 GB	1.5 TB	1 TB
16C 64 GB	50	640 GB	2 TB	2 TB

Calculate shard size

Both the number of shards and the size of each shard contribute to the stability and performance of an Alibaba Cloud Elasticsearch cluster. An Elasticsearch cluster creates five shards for each index by default. However, you may need to modify the default settings to improve performance in most cases.

- For small Elasticsearch nodes, the size of each shard must be smaller than or equal to 30 GB. For large Elasticsearch nodes, the size of each shard must be smaller than or equal to 50 GB.
- For log analysis or extremely large indexes, the size of each shard must be smaller than or equal to 100 GB.
- The number of shards, including replicas, must be equal to the number of nodes or equal to an integer multiple of the number of nodes.
- We recommend that you create a maximum of five shards for one index on a node.



Note:

The calculation methods described in this topic may vary depending on the data structure, query requirements, data size, performance, and data changes. You must determine the number of shards and calculate the size of the shards based on your actual situation.

We recommend that you run tests and determine a best practice that satisfies your requirements based on the actual data and usage scenario. You can use the elastic scaling feature of Alibaba Cloud Elasticsearch to make a plan based on this topic. You can increase disk size, add more nodes, and upgrade node specifications based on your actual needs.

3 Activate Alibaba Cloud Elasticsearch

This topic describes how to activate Alibaba Cloud Elasticsearch.

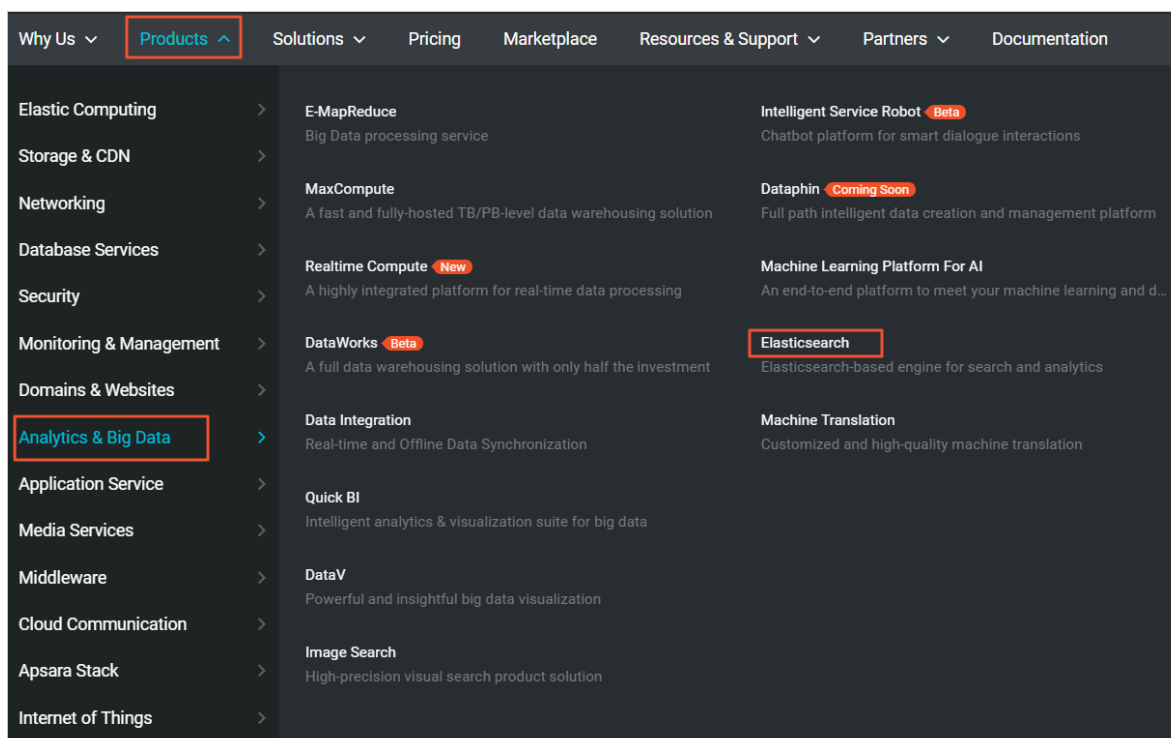
Prerequisites

Before you purchase Alibaba Cloud Elasticsearch, make sure that you have completed the following tasks:

- You have created an Alibaba Cloud account. To create an Alibaba Cloud account, click [Create an Alibaba Cloud account](#).
- You have created a VPC network and VSwitch. For more information, see [Create a VPC and a VSwitch](#).

Procedure

1. Log on to the [Alibaba Cloud International site](#).
2. Select Products > Analytics & Big Data, and then click Elasticsearch.



3. Click Buy Now, set the parameters on the buy page, and then click Buy Now.

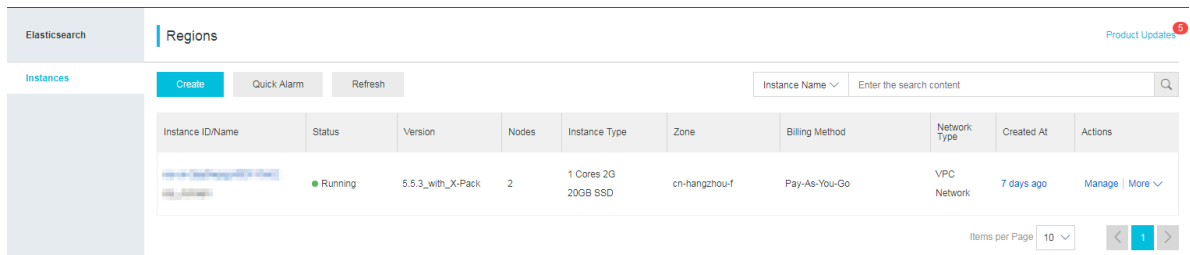
For more information about the parameters, see [Buy page parameters](#).

4. On the Confirm Order page, confirm the instance information, select the check box to agree with the terms of service, and then click Activate.

Confirm Order

5. After the service is activated, click Console to log on to the Alibaba Cloud Elasticsearch console.

6. You can then view the purchased Elasticsearch instance in the Alibaba Cloud Elasticsearch console. Wait for the instance to be activated.



Instance ID/Name	Status	Version	Nodes	Instance Type	Zone	Billing Method	Network Type	Created At	Actions
elasticsearch-xxxxxx	Running	5.5.3_with_X-Pack	2	1 Cores 2G 20GB SSD	cn-hangzhou-f	Pay-As-You-Go	VPC Network	7 days ago	Manage More


When the status of the instance changes to Active, the instance is activated. You can then use the Elasticsearch instance to perform [Elasticsearch access testing](#), [configure Elasticsearch features](#), import data, and query data.



