

Elasticsearch Service

Product Introduction

Product Documentation



Copyright Notice

©2013-2019 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.

Contents

Product Introduction

Overview

Features

Performance

Strengths

Scenarios

Elastic Stack (X-Pack)

Capabilities and Restrictions

Related Concepts

Product Introduction

Overview

Last updated : 2020-02-26 21:55:27

Tencent Cloud Elasticsearch Service (ES) is a cloud-managed Elasticsearch service that is highly available and scalable. It is built on Elasticsearch, a RESTful API-based distributed search and analysis engine that can be used to search and analyze massive data.

While retaining Elasticsearch's compatibility and openness, ES incorporates Tencent Cloud's computing, storage, and security technologies. It has various cluster management functions and is secure, elastic, and highly available. Additionally, ES integrates the official Elastic Stack features (formerly X-Pack), which adds permission management, SQL, machine learning, and alert features to the open source foundation. These features simplify basic OPS tasks like cluster deployment and operation management, letting you focus on the business.

ES allows you to quickly build applications like massive data storage and search, and real-time log analysis, website search and navigation, search for enterprises, log monitoring, and click analysis.

Main Components

- Elasticsearch

Elasticsearch is a distributed search engine which stores massive data, search the whole text and analyzes statistics. Its RESTful APIs and different programming language clients make it easy to develop based on your business needs.

- Kibana

This is a data visualization tool that makes it easy to query and analyze the data stored in an Elasticsearch cluster.

- Elastic Stack (formerly X-Pack)

Elastic Stack is Elasticsearch's official plugin which has various advanced features. Its data permission management can be applied to the field level. It can be easily integrated to your existing business via SQL and JDBC connection. Its machine learning and alerts can analyze cluster data and fluctuations to predict data trends and send out huge fluctuation alarms.

Features

Last updated : 2020-02-17 14:01:08

Real-time logs of other cloud products such as CVM, CDB and TKE and the stocked and incremental business data can be aggregated and transferred to the ES cluster for distributed data storage, query and analysis.

Data collection and synchronization

- You can use the Beats feature in ES to transfer data to ES cluster for storage, or to Logstash for custom conversion and parse before transferring them to ES cluster.
- ES provides easy-to-use RESTful API for you to develop your own client, and you can call the data storage API to store the data in ES clusters.
- ES is built in a VPC, and you can easily use various data synchronization plug-ins to sync the data of existing cloud products into ES clusters.

Data storage

- ES provides different types of nodes and high-performance SSDs, ensuring the data read/write performance.
- It can be elastically scaled to hundreds of nodes for data storage at the petabyte level, satisfying the needs of different business scenarios.
- It can detect and replace faulty nodes, ensuring high cluster availability.
- It features full-text search.

Data query, analysis and visualization

- ES features full-text search, structured query, data filtering, metric statistics etc., which is applicable to information search, data analysis and many other scenarios.
- ES provides easy-to-use RESTful API and clients in various languages for you to build your own search services.
- With Kibana, you can easily search and statistically analyze cluster data in a browser.

Performance

Last updated : 2019-11-08 17:43:59

The Rally stress test tool provided by Elasticsearch and the GeoNames dataset (3.2 GB in size with 11,523,468 documents) are used to perform stress test on the ES clusters of two different specifications built by Tencent Cloud Elasticsearch Service. The performance test results are as shown below:

2 cores, 8 GB memory, 3 nodes

Metric	Task	Value	Unit
Total indexing time	-	18.5837	min
Min indexing time per shard	-	0.000533333	min
Median indexing time per shard	-	3.4134	min
Max indexing time per shard	-	4.21195	min
Total refresh time	-	2.40147	min
Min refresh time per shard	-	0.0009	min
Median refresh time per shard	-	0.45715	min
Max refresh time per shard	-	0.5655	min
Total flush time	-	0.328317	min
Min flush time per shard	-	0.0004	min
Median flush time per shard	-	0.0643583	min
Max flush time per shard	-	0.0972167	min
Total Young Gen GC	-	127.992	s
Total Old Gen GC	-	1.327	s
Store size	-	3.22145	GB
Translog size	-	2.80E-07	GB
Heap used for segments	-	10.8161	MB

