

# Alibaba Cloud

Elasticsearch  
AliES Kernel

Document Version: 20220215

# 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
 <b>Danger</b>	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 <b>Danger:</b> Resetting will result in the loss of user configuration data.
 <b>Warning</b>	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 <b>Warning:</b> Restarting will cause business interruption. About 10 minutes are required to restart an instance.
 <b>Notice</b>	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	 <b>Notice:</b> If the weight is set to 0, the server no longer receives new requests.
 <b>Note</b>	A note indicates supplemental instructions, best practices, tips, and other content.	 <b>Note:</b> You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click <b>Settings</b> > <b>Network</b> > <b>Set network type</b> .
<b>Bold</b>	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click <b>OK</b> .
<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.AliES release notes -----	05
2.Use the pruning feature for a time series index -----	07
3.Use the slow query isolation feature -----	10

# 1. AliES release notes

AliES is a highly tailored kernel for Alibaba Cloud Elasticsearch. AliES supports all the features provided by the open source Elasticsearch kernel and provides additional features, such as metric optimization, thread pooling, circuit breaking optimization, and query and write performance optimization. These additional features are developed based on the abundant experience of the Alibaba Cloud Elasticsearch team in multiple scenarios. These features not only improve cluster stability and performance but also reduce costs and extend the scope of monitoring and O&M. This topic describes the optimized features in each version of AliES.

## Elasticsearch V7.10.0 of the Standard Edition

Kernel version	Description
V1.4.0	<ul style="list-style-type: none"> <li>The <b>aliyun-knn</b> plug-in is updated to improve write performance. The plug-in supports script queries and is integrated with the optimized capabilities of the related hardware to improve the <b>vector search</b> feature.</li> <li>The <b>aliyun-qos</b> plug-in is optimized to improve cluster-level throttling. When you use this plug-in, you do not need to focus on the topology and load of the nodes in your Elasticsearch cluster. Traffic is automatically distributed to the nodes. This improves cluster usability and stability.</li> </ul>
V1.3.0	<ul style="list-style-type: none"> <li>The <b>slow query isolation</b> feature is provided to reduce the impact of anomalous queries on cluster stability.</li> <li>The <b>gig plug-in</b> is provided to perform a switchover within seconds after an exception occurs on a cluster. This plug-in prevents query jitters caused by anomalous nodes.</li> </ul> <div> <p> <b>Note</b> For Elasticsearch V7.10.0 clusters of the Standard Edition, the gig plug-in is integrated into the aliyun-qos plug-in. The aliyun-qos plug-in is installed by default.</p> </div> <ul style="list-style-type: none"> <li>The <b>physical replication</b> feature is provided to improve the write performance of indexes that have replica shards.</li> <li>The <b>pruning</b> feature is provided for time series indexes to improve the query performance of the indexes.</li> <li>The <b>access logs</b> of clusters can be viewed. These logs contain fields such as Time, Node IP, and Content. You can use these logs to troubleshoot issues and analyze requests.</li> <li>The <b>scheduling performance of dedicated master nodes is improved by 10 times</b>. Each dedicated master node can schedule more shards.</li> </ul>

## Elasticsearch V6.7.0 of the Standard Edition

Kernel version	Description
----------------	-------------

Kernel version	Description
V1.3.0	<ul style="list-style-type: none"> <li>The <b>slow query isolation</b> feature is provided to reduce the impact of anomalous queries on cluster stability.</li> <li>The <b>gig plug-in</b> is provided to perform a switchover within seconds after an exception occurs on a cluster. This plug-in prevents query jitters caused by anomalous nodes.</li> </ul> <div>  <b>Notice</b> Before you use the preceding features, make sure that the kernel version of your Elasticsearch cluster is V1.3.0. Otherwise, upgrade the kernel. You can upgrade only the kernels of Standard Edition clusters whose kernel versions are V0.3.0, V1.0.2, or V1.3.0. </div>
V1.2.0	<ul style="list-style-type: none"> <li>The <b>physical replication</b> feature is provided to improve the write performance of indexes that have replica shards.</li> <li>The <b>pruning</b> feature is provided for time series indexes to improve the query performance of the indexes.</li> <li>Primary key-based data deduplication is optimized during queries. This improves the write performance of documents that contain primary keys by 10%.</li> <li>Finite state transducers (FSTs) that do not occupy heap memory are supported. A single node can store a maximum of 20 TiB of index data.</li> </ul>
V1.0.2	<p>The <b>access logs</b> of clusters can be viewed. These logs contain fields such as Time, Node IP, and Content. You can use these logs to troubleshoot issues and analyze requests.</p>
V1.0.1	<p>Circuit breaking policies can be configured for Java virtual machines (JVMs). When the usage of the JVM heap memory of your cluster reaches 95%, the system rejects requests to protect the cluster. The following parameters are used to configure the policies:</p> <ul style="list-style-type: none"> <li><code>indices.breaker.total.use_real_memory</code>: The default value is false.</li> <li><code>indices.breaker.total.limit</code>: The default value is 95%.</li> </ul>
V0.3.0	<ul style="list-style-type: none"> <li>The <b>scheduling performance of dedicated master nodes is improved by 10 times</b>. Each dedicated master node can schedule more shards.</li> <li>Write performance is improved by 10%, and the overheads of translog flush are reduced.</li> </ul>