Notice:

Before you import data to the activated Elasticsearch instance, you must manually create indexes and mappings. For example, if you import data to the instance while the system is automatically creating mappings, the created mappings may not be able to meet your requirements and an error may occur.

To avoid this issue, Elasticsearch disables the auto indexing feature by default. Therefore, you must create indexes and mappings before you import data to the instance.



If you must use the auto indexing feature, go to the [Cluster Configuration](#) page in the Elasticsearch console, enable the auto indexing feature under YML Configuration, and then restart the instance to apply the change.

4 Buy page parameters

This topic describes the parameters on the Alibaba Cloud Elasticsearch buy page.

Billing method

Alibaba Cloud Elasticsearch supports both the subscription and pay-as-you-go billing methods.

- **Subscription:** Currently, Alibaba Cloud Elasticsearch offers promotional discount based on the subscription duration: monthly or annual. Alibaba Cloud Elasticsearch subscription instances purchased on the Alibaba Cloud International site do not support conditional refunds or unconditional refunds within five days. If you want to cancel your subscription, make sure that all data on the subscription instance has been backed up. Sign in on the Alibaba Cloud International site, and select Console > Billing Management > Renew, and manually disable the auto-renewal feature. Before the end of the current billing cycle, you can still use the Elasticsearch instance. However, the subscription fee is not refunded. Alibaba Cloud will stop renewing your instance in the next billing cycle.

| Elasticsearch (Subscription)



- If the auto-renewal feature is disabled, you must log on to the Elasticsearch console and manually renew the overdue instances. For more information, see [Overdue payments](#).
- Elasticsearch instances cannot be manually released in the console.
- Auto-renewal is supported and disabled by default. For more information, see the [Auto-renewal description](#) in this topic.

- **Pay-as-you-go:** We recommend that you purchase Alibaba Cloud Elasticsearch pay-as-you-go instances for testing purposes at the development and testing stages.


Elasticsearch (Pay-As-You-Go)



Subscription



Pay-As-You-Go

- Auto-renewal is supported. For more information, see [Overdue payments](#).
- You can log on to the Elasticsearch console, click More, and then select Release to manually release a pay-as-you-go instance.

Parameters

Parameter	Description
Region	For more information, see Regions and zones .
Zone	For more information, see Regions and zones .
Resource Group	We recommend that you select the Default Resource Group.
Version	Elasticsearch 6.7, 6.3, and 5.5.3 are supported. Elasticsearch 6.3.2 Release Notes . Elasticsearch 5.5.3 Release Notes .
Network Type	Currently, only VPC is supported.
Number of Available Zones	1-AZ, 2-AZ, and 3-AZ are supported. For more information, see Number of zones .
VPC	Select a VPC network in the corresponding region.  Note: If you want to use an Alibaba Cloud ECS instance to access your Elasticsearch instance, then the ECS instance must be connected to the same VPC network as your Elasticsearch instance.
VSwitch	After you have specified a VPC network, only VSwitches in the zones supported by the region of your Elasticsearch instance are displayed.
Specification Family	For more information, see Instance types and specifications .
Instance Type	For more information, see Instance types and specifications .

Parameter	Description
Nodes	<p>The number of data nodes that you want to purchase.</p> <div style="background-color: #f0f0f0; padding: 10px;">  Notice: <ul style="list-style-type: none"> · You must purchase a minimum of two data nodes. Use caution because a cluster that contains only two data nodes may have the split-brain syndrome. · The default number of data nodes is 3. Valid values: from 2 to 50. </div>
Dedicated Master Node	For more information, see Dedicated master nodes .
Client Node	For more information, see Client nodes .
Warm Node	For more information, see Warm nodes .
Disk Type	<p>SSD Cloud Disk and Ultra Disk are supported.</p> <ul style="list-style-type: none"> · SSD Cloud Disk (default): Supports up to 2 TiB of storage space. SSD cloud disks are suitable for online data analysis and search that require high throughput and fast response. · Ultra Disk: Supports up to 5 TiB of storage space. Ultra disks are cost-effective and can be used in scenarios such as logging and analyzing large amounts of data. <div style="background-color: #f0f0f0; padding: 10px;">  Notice: Ultra disks larger than 2.5 TiB cannot be expanded because they run in disk arrays or RAID 0. </div>
Node Storage	<p>The storage space per data node varies depending on the type of the disk.</p> <ul style="list-style-type: none"> · SSD Cloud Disk: Supports up to 2,048 GiB (2 TiB) of storage space. Unit: GiB. · Ultra Disk: <ul style="list-style-type: none"> - An ultra disk can provide up to 5,120 GiB (5 TiB) of storage space. 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. - For purchased Alibaba Cloud Elasticsearch instances, if the disks are smaller than 2 TiB, you can only expand them up to 2 TiB. If the disks are larger than 2 TiB, you cannot expand the disk size.

Parameter	Description
Username	<p>By default, the elastic account is used to log on to Alibaba Cloud Elasticsearch instances and the Kibana console.</p> <div style="background-color: #f0f0f0; padding: 10px;">  Notice: If you use the elastic account to access your Alibaba Cloud Elasticsearch service, after you reset the password of the elastic account, it may take a period of time for the new password to take effect. When Elasticsearch is applying the new password, your services will be affected. We recommend that you do not use the elastic account to access your Elasticsearch service. We recommend that you log on to the Kibana console and create a user account with the required role to access Elasticsearch. </div>
Password	The password of the elastic account. You must enter a password.
Duration	Only subscription instances support this parameter. The default subscription duration is one month. You can slide to select a subscription duration to meet your business needs. Supported subscription duration: from 1 to 9 months or from 1 to 3 years.
Auto Renew	<ul style="list-style-type: none"> · To enable this feature, select Auto Renew on the buy page. · For purchased Alibaba Cloud Elasticsearch subscription instances, you can enable the auto-renewal feature for these instances in Renew. <div style="background-color: #f0f0f0; padding: 10px;">  Notice: <ul style="list-style-type: none"> - Monthly subscription: The auto-renewal cycle is one month. - Annual subscription: The auto-renewal cycle is one year · </div>

Regions and zones

Alibaba Cloud Elasticsearch supports the following [Regions and zones](#):

- China (Hangzhou): Zone B, Zone F, Zone G, Zone H, Zone I, and Zone E.
- China (Beijing): Zone E, Zone A, Zone C, and Zone D.
- China (Shanghai): Zone B and Zone D.
- China (Shenzhen): Zone C.

