

Elasticsearch Service

User Guide

Product Documentation



Copyright Notice

©2013-2024 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

User Guide

Last updated : 2023-12-20 21:23:25

This document describes how to get started with Elasticsearch Service (ES).

1. Basic knowledge of ES

- [Why ES?](#)
- [What features does ES have?](#)
- [Performance of ES](#)
- [Use cases of ES](#)
- [Elastic Stack \(X-Pack\)](#)
- [Capabilities and restrictions of ES](#)

2. ES billing mode

ES is pay-as-you-go. For more information, please see [Billing Overview](#).

3. Getting started

3.1. Evaluation of cluster specification and capacity configuration

Before purchasing a cluster, you need to evaluate your specific business according to the actual situation to ensure that the created cluster meets your actual needs. For more information, please see [Evaluation of Cluster Specification and Capacity Configuration](#).

3.2. ES cluster purchase

Before using ES, you need to sign up for a Tencent Cloud account first and then click **Buy Now** on the [purchase page](#) to create a cluster. For more information, please see [Creating Clusters](#).

3.3. Cluster access

After successfully creating a cluster, you can start to access it through the [API](#), [client](#), or [Kibana](#). Enabling [ES cluster user authentication](#) can improve the security of cluster access.

4. Overview of console features

If you want to	You can read
View cluster status information in the cluster list	Cluster Status
Restart a cluster as needed	Restarting Clusters
Terminate a cluster	Terminating Clusters
Create a cluster that supports multi-AZ	Multi-AZ Cluster Deployment
Understand the concepts and principles of cluster configuration adjustment and scale a cluster	Adjusting Configuration and Suggestions and Principles for Cluster Specification Adjustment
Configure synonyms and YML in a cluster	Synonym Configuration and YML File Configurations
Configure plugins in a cluster	Plugin List , IK Analysis Plugin , and QQ Analysis Plugin
Monitor the status of a running cluster	Viewing Monitors , Configuring Alarms , and Suggestions for Configuring Monitors and Alarms
Query cluster logs	Querying Cluster Logs
Back up data	Automatic Snapshot Backup and Using COS for Backup and Restoration
Upgrade a cluster and use advanced features	Upgrading ES Clusters and ES Version Upgrade Check

5. Best practice

5.1. Data migration and sync

1. Data migration

If you want to migrate your data to ES, you can choose a suitable migration solution based on your business needs, such as COS snapshot, Logstash, and elasticsearch-dump. For more information, please see [Data Migration](#).

2. Data ingestion into ES

You can connect your data source of different types to ES through the Logstash and Beats components. For more information, please see [Data Ingestion into ES](#).

3. Real-time MySQL data sync to ES

You can sync data to ES in real time by syncing MySQL binlog. For more information, please see [Syncing MySQL Data to ES in Real Time](#).

5.2. Use case construction

1. Build a log analysis system

You can import your logs into ES and access Kibana from a browser to perform query and analysis (for example, through the most typical log analysis architectures Filebeat + Elasticsearch + Kibana and Logstash + Elasticsearch + Kibana). For more information, please see [Building a Log Analysis System](#).

5.3. Index settings

1. Default index template description and adjustment

You can describe and adjust the default template. For more information, please see [Default Index Template Description and Adjustment](#).

2. Index management with Curator

By managing indexes with Curator, you can clear indexes created 7 days ago, back up specified indexes regularly every day, and migrate indexes from a hot node to a warm node regularly. For more information, please see [Managing Indexes with Curator](#).

3. Hot/Warm architecture and index lifecycle management

You can specify the specifications of hot and warm nodes based on your business needs to quickly build an ES cluster in the hot/warm architecture. For more information, please see [Hot/Warm Architecture and Index Lifecycle Management](#).

5.4. SQL support

ES supports SQL instead of DSL as the query language. For those engaged in product operations and data analysis and new ES users, using SQL for queries can reduce their learning costs for getting started with ES. For more information, please see [SQL Support](#).

6. FAQs for beginners

6.1. Product

- [What business scenarios is ES suitable for?](#)
- [I have an unpaid switch order. Will the order still be valid if I upgrade the cluster configuration?](#)
- [Can I change the cloud disk type after a successful purchase?](#)

6.2 Cluster exceptions

- [Exceptional Cluster Health Status \(Red and Yellow\)](#)
- [Memory Cleanup Methods](#)
- [Bulk Rejection/Search Rejection](#)
- [High Cluster CPU Utilization](#)
- [High Cluster Disk Utilization and read_only Status](#)
- [Uneven Cluster Load](#)

7. Feedback and suggestions

If you have any questions or suggestions about ES, you can send your feedback through the following channels, and we will get back to you accordingly:

- If you find issues with product documentation, such as links, contents, and APIs, you can click **Send Feedback** on the right of the document page and select the specific issues.
- If you encounter problems when using the product, please [submit a ticket](#) for assistance.