## 2. Use the pruning feature for a time series index

When you query data from a time series index, you can specify a time range to filter the data. The larger the data volume, the longer the query latency. You can use the pruning feature to improve query performance. This topic describes how to use the pruning feature.


### Prerequisites

An Alibaba Cloud Elasticsearch V6.7.0 or V7.10.0 cluster is created. The kernel version of the V6.7.0 cluster is 1.2.0 or later. For more information about how to create a cluster, see [Create an Alibaba Cloud Elasticsearch cluster](#).

### Context

The pruning feature improves segment merging and query performance. The following table compares open source Elasticsearch with Alibaba Cloud Elasticsearch in terms of the segment merging policy and query performance.

Service	Segment merging policy	Query performance
Open source Elasticsearch	Index segments are merged based on their sizes. Index segments that have similar sizes are merged. This method is efficient but cannot ensure data continuity.	All data is scanned for a query, which causes serious performance loss.
Alibaba Cloud Elasticsearch (with the pruning feature enabled)	A time series field is added when an index is created. Index segments are merged based on their sizes and the time series field. Index segments that have similar sizes and contain data in adjacent time periods are merged. This improves data continuity.	Data is pruned based on the time range specified in a query, which improves query performance by 40%.

 **Note** You can run all the commands that are provided in this topic in the Kibana console. For more information, see [Log on to the Kibana console](#).

### Precautions

- The pruning feature is available only for Alibaba Cloud Elasticsearch V6.7.0 clusters whose kernel versions are V1.2.0 or later and V7.10.0 clusters.
- You must enable the pruning feature for an index when you create the index. If you enable the feature after an index is created, the expected query performance cannot be achieved.
- After you disable the pruning feature for an index, we recommended that you do not enable the feature again for the index. If you enable it again and the segment merging of the index covers non-time series data, the expected query performance cannot be achieved.

### Enable the pruning feature for an index

To enable the pruning feature, specify a time series field when you create an index. The following code uses `timestamp` as the time series field:

```
PUT index-1/_settings
{
  "index" : {
    "merge.policy.time_series_field" : "timestamp"
  }
}
```



**Notice** The data type of the time series field must be DATE or LONG.

## Query data based on the specified time series field

After you run the following command, the system filters data based on the timestamp field and then searches for data.

```
POST index-1/_search
{
  "query": {
    "bool": {
      "filter": [
        {
          "range": {
            "timestamp": {
              "format": "yyyy-MM-dd HH:mm:ss",
              "gte": "2020-06-01 23:00:00",
              "lt": "2020-06-06 23:05:00",
              "time_zone": "+08:00"
            }
          }
        },
        {
          "terms": {
            "region": [
              "sh"
            ]
          }
        }
      ]
    }
  }
}
```

## Disable the pruning feature for an index

1. Disable the index.

```
POST index-1/_close
```

2. Update the `settings` of the index to disable the pruning feature.



```
PUT index-1/_settings
{
  "index" : {
    "merge.policy.time_series_field" : null
  }
}
```

### 3. Enable the index again.

```
POST index-1/_open
```

## 3. Use the slow query isolation feature


When you use Elasticsearch for queries, you may encounter the following issue: You send a query request to an Elasticsearch cluster, but the query is defined as a slow query. As a result, all the resources on the nodes in the cluster are used for the query, which affects your online business. To address this issue, the Alibaba Cloud Elasticsearch team develops the slow query isolation feature. This feature can be used to track the overheads for a single query request and implement logical separation. If the overheads for the request exceed a specific threshold, the system considers the query as an anomalous query and suspends it. This prevents exceptions caused by a single anomalous query in the cluster and improves cluster stability. This topic describes how to use the slow query isolation feature.