- **India (Mumbai): Zone A.**
- **Singapore: Zone A and Zone B.**
- **China(Hong Kong): Zone B and Zone C.**
- **US (Silicon Valley): Zone A and Zone B.**
- **Malaysia (Kuala Lumpur): Zone A and Zone B.**
- **Germany (Frankfurt): Zone A and Zone B.**
- **Japan (Tokyo): Zone A.**
- **Australia (Sydney): Zone A.**
- **Indonesia (Jakarta): Zone A.**
- **China (Qingdao): Zone B and Zone C.**

Number of zones

You can specify the number of zones when you purchase an Alibaba Cloud Elasticsearch instance. If you choose two or three zones, the Elasticsearch instance is deployed across zones. Currently, you can deploy an Elasticsearch instance across zones in the China (Hangzhou), China (Beijing), China (Shanghai), or China (Shenzhen) region. You can use the following methods to deploy an Elasticsearch instance:

- **One zone:** the default deployment for handling common tasks.
- **Across two zones:** a disaster recovery deployment for production.
- **Across three zones:** a deployment for implementing high availability. We recommend that you use this deployment for production where high availability is required.

If you choose to deploy an Elasticsearch instance across zones, you do not need to manually specify the zones. The system will select the zones for you. When you purchase and use an Alibaba Cloud Elasticsearch instance deployed across zones, follow these guidelines:

- **Nodes**
 - You must purchase dedicated master nodes.
 - The number of data nodes, warm nodes, or client nodes must be a multiple of the number of the zones. For more information about zones, see [Regions and zones](#).

- Index replicas
 - When an Elasticsearch instance is deployed across two zones and one zone becomes unavailable, the other zone is used to provide services. Therefore, the number of index replicas must be greater than or equal to 1.



Note:

By default, an Elasticsearch instance has five primary shards and one replica. If you do not have specific requirements on the throughput of the instance, you can use the default setting.

- When an Elasticsearch instance is deployed across three zones and one or two of the zones become unavailable, the remaining zones are used to provide services. Therefore, the number of index replicas must be greater than or equal to 2.



Note:

By default, an Elasticsearch instance has five primary shards and one replica. Therefore, you need to manually modify the index template to adjust the default number of replicas. For more information, see [Index templates](#).

The following sample code shows how to modify the index template to set the number of replicas to 2.

```
PUT _template / template_1
{
  "template " : "*",
  "settings " : {
    "number_of_ replicas " : 2
  }
}
```




```
}

```

- Cluster settings

The system will automatically configure cluster settings for shard allocation awareness for Alibaba Cloud Elasticsearch instances deployed across zones. For more information, see [Shard allocation awareness](#).

- If an Alibaba Cloud Elasticsearch instance is deployed in the cn-hangzhou-f and cn-hangzhou-g zones, the cluster settings are as follows:

Parameter	Description	Value
<pre>cluster.routing allocation awareness.attributes</pre>	<p> Notice: Do not modify this parameter by using the Elasticsearch API. Otherwise, an exception may occur.</p> <p>This parameter specifies the node attributes that are specified as shard allocation awareness attributes for Alibaba Cloud Elasticsearch. The Alibaba Cloud Elasticsearch instance identifies the zone of a node by adding <code>node.attr.zone_id</code> to the startup configuration of the node. Therefore, this parameter is set to <code>zone_id</code>.</p>	"zone_id"

Parameter	Description	Value
<code>cluster.routing.allocation.awareness.force.zone_id.values</code>	<p>Enables or disables forced awareness for replica allocation. Forced awareness prevents a zone from being overloaded when other zones fail. For example, an index has one primary shard and three replica shards. The Elasticsearch instance is deployed across the cn-hangzhou-f and cn-hangzhou-g zones. According to the shard allocation awareness policy, two shards are allocated to the zone cn-hangzhou-f and two shards are allocated to the zone cn-hangzhou-g. If the <code>cluster.routing.allocation.awareness.force.zone_id.values</code> parameter is set, when the zone cn-hangzhou-f fails, forced awareness can prevent the system from reallocating the shards of the zone cn-hangzhou-f to the zone cn-hangzhou-g. Otherwise, the zone cn-hangzhou-g has to host all the shards, which may be overloaded.</p> <div data-bbox="686 1848 1053 2096" style="background-color: #f0f0f0; padding: 5px;">  Note: By default, forced awareness is disabled. You can change the setting as needed. </div>	<code>["cn-hangzhou-f", "cn-hangzhou-g"]</code>

- Alibaba Cloud Elasticsearch instances deployed across zones will add `- Enode . attr . zone_id` to the startup configuration of the nodes. For example, if a node is deployed in the zone `cn-hangzhou-g`, then `- Enode . attr . zone_id = cn - hangzhou - g` is added to the startup configuration of the node.
- Switch over and back

Alibaba Cloud Elasticsearch instances deployed across zones support the switch over and switch back operations.

- Switch over



Notice:

To make sure that the Elasticsearch instance can read and write data normally after the switchover is complete, we recommend that the Elasticsearch instance has at least one replica.

When one or more nodes of the Elasticsearch instance is a zone fail, you can switch the network traffic to other zones. On the Basic Information page of the Elasticsearch console, click Node Visualization, move your mouse pointer to the zone where the nodes fail, and click Switch Over. Follow the instructions to complete the switchover operation.



Note:

- The status of the specified zone changes from Enabled to Disabled.
- Network traffic sent from clients to the zone is then forwarded to other enabled zones and the nodes in this zone are removed from the cluster.
- To ensure that the indexes can be read and written normally and the computing resources are sufficient, Alibaba Cloud Elasticsearch will add new nodes to other enabled zones. These nodes may include dedicated master nodes, client nodes, data nodes, and Kibana node.


Before you perform the switch over operation, if your index has one replica, then the status of the cluster displays yellow after the switchover process is complete. After the switchover process is complete, you can reference the following example to set the cluster parameters so that the shards in the

disabled zone can be allocated to other zones. After the shards are allocated, the status of the cluster displays green.

```
PUT /_cluster/settings
{
  "persistent": {
    "cluster.routing.allocation.awareness.force.zone_id.values": {"0": null, "1": null, "2": null}
  }
}
```

- Switch back

On the Basic Information page of the Elasticsearch console, click Data Visualization, move your mouse pointer to the zone that you want to recover, and click Switch Back. Follow the instructions to complete the switchback operation.