Metric	Task	Value	Unit
Heap used for doc values	-	0.0831299	MB
Heap used for terms	-	9.62618	MB
Heap used for norms	-	0.0795288	MB
Heap used for points	-	0.217488	MB
Heap used for stored fields	-	0.809799	MB
Segment count	-	102	-
Min Throughput	index-append	41895.6	docs/s
Median Throughput	index-append	42562	docs/s
Max Throughput	index-append	43352	docs/s
50th percentile latency	index-append	757.636	ms
90th percentile latency	index-append	1388.11	ms
99th percentile latency	index-append	2650.44	ms
99.9th percentile latency	index-append	5270.13	ms
100th percentile latency	index-append	6535.29	ms
50th percentile service time	index-append	757.636	ms
90th percentile service time	index-append	1388.11	ms
99th percentile service time	index-append	2650.44	ms
99.9th percentile service time	index-append	5270.13	ms
100th percentile service time	index-append	6535.29	ms
error rate	index-append	0	%
Min Throughput	index-stats	99.99	ops/s
Median Throughput	index-stats	100.04	ops/s
Max Throughput	index-stats	100.06	ops/s
50th percentile latency	index-stats	6.04131	ms

Metric	Task	Value	Unit
90th percentile latency	index-stats	6.56561	ms
99th percentile latency	index-stats	10.7941	ms
99.9th percentile latency	index-stats	22.6768	ms
100th percentile latency	index-stats	24.5623	ms
50th percentile service time	index-stats	5.9341	ms
90th percentile service time	index-stats	6.42644	ms
99th percentile service time	index-stats	7.56809	ms
99.9th percentile service time	index-stats	22.5948	ms
100th percentile service time	index-stats	24.4467	ms
error rate	index-stats	0	%
Min Throughput	node-stats	99.93	ops/s
Median Throughput	node-stats	100.06	ops/s
Max Throughput	node-stats	100.33	ops/s
50th percentile latency	node-stats	6.74088	ms
90th percentile latency	node-stats	7.28822	ms
99th percentile latency	node-stats	8.62256	ms
99.9th percentile latency	node-stats	13.271	ms
100th percentile latency	node-stats	13.9379	ms
50th percentile service time	node-stats	6.64634	ms
90th percentile service time	node-stats	7.18403	ms
99th percentile service time	node-stats	8.34209	ms
99.9th percentile service time	node-stats	13.1784	ms
100th percentile service time	node-stats	13.8411	ms
error rate	node-stats	0	%

Metric	Task	Value	Unit
Min Throughput	default	46.85	ops/s
Median Throughput	default	47.58	ops/s
Max Throughput	default	47.7	ops/s
50th percentile latency	default	1023.59	ms
90th percentile latency	default	1890.67	ms
99th percentile latency	default	2017.58	ms
99.9th percentile latency	default	2026.21	ms
100th percentile latency	default	2026.21	ms
50th percentile service time	default	20.2853	ms
90th percentile service time	default	24.3425	ms
99th percentile service time	default	33.3526	ms
99.9th percentile service time	default	54.702	ms
100th percentile service time	default	74.2832	ms
error rate	default	0	%
Min Throughput	term	199.82	ops/s
Median Throughput	term	200.02	ops/s
Max Throughput	term	200.05	ops/s
50th percentile latency	term	4.54844	ms
90th percentile latency	term	9.97587	ms
99th percentile latency	term	18.3546	ms
99.9th percentile latency	term	21.2112	ms
100th percentile latency	term	21.5267	ms
50th percentile service time	term	4.39147	ms
90th percentile service time	term	4.84527	ms

Metric	Task	Value	Unit
99th percentile service time	term	6.62189	ms
99.9th percentile service time	term	19.4932	ms
100th percentile service time	term	21.4435	ms
error rate	term	0	%
Min Throughput	phrase	199.68	ops/s
Median Throughput	phrase	200.03	ops/s
Max Throughput	phrase	200.11	ops/s
50th percentile latency	phrase	3.57488	ms
90th percentile latency	phrase	4.62139	ms
99th percentile latency	phrase	19.6935	ms
99.9th percentile latency	phrase	24.9076	ms
100th percentile latency	phrase	25.0486	ms
50th percentile service time	phrase	3.4742	ms
90th percentile service time	phrase	4.0265	ms
99th percentile service time	phrase	7.57333	ms
99.9th percentile service time	phrase	18.9011	ms
100th percentile service time	phrase	23.8045	ms
error rate	phrase	0	%
Min Throughput	country_agg_uncached	4.99	ops/s
Median Throughput	country_agg_uncached	5	ops/s
Max Throughput	country_agg_uncached	5	ops/s
50th percentile latency	country_agg_uncached	197.534	ms
90th percentile latency	country_agg_uncached	217.842	ms
99th percentile latency	country_agg_uncached	271.988	ms