### Context

To use the slow query isolation feature, you must configure a resource isolation pool that has a fixed memory size. If the size of the memory used for a single query request exceeds a specific threshold, the query request is directed to the isolation pool for management. If the total size of the memory used by the query requests in the pool exceeds a specific threshold, the system suspends the query requests that consume the most memory based on a priority policy. The priority policy can be adopted by users based on their business requirements.

### Precautions

- The slow query isolation feature is available for Alibaba Cloud Elasticsearch V6.7.0 clusters whose kernel versions are V1.3.0 and Alibaba Cloud Elasticsearch V7.10.0 clusters.


 **Note** If the version of your Elasticsearch cluster is V6.7.0, you must make sure that the kernel version of the cluster is V1.3.0 before you use the slow query isolation feature. If the kernel version of the cluster is not V1.3.0, upgrade the kernel. You can upgrade the kernels only of Standard Edition clusters whose kernel versions are V0.3.0, V1.0.2, or V1.2.0. If the version of your Elasticsearch cluster is V7.10.0, you can directly use this feature.

- The slow query isolation feature is disabled by default. You must enable the feature before you use it.
- All the commands provided in this topic can be run in the Kibana console. For more information about how to log on to the Kibana console, see [Log on to the Kibana console](#).

### Procedure

1. Enable the slow query isolation feature.

```
PUT _cluster/settings
{
  "persistent": {
    "search.isolator.enabled": true
  }
}
```

 **Note** If you want to disable the feature, set `search.isolator.enabled` to null or false.

2. Configure thresholds to intercept query requests. If the size or latency of a query request exceeds the related threshold, the query request is directed to the slow query isolation pool.

```
PUT _cluster/settings
{
  "persistent": {
    "search.isolator.trigger.task.mem_cost": "500mb",
    "search.isolator.trigger.task.latency": "10s"
  }
}
```

Parameter	Default value	Description
search.isolator.trigger.task.mem_cost	100mb	The threshold for the size of the memory that can be used for a single query request. If the size of the memory that is used for a query request exceeds the threshold, the system directs the query request to the slow query isolation pool.
search.isolator.trigger.task.latency	10s	The threshold for the latency of a query request. If the time spent on a query request exceeds the threshold, the system directs the query request to the slow query isolation pool.

3. Configure the thresholds for the total size of the memory that can be used for the query requests in the slow query isolation pool and the maximum number of query requests that can be processed at the same time in the slow query isolation pool. If the total size of the memory used by the query requests in the slow query isolation pool or the number of query requests that are processed at the same time in the slow query isolation pool exceeds the related threshold, the system suspends the query requests that consume the most memory in the isolation pool.

```
PUT _cluster/settings
{
  "persistent": {
    "search.isolator.total.mem.limit": "60%",
    "search.isolator.total.heap.usage.limit": "75%",
    "search.isolator.total.tasks.limit": 1000
  }
}
```

Parameter	Default value	Description
search.isolator.total.mem.limit	60%	The threshold for the proportion of the heap memory that is consumed by the query requests in the slow query isolation pool to the memory of the whole cluster. The default value is 60%. This value indicates that the query requests in the slow query isolation pool are suspended if the proportion reaches 60%.

Parameter	Default value	Description
search.isolator.total.heap.usage.limit	75%	The threshold for the heap memory usage of the cluster. The default value is 75%. This value indicates that the query requests in the slow query isolation pool are suspended if the usage reaches 75%.
search.isolator.total.tasks.limit	1000	The maximum number of query requests that can be processed at the same time in the slow query isolation pool. The default value is 1000. This value indicates that the query requests in the slow query isolation pool are suspended if the number of query requests that are processed at the same time exceeds 1,000.

#### 4. View the query requests in the slow query isolation pool.

```
GET _tasks/isolator?detailed=true
```

#### 5. Cancel a query request.

```
POST _tasks/<taskId>/_cancel
```

 **Note** *taskId*: the ID of the query request.