 **Note:**

- The status of the zone changes from Disabled to Enabled.
- Network traffic sent from clients is then forwarded to the recovered zone and nodes in this zone are added to the cluster again.
- During the switchback process, Alibaba Cloud Elasticsearch will remove the nodes that are added during the switchover process. These nodes may include dedicated master nodes, client nodes, and data nodes. When Elasticsearch removes data nodes, it will migrate the data on the data nodes to other data nodes.

Instance types and specifications

The following table shows the supported instance types and relevant specifications. For more information about prices of different instance types, see [Pricing](#).

Instance type	Specification
Cloud disk type	1-core 2 GB (testing only), 2-core 4 GB, 2-core 8 GB, 4-core 16 GB, 8-core 32 GB, 16-core 32 GB, and 16-core 64 GB.
Local SSD type	8-core 32 GB (894 GiB of disk space) and 16-core 64 GB (1,788 GiB of disk space).
Local SATA type	8-core 32 GB (22,000 GiB of disk space) and 16-core 64 GB (44,000 GiB of disk space).

**Notice:**

The 1-core 2 GB instances are only for testing purposes. Do not use them for production purposes. The SLA does not cover these instances.

Dedicated master nodes**Notice:**

- If you have purchased 10 or more data nodes, the dedicated master node feature is automatically disabled. You must manually purchase dedicated master nodes.
- The default number of dedicated master nodes is 3 and cannot be changed.
- The default dedicated master node specification is 2-core 8 GB. You can change the specification based on your business needs.
- The default storage type of dedicated master nodes is SSD cloud disk and cannot be changed.
- The default storage space specified for each dedicated master node is 20 GiB and cannot be changed.
- You cannot release your purchased dedicated master nodes.
- You cannot downgrade your purchased dedicated master nodes.

We recommend that you purchase dedicated master nodes to improve the stability of your services. When you purchase an Alibaba Cloud Elasticsearch instance, click **Yes** on the right side of **Dedicated Master Node** to purchase dedicated master nodes. You can also purchase or upgrade dedicated master nodes on the [Cluster upgrade](#) page. The dedicated master nodes will be billed based on the new specification.

Alibaba Cloud Elasticsearch supports the following types of dedicated master nodes. For more information about pricing, see [Pricing](#).

- 1-core 2 GB (currently unavailable)
- 2-core 8 GB (default)
- 4-core 16 GB
- 8-core 32 GB
- 16-core 32 GB
- 16-core 64 GB

**Note:**

- The minimum specification of dedicated master nodes that you can purchase on the Elasticsearch instance buy page or upgrade page is 2-core 8 GB. If you have purchased dedicated master nodes with a specification higher than 1-core 2 GB, then the Yes option is selected for Dedicated Master Node on the cluster upgrade page.
- On the Upgrade page, click Yes on the right side of Dedicated Master Node to purchase or upgrade dedicated master nodes. The nodes will be billed based on the new specification. If your dedicated master nodes are free nodes provided by Elasticsearch, then we will start charging these dedicated master nodes after you upgrade them.
- If you have already purchased dedicated master nodes when you purchase an Alibaba Cloud Elasticsearch instance and the Dedicated Master Node parameter on the cluster upgrade page displays No, this indicates that the specification of the purchased dedicated master nodes is 1-core 2 GB.

Client nodes



Notice:

- The default number of client nodes is 2. Valid values: from 2 to 25.
- The default client node specification is 2-core 8 GB and cannot be changed.
- The default storage type of client nodes is ultra disk and cannot be changed.
- The default storage space specified for each client node is 20 GiB and cannot be changed.
- You cannot release your purchased client nodes.
- You cannot downgrade your purchased client nodes.

You can purchase client nodes to share the CPU overheads of data nodes in order to improve the computing performance and service stability. For CPU-intensive services, we recommend that you purchase client nodes. For example, you can use client nodes to share the overheads if too many aggregation operations are performed. For more information, see [Elasticsearch official node types](#).

When you purchase an Alibaba Cloud Elasticsearch instance, click Yes on the right side of Client Node to purchase client nodes. You can also purchase or upgrade client nodes on the [Cluster upgrade](#) page. The client nodes will be charged based on the new specification.

Alibaba Cloud Elasticsearch supports the following types of client nodes. For more information about pricing, see [Pricing](#).

- 2-core 8 GB (default)
- 4-core 16 GB
- 8-core 32 GB
- 16-core 32 GB
- 16-core 64 GB

Warm nodes



Notice:

- The default number of warm nodes is 2. Valid values: from 2 to 50.
- The default warm node specification is 2-core 8 GB. You can change the specification based on your business needs.
- The default storage type of warm nodes is ultra disk and cannot be changed.
- The default storage space specified for each warm node is 500 GiB and cannot be changed.
- You cannot release your purchased warm nodes.
- You cannot downgrade your purchased warm nodes.

If your business includes the following types of indexes at the same time, we recommend that you purchase warm nodes to implement the hot-warm architecture. The hot-warm architecture can improve the computing performance and service stability of Elasticsearch. For more information, see [ES 5.x hot-warm architecture](#).

- Frequently queried or written indexes
- Infrequently queried or written indexes, typically indexes of records.

When you purchase an Alibaba Cloud Elasticsearch instance, click Yes on the right side of Warm Node to purchase warm nodes. You can also purchase or upgrade warm nodes on the [Cluster upgrade](#) page. The warm nodes will be billed based on the new specification.

Hot - warm architecture

After you purchase warm nodes, the system will add `- Enode . attr . box_type` to the startup configuration of the nodes as follows:

Node type	Startup parameter
Data node	-Enode.attr.box_type=hot
Warm node	-Enode.attr.box_type=warm

Alibaba Cloud Elasticsearch supports the following types of warm nodes. For more information about pricing, see [Pricing](#).

- 2-core 8 GB (default)
- 4-core 16 GB
- 8-core 32 GB



Note:

- Each ultra disk can provide up to 5 TiB of storage space. Ultra disks are cost-effective and can be used in scenarios such as logging and analyzing large amounts of data.

Ultra disks larger than 2 TiB cannot be expanded because they run in disk arrays or RAID 0.

- An ultra disk provides up to 5,120 GiB (5 TiB) of storage space. 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.

Notes

- Node types

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

Node type	Description
Data node	Used as data nodes if you have purchased dedicated master nodes. If no dedicated master node has been purchased, data nodes are also used as dedicated master nodes.
Dedicated master node	Used as dedicated master nodes.
Client node	Used as client nodes.

Node type	Description
Warm node	Used as data nodes and dedicated master nodes if no dedicated master node has been purchased. If you have purchased dedicated master nodes, warm nodes are only used as data nodes.

- Node types and specification families

The following table lists the node types and corresponding specification families supported by Alibaba Cloud Elasticsearch.