Metric	Task	Value	Unit
100th percentile latency	country_agg_uncached	275.963	ms
50th percentile service time	country_agg_uncached	194.061	ms
90th percentile service time	country_agg_uncached	209.086	ms
99th percentile service time	country_agg_uncached	216.432	ms
100th percentile service time	country_agg_uncached	217.275	ms
error rate	country_agg_uncached	0	%
Min Throughput	country_agg_cached	99.97	ops/s
Median Throughput	country_agg_cached	100.05	ops/s
Max Throughput	country_agg_cached	100.08	ops/s
50th percentile latency	country_agg_cached	4.9212	ms
90th percentile latency	country_agg_cached	5.44065	ms
99th percentile latency	country_agg_cached	7.15509	ms
99.9th percentile latency	country_agg_cached	16.9407	ms
100th percentile latency	country_agg_cached	17.8111	ms
50th percentile service time	country_agg_cached	4.81515	ms
90th percentile service time	country_agg_cached	5.29377	ms
99th percentile service time	country_agg_cached	6.38482	ms
99.9th percentile service time	country_agg_cached	16.8318	ms
100th percentile service time	country_agg_cached	17.7311	ms
error rate	country_agg_cached	0	%
Min Throughput	scroll	25.02	pages/s
Median Throughput	scroll	25.02	pages/s
Max Throughput	scroll	25.03	pages/s
50th percentile latency	scroll	760.634	ms

Metric	Task	Value	Unit
90th percentile latency	scroll	794.699	ms
99th percentile latency	scroll	864.897	ms
100th percentile latency	scroll	874.768	ms
50th percentile service time	scroll	760.32	ms
90th percentile service time	scroll	794.397	ms
99th percentile service time	scroll	864.658	ms
100th percentile service time	scroll	874.556	ms
error rate	scroll	0	%
Min Throughput	expression	2	ops/s
Median Throughput	expression	2	ops/s
Max Throughput	expression	2	ops/s
50th percentile latency	expression	382.483	ms
90th percentile latency	expression	414.775	ms
99th percentile latency	expression	455.236	ms
100th percentile latency	expression	473.181	ms
50th percentile service time	expression	382.298	ms
90th percentile service time	expression	414.577	ms
99th percentile service time	expression	455.11	ms
100th percentile service time	expression	472.998	ms
error rate	expression	0	%
Min Throughput	painless_static	1.5	ops/s
Median Throughput	painless_static	1.5	ops/s
Max Throughput	painless_static	1.5	ops/s
50th percentile latency	painless_static	480.188	ms

Metric	Task	Value	Unit
90th percentile latency	painless_static	505.003	ms
99th percentile latency	painless_static	529.066	ms
100th percentile latency	painless_static	547.199	ms
50th percentile service time	painless_static	479.938	ms
90th percentile service time	painless_static	504.731	ms
99th percentile service time	painless_static	528.857	ms
100th percentile service time	painless_static	546.954	ms
error rate	painless_static	0	%
Min Throughput	painless_dynamic	1.5	ops/s
Median Throughput	painless_dynamic	1.5	ops/s
Max Throughput	painless_dynamic	1.5	ops/s
50th percentile latency	painless_dynamic	469.434	ms
90th percentile latency	painless_dynamic	508.615	ms
99th percentile latency	painless_dynamic	581.127	ms
100th percentile latency	painless_dynamic	621.998	ms
50th percentile service time	painless_dynamic	469.178	ms
90th percentile service time	painless_dynamic	508.349	ms
99th percentile service time	painless_dynamic	580.819	ms
100th percentile service time	painless_dynamic	621.799	ms
error rate	painless_dynamic	0	%
error rate	large_terms	0	%
Min Throughput	large_filtered_terms	1.52	ops/s
Median Throughput	large_filtered_terms	1.52	ops/s
Max Throughput	large_filtered_terms	1.52	ops/s

Metric	Task	Value	Unit
50th percentile latency	large_filtered_terms	39664.6	ms
90th percentile latency	large_filtered_terms	46001.3	ms
99th percentile latency	large_filtered_terms	47328	ms
100th percentile latency	large_filtered_terms	47488.6	ms
50th percentile service time	large_filtered_terms	651.731	ms
90th percentile service time	large_filtered_terms	673.319	ms
99th percentile service time	large_filtered_terms	715.941	ms
100th percentile service time	large_filtered_terms	723.06	ms
error rate	large_filtered_terms	0	%
Min Throughput	large_prohibited_terms	1.55	ops/s
Median Throughput	large_prohibited_terms	1.56	ops/s
Max Throughput	large_prohibited_terms	1.57	ops/s
50th percentile latency	large_prohibited_terms	35606.8	ms
90th percentile latency	large_prohibited_terms	40847.6	ms
99th percentile latency	large_prohibited_terms	42170.5	ms
100th percentile latency	large_prohibited_terms	42329.2	ms
50th percentile service time	large_prohibited_terms	648.82	ms
90th percentile service time	large_prohibited_terms	672.114	ms
99th percentile service time	large_prohibited_terms	722.666	ms
100th percentile service time	large_prohibited_terms	733.307	ms
error rate	large_prohibited_terms	0	%