Node type	Specification family
Data node	Cloud disk type, local SSD type, and local SATA type.
Dedicated master node	Cloud disk type with a minimum configuration of 2-core and 8 GB.
Client node	Cloud disk type with a minimum configuration of 2-core and 8 GB.
Warm node	Cloud disk type with a minimum configuration of 2-core and 8 GB.

5 Elasticsearch access test

After you have created an Alibaba Cloud Elasticsearch instance, you can log on to the Alibaba Cloud Elasticsearch console, go to the Kibana console, and test the instance on the Dev Tools page, or you can run the `curl` command on an ECS instance that meets the requirements.

Elasticsearch has provided other Elasticsearch clients to run the test. For more information, see [Elasticsearch client](#).

Username and password

You must specify a username and password to log on to the Alibaba Cloud Elasticsearch instance.

- **Username:** The account that is used to access the Alibaba Cloud Elasticsearch instance. we recommend that you do not use the elastic account.
- **password:** The password that you specified when purchasing an Alibaba Cloud Elasticsearch instance or initialized Kibana.



Note:

- You can use the elastic account to access the Alibaba Cloud Elasticsearch instance. However, it takes time for the new password that you specified when modifying the Elasticsearch account to take effect. Services may become unavailable during this period. Therefore, we recommend that you do not use the elastic account to access the Alibaba Cloud Elasticsearch instance.
- If the version of the Alibaba instance that you have created contains information about `with_X-Pack`, you must specify the username and password before you can access the Alibaba Cloud Elasticsearch instance.

Prerequisites for using an ECS instance to access Elasticsearch

- The Alibaba Cloud Elasticsearch instance and the Alibaba Cloud ECS instance must be deployed in the same VPC.
- Alibaba Cloud ECS instances are deployed in the classic network. If your ECS instance needs to access the Elasticsearch instance that is in a VPC, see [Classic network errors](#).

curl testing



Note:

If no `filebeat` index exists on the Alibaba Cloud Elasticsearch instance, run the `PUT filebeat` command to create a corresponding index, or change the YML file to Allow Automatic Indexing. Automatic indexing is disabled by default. Otherwise, the `index_not_found_exception` error occurs when you run the following command.

For Linux

You can run the cURL command to access port `9200` of the Alibaba Cloud Elasticsearch instance from a Linux environment.

Specify a username and password to access the instance, for example:

```
curl -XPOST -u username : password 'http ://< HOST > : 9200 /
filebeat / my_type /'? pretty -d '{" title " : " One ", " tags " : [ "
ruby "]"}'
```

- `< HOST >`: The private or public IP address of the Alibaba Cloud Elasticsearch instance. For more information, see the information page of the Alibaba Cloud Elasticsearch instance.

Response:

```
{
  "_index " : " filebeat ",
  "_type " : " my_type ",
  "_id " : " AV - bTkaTwdiHx faSqlAt ",
  "_version " : 1 ,
  " result " : " created ",
  "_shards " : {
    " total " : 2 ,
    " successful " : 2 ,
    " failed " : 0
  },
  " created " : true
}
```

For Windows

You can run the cURL command to access port `9200` of the Alibaba Cloud Elasticsearch instance from a Windows environment.

Specify a username and password to access the instance, for example:

```
curl -XPOST -u username : password " http ://< HOST > : 9200 /
filebeat / my_type /"? pretty - d '{"title " : " One "', "tags " : [" ruby "]}'"
```

- **< HOST >**: The private or public IP address of the Alibaba Cloud Elasticsearch instance. For more information, see the information page of the Alibaba Cloud Elasticsearch instance.

Response:

```
{
  "_index " : " filebeat ",
  "_type " : " my_type ",
  "_id " : " AWVIU5lY4s iSsiAh0Td6 ",
  "_version " : 1 ,
  "result " : " created ",
  "_shards " : {
    " total " : 2 ,
    " successful " : 2 ,
    " failed " : 0
  },
  " created " : true
}
```

Create a document

Use the following **HTTP POST** method to create a document:

```
curl http ://< HOST > : 9200 / my_index / my_type -XPOST - d '{"
title " : " One ", " tags " : [" ruby "]}'"
```

- **my_index** : The name of the index.
- **< HOST >**: The private or public IP address of the Alibaba Cloud Elasticsearch instance. For more information, see the information page of the Alibaba Cloud Elasticsearch instance.
- Each document has its own **ID** and **type** that are contained in the response. The system will randomly generate an **ID** and **type** if you do not specify them when you create the document.



Note:

If you have enabled automatic indexing (disabled by default) and the specified **index** name does not exist, the system will automatically create an **index** when you create the document.

Sample response when you successfully create a document:

```
{
  "_index": " my_index ",
  "_type": " my_type ",
  "_id": " AV4JIVi15n y3i8DCdK1H ",
  "_version": 1,
  "result": " created ",
  "_shards": {
    "total": 2,
    "successful": 1,
    "failed": 0
  },
  "created": true
}
```

Update a document

You can use the following statements to update an existing document on Elasticsearch:

```
` http ://< HOST >: 9200 / my_index / my_type /< doc_id >`
```

- **< HOST >**: The private or public IP address of the Alibaba Cloud Elasticsearch instance. For more information, see the information page of the Alibaba Cloud Elasticsearch instance.
- **< doc_id >**: The ID of the document.

```
$ curl http://<HOST>:9200/my_index/my_type/AV4JIV i15 ny3 i8 DCdK1H -
XPOST -d '{"title": "Four updated", "tags": ["ruby", "php"]}'
```

Sample response when you successfully update a document:

```
{
  "_index": " my_index ",
  "_type": " my_type ",
  "_id": " AV4JIVi15n y3i8DCdK1H ",
  "version": 2,
  "result": " updated ",
  "_shards": {
    "total": 2,
    "successful": 1,
    "failed": 0
  },
  "created": false
}
```



Note:

You can also use an API to perform a batch update operation on the documents.

Retrieve a document

You can use the following `HTTP GET` method to retrieve a document:

```
$ curl http://<HOST>:9200/my_index/my_type/AV4JIvi15n
y3i8DCdK1H
{
  "_index" : "my_index",
  "_type" : "my_type",
  "_id" : "_b - kbI1MREmi9 SeixFNEVw",
  "_version" : 2,
  "exists" : true,
  "_source" : { "title" : "Four updated", "tags": ["ruby", "
php"] }
}
```

- `<HOST>`: The private or public IP address of the Alibaba Cloud Elasticsearch instance. For more information, see the information page of the Alibaba Cloud Elasticsearch instance.

Search for a document

You can use the `HTTP GET` or `HTTP POST` method to search for a document, and use the URI parameter to specify the object that you want to search for, such as:

```
http://:9200/_search
http://:9200/{index_name}/_search
http://:9200/{index_name}/{type_name}/_search
```

- `<HOST>`: The private or public IP address of the Alibaba Cloud Elasticsearch instance. For more information, see the information page of the Alibaba Cloud Elasticsearch instance.

Example:

```
$ curl http://<HOST>:9200/my_index/my_type/_search?q=
title:T*
```

- `<HOST>`: The private or public IP address of the Alibaba Cloud Elasticsearch instance. For more information, see the information page of the Alibaba Cloud Elasticsearch instance.

Complex search

You must use the following `HTTP POST` method for complex search:

```
$ curl http://:9200/my_index/my_type/_search?pretty=
true -XPOST -d '{
  "query": {
    "query_string": {"query": "*"}
  },
}
```

```
" facets ": {
  " tags ": {
    " terms ": {" field ": " tags " }
  }
}'
```

- **< HOST >**: The internal or public IP address of the Alibaba Cloud Elasticsearch instance. For more information, see the information page of the Alibaba Cloud Elasticsearch instance.

Delete a document

```
$ curl http://< HOST >: 9200 /{ index }/{ type }/{ id } - XDELETE
```

- **< HOST >**: The private or public IP address of the Alibaba Cloud Elasticsearch instance. For more information, see the information page of the Alibaba Cloud Elasticsearch instance.

Delete the specified type of documents

```
$ curl http://< HOST >: 9200 /{ index }/{ type } - XDELETE
```

- **< HOST >**: The private or public IP address of the Alibaba Cloud Elasticsearch instance. For more information, see the information page of the Alibaba Cloud Elasticsearch instance.

Delete an index

```
$ curl http://< HOST >: 9200 /{ index } - XDELETE
```

- **< HOST >**: The private or public IP address of the Alibaba Cloud Elasticsearch instance. For more information, see the information page of the Alibaba Cloud Elasticsearch instance.

6 Configure Elasticsearch features

This topic describes how to use ECS to access Elasticsearch, configure the VPC and public network whitelists, and customize monitoring and alerting settings.

Use ECS to access Elasticsearch

Alibaba Cloud Elasticsearch allows you to use the following methods to access Elasticsearch:

- Log on to the [Kibana console](#) that has been integrated into Alibaba Cloud Elasticsearch, and then access Elasticsearch from the Dev Tools page.
- Use an application or the cURL tool installed on an ECS instance to access Elasticsearch. The ECS instance and Elasticsearch instance must be created in the same region and VPC network.
- Use the public network address of the Elasticsearch instance to access Elasticsearch.
- Use other [Elasticsearch clients](#) listed on the official Elasticsearch site to access Elasticsearch.

Use ECS to access Elasticsearch (optional)



Notice:

- If your ECS instance is connected to the classic network, then you must create a ClassicLink connection to access Elasticsearch. For more information, see [Classic network errors](#).
- The ClassicLink connection is a unidirectional connection originated from the classic network to the VPC network. You cannot use this connection to access the classic network from the VPC network.

If you already have an ECS instance deployed in the same VPC network as your Elasticsearch instance, then you can use this ECS instance as a client to access Elasticsearch. You only need to deploy your application or install the cURL tool on the ECS instance.

If your ECS instance does not meet the preceding requirements, then you must purchase an ECS instance in the same VPC network as your Elasticsearch instance. For more information, see [Create an instance by using the wizard](#).

**Note:**

- The purchased ECS instance is used as a client. You can deploy applications on the ECS instance and then use it to access Elasticsearch.
- If you do not have an ECS instance in the same VPC network as your Elasticsearch instance, you can also use the Kibana console to send test queries to Elasticsearch.

Use ECS and cURL to access Elasticsearch (optional)

1. Log on to the ECS instance that is connected to the same VPC network as your Elasticsearch instance, and then install the cURL tool.
2. Add the cURL path to the path environment variable of the ECS instance so that you can use cURL to access Elasticsearch.

Configure the VPC and public network whitelists

Alibaba Cloud Elasticsearch is integrated with the X-Pack plug-in enterprise edition. The X-Pack plug-in is used to protect data transmission when you use the public or internal network address of your Elasticsearch instance to access Elasticsearch. You can configure the VPC or public network whitelist to implement access control for your Elasticsearch instance. For more information, see [Security configuration](#).

Customize monitoring and alerting settings

Alibaba Cloud Elasticsearch supports monitoring Elasticsearch instances and generating alerts by sending SMS messages. The monitoring data includes the following metrics. You can customize the thresholds for triggering alerts. For more information, see [ES CloudMonitor alarm](#).

- Cluster Status
- Cluster Index Queries (QPS)
- Cluster Write Queries (QPS)
- Node CPU Usage (%)
- Node Disk Space Usage (%)
- Node Heap Memory Usage (%)
- Node Workload Within One Minute

7 Elasticsearch clients

This topic describes how to use the PHP, Python, and Java clients of Elasticsearch to access Alibaba Cloud Elasticsearch. Client examples and usage guidelines are also provided for your reference.

PHP client



Warning:

- The default connection pool provided by the PHP client of Elasticsearch is not suitable for accessing services on the cloud. Alibaba Cloud Elasticsearch handles requests sent from clients based on load balancing. Therefore, the PHP client must use the SimpleConnectionPool method to create a connection pool. Otherwise, when the Alibaba Cloud Elasticsearch is restarted, a connection error occurs. This may adversely affect your businesses.
- However, even if the connection pool is created by using the SimpleConnectionPool method, a connection error may occur when the Elasticsearch instance is restarted. For example, the "No enabled connection" may occur. To resolve this issue, make sure that the PHP client supports auto reconnection.

To test the connectivity to your Elasticsearch instance, connect the PHP client to port 9200 on your instance. The client example is as follows. For more information, see [elasticsearch-php](#).

```
<? php
require 'vendor / autoload . php ';
use Elasticsearch \ ClientBuilder ;

$ client = ClientBuilder :: create ()-> setHosts ([
    [
        ' host ' => '< HOST >',
        ' port ' => ' 9200 ',
        ' scheme ' => ' http ',
        ' user ' => '< USER NAME >',
        ' pass ' => '< PASSWORD >'
    ]
])-> setConnectionPool ('\ Elasticsearch \ Connection Pool \
SimpleConnectionPool ', [])
-> setRetries ( 10 )-> build ();

$ indexParams = [
    ' index ' => ' my_index ',
    ' type ' => ' my_type ',
    ' id ' => ' 1 ',
    ' body ' => [ ' testField ' => ' abc ' ],
    ' client ' => [
```

```

    ' timeout ' => 10 ,
    ' connect_timeout ' => 10
  ]
];
$ indexResponse = $ client -> index ($ indexParams );
print_r ($ indexResponse );

$ searchParams = [
  ' index ' => ' my_index ',
  ' type ' => ' my_type ',
  ' body ' => [
    ' query ' => [
      ' match ' => [
        ' testField ' => ' abc '
      ]
    ]
  ],
  ' client ' => [
    ' timeout ' => 10 ,
    ' connect_timeout ' => 10
  ]
];
$ searchResponse = $ client -> search ($ searchParams );
print_r ($ searchResponse );
? >

```

- **< USER_NAME >**: replace USER NAME with the username of your Elasticsearch instance.
- **< PASSWORD >**: replace PASSWORD with the password of your Elasticsearch instance.
- **< HOST >**: replace HOST with the public or internal network address on the [Basic information](#) page of your Elasticsearch instance.

Python client

To test the connectivity to your Elasticsearch instance, connect the Python client to port 9200 on your instance. The client example is as follows. For more information, see [elasticsearch-py](#).

```

from elasticsearch import Elasticsearch, RequestsHttpConnection
import certifi

es = Elasticsearch (
    ['< HOST >'],
    http_auth=(' username ', ' password '),
    port = 9200 ,
    use_ssl = False
)

res = es . index ( index =" my_index ", doc_type =" my_type ", id
= 1 , body ={" title ": " One ", " tags ": [" ruby "]})
print ( res [' created '])
res = es . get ( index =" my - index ", doc_type =" my - type ",
id = 1 )

```

```
print ( res [' _source '])
```

Java client



Notice:

- The official Elasticsearch team no longer maintains the TransportClient. We recommend that you do not use the TransportClient to access Alibaba Cloud Elasticsearch. A `NoNodeAvailableException` error occurs when you use TransportClient 5.5.3 to access Alibaba Cloud Elasticsearch. We recommend that you use the [Java low level REST client](#) to access Alibaba Cloud Elasticsearch.
- The Java REST client described in this topic is only compatible with Alibaba Cloud Elasticsearch 5.5.3. It is incompatible with Alibaba Cloud Elasticsearch 6.3.2. If you are using Alibaba Cloud Elasticsearch 6.3.2, reference the official Elasticsearch document [Java REST Client 6.3.2](#).
- The version of the Java REST client must be same as the Elasticsearch instance version.

Preparations

- [Create an Alibaba Cloud Elasticsearch instance](#), select Elasticsearch 5.5.3, and then [enable auto-indexing](#). If you have manually created indexes and mappings, then you do not need to enable auto-indexing.
- Install the JDK and configure environment variables. The JDK version must be 1.8 or later.

To test the connectivity to your Elasticsearch instance, connect the Java REST client to port 9200 on your instance. The client example is as follows:

```
import org.apache.http.HttpEntity ;
import org.apache.http.HttpHost ;
import org.apache.http.auth.AuthScope ;
import org.apache.http.auth.UsernamePasswordCredentials ;
import org.apache.http.client.CredentialsProvider ;
import org.apache.http.entity.ContentType ;
import org.apache.http.impl.client.BasicCredentialsProvider ;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder ;
import org.apache.http.nio.entity.NStringEntity ;
import org.apache.http.util.EntityUtils ;
import org.elasticsearch.client.Response ;
import org.elasticsearch.client.RestClient ;
import org.elasticsearch.client.RestClientBuilder ;
import java.io.IOException ;
import java.util.Collections ;
```

```

public class RestClient Test {
    public static void main ( String [] args ){
        final Credential sProvider credential sProvider = new
        BasicCredentialsProvider ();
        credential sProvider . setCredentials ( AuthScope . ANY ,
        new UsernamePasswordCredentials ("< USER NAME
>", "< PASSWORD >"));
        RestClient restClient = RestClient . builder ( new
        HttpHost ("< HOST >", 9200 ))
        . setHttpClientConfigCallback ( new RestClient
        Builder . HttpClientConfigCallback () {
            @Override
            public HttpAsyncClientBuilder customizeH
            ttpClient ( HttpAsyncClientBuilder httpClient Builder ) {
                return httpClient Builder . setDefault
                Credential sProvider ( credential sProvider );
            }
        }). build ();
        try {
            // index a document
            HttpEntity entity = new NStringEntity ("{\n \"
user \" : \" kimchy \"\n }", ContentType . APPLICATIO N_JSON );
            Response indexResponse = restClient . performReq
            uest (
                " PUT ",
                "/ index / type / 123 ",
                Collections . < String , String > emptyMap (),
                entity );
            // search a document
            Response response = restClient . performReq uest ("
GET ", "/ index / type / 123 ",
                Collections . singletonM ap (" pretty ", "
true "));
            System . out . println ( EntityUtils . toString (
            response . getEntity ());
        } catch ( IOException e ) {
            e . printStack Trace ();
        }
    }
}

```

- `< USER NAME >`: replace `USER NAME` with the username of your Elasticsearch instance.
- `< PASSWORD >`: replace `PASSWORD` with the password of your Elasticsearch instance.
- `< HOST >`: replace `HOST` with the public or internal network address shown on the Basic Information page of your Elasticsearch instance.

For more client examples for other languages, see [HTTP/REST Clients and Security](#).

8 Add Elasticsearch Head to Google Chrome

This topic describes how to add Elasticsearch Head to Google Chrome. With Elasticsearch Head, you can use the public network address of an Alibaba Cloud Elasticsearch instance to access the instance and perform operations.

Prerequisites

Before you add Elasticsearch Head to Google Chrome, make sure that you can access the domain `chrome.google.com`.

Context




Note:

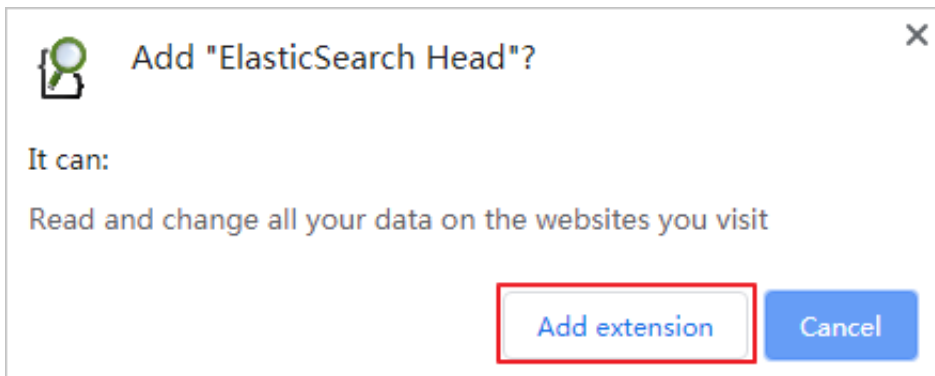
- Elasticsearch Head is a third-party extension.
- In a public network, you cannot use the internal network address and port of an Alibaba Cloud Elasticsearch instance to access the instance with Elasticsearch Head.

Procedure

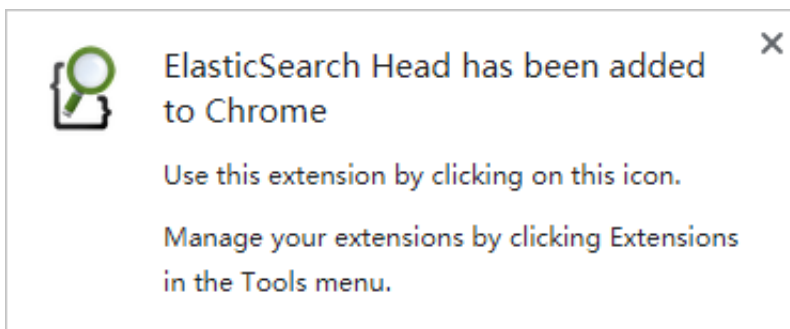
1. Enter the Elasticsearch Head link <https://chrome.google.com/webstore/detail/elasticsearch-head/ffmkiejjmecolpfloofpjologoblkegm> into the address bar of Google Chrome, and then click Add to Chrome.



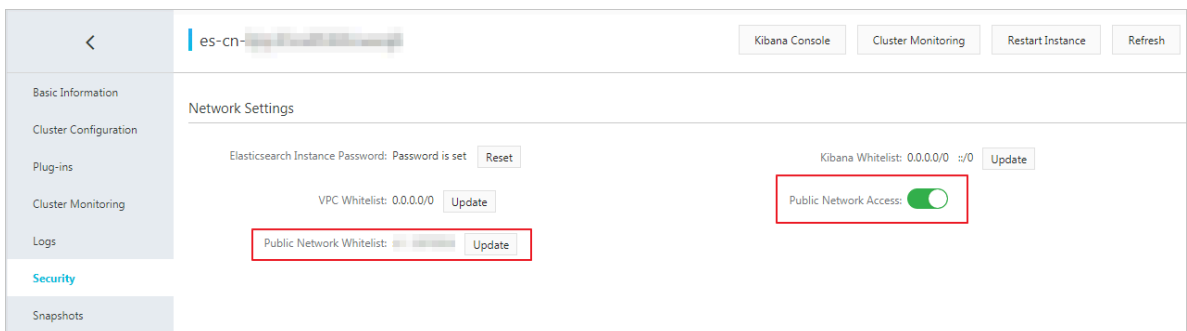
2. Click Add extension in the dialog box.



The system then downloads and installs Elasticsearch Head. After the installation is complete, the system prompts a message indicating that Elasticsearch Head has been installed.




3. Log on to the [Alibaba Cloud Elasticsearch console](#), enable **public address** for your Elasticsearch instance, and add the public network IP address of your host to **public IP addresses with whitelist**.




Note:

- You can enter IP into the search bar on the Google homepage, and then click the **What Is My IP Address** link to view the public network IP address of your host.
- By default, the public network access function forbids all IPv4 addresses.

- 4. Click the Elasticsearch Head icon on the right side of the Google Chrome address bar to open the Elasticsearch cluster connection page.



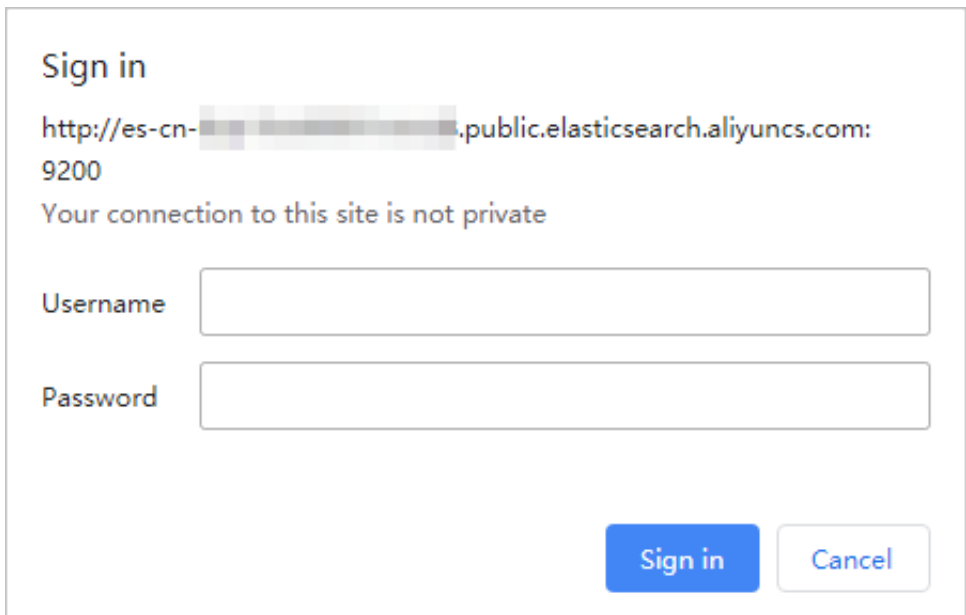
- 5. Enter `http://Elasticsearch Instance Public Network Address: Port Number/` into the address bar, and then click Connect.



You can log on to the Elasticsearch console and view the public network address and port number of the Elasticsearch instance on the basic information page. The default port number is 9200. The following is a sample connection address:

```
http://es-cn-45xxxxxxxxx01xw6w.public.elasticsearch.aliyuncs.com:9200/
```

- 6. In the Sign in dialog box, enter the username and password that are used to log on to the Kibana console of the Elasticsearch instance, and then click Log on.



Note:

The Alibaba Cloud Elasticsearch with Commercial Feature is integrated with X-Pack for security purposes. Therefore, you must enter the username and password for authentication before you can log on to the instance. If the Sign in dialog box is not displayed, verify that the public network whitelist of Alibaba Cloud Elasticsearch contains the public network IP address of your host, or clear the cache of your Web browser and then try again.

7. After you log on to the Elasticsearch instance, you can then perform other operations.