4 Cores, 16 GB Memory, 3 Nodes

Metric	Task	Value	Unit
--------	------	-------	------

Metric	Task	Value	Unit
Total indexing time	-	20.1957	min
Min indexing time per shard	-	0.0007333333	min
Median indexing time per shard	-	3.77953	min
Max indexing time per shard	-	4.63752	min
Total merge time	-	1.57487	min
Min merge time per shard	-	0	min
Median merge time per shard	-	0.176658	min
Max merge time per shard	-	0.634067	min
Total merge throttle time	-	0.55105	min
Min merge throttle time per shard	-	0	min
Median merge throttle time per shard	-	0.065	min
Max merge throttle time per shard	-	0.217033	min
Total refresh time	-	1.41135	min
Min refresh time per shard	-	0.00106667	min
Median refresh time per shard	-	0.269958	min
Max refresh time per shard	-	0.345733	min
Total flush time	-	0.533133	min
Min flush time per shard	-	0.000566667	min
Median flush time per shard	-	0.115592	min
Max flush time per shard	-	0.136683	min
Total Young Gen GC	-	70.747	s
Total Old Gen GC	-	0.92	s
Store size	-	3.31581	GB
Translog size	-	2.80E-07	GB

Metric	Task	Value	Unit
Heap used for segments	-	11.0486	MB
Heap used for doc values	-	0.100529	MB
Heap used for terms	-	9.84413	MB
Heap used for norms	-	0.0755005	MB
Heap used for points	-	0.216421	MB
Heap used for stored fields	-	0.811981	MB
Segment count	-	97	-
Min Throughput	index-append	74421.1	docs/s
Median Throughput	index-append	75636.9	docs/s
Max Throughput	index-append	76877.4	docs/s
50th percentile latency	index-append	377.922	ms
90th percentile latency	index-append	663.055	ms
99th percentile latency	index-append	3068.99	ms
100th percentile latency	index-append	5554.97	ms
50th percentile service time	index-append	377.922	ms
90th percentile service time	index-append	663.055	ms
99th percentile service time	index-append	3068.99	ms
100th percentile service time	index-append	5554.97	ms
error rate	index-append	0	%
Min Throughput	index-stats	99.93	ops/s
Median Throughput	index-stats	100.04	ops/s
Max Throughput	index-stats	100.06	ops/s
50th percentile latency	index-stats	6.62305	ms
90th percentile latency	index-stats	7.35102	ms

Metric	Task	Value	Unit
99th percentile latency	index-stats	18.0909	ms
99.9th percentile latency	index-stats	24.5381	ms
100th percentile latency	index-stats	24.7431	ms
50th percentile service time	index-stats	6.50957	ms
90th percentile service time	index-stats	7.18652	ms
99th percentile service time	index-stats	9.38455	ms
99.9th percentile service time	index-stats	24.4424	ms
100th percentile service time	index-stats	24.6576	ms
error rate	index-stats	0	%
Min Throughput	node-stats	99.92	ops/s
Median Throughput	node-stats	100.04	ops/s
Max Throughput	node-stats	100.25	ops/s
50th percentile latency	node-stats	7.15655	ms
90th percentile latency	node-stats	7.96104	ms
99th percentile latency	node-stats	10.2362	ms
99.9th percentile latency	node-stats	25.7397	ms
100th percentile latency	node-stats	29.1573	ms
50th percentile service time	node-stats	7.04389	ms
90th percentile service time	node-stats	7.84655	ms
99th percentile service time	node-stats	9.13249	ms
99.9th percentile service time	node-stats	10.7357	ms
100th percentile service time	node-stats	29.072	ms
error rate	node-stats	0	%
Min Throughput	default	41.88	ops/s

Metric	Task	Value	Unit
Median Throughput	default	42.18	ops/s
Max Throughput	default	42.52	ops/s
50th percentile latency	default	3789.31	ms
90th percentile latency	default	5170.79	ms
99th percentile latency	default	5582.03	ms
99.9th percentile latency	default	5610.7	ms
100th percentile latency	default	5618.05	ms
50th percentile service time	default	23.1496	ms
90th percentile service time	default	25.8865	ms
99th percentile service time	default	33.249	ms
99.9th percentile service time	default	45.493	ms
100th percentile service time	default	62.4174	ms
error rate	default	0	%
Min Throughput	term	199.08	ops/s
Median Throughput	term	200	ops/s
Max Throughput	term	200.03	ops/s
50th percentile latency	term	5.02391	ms
90th percentile latency	term	21.18	ms
99th percentile latency	term	35.2251	ms
99.9th percentile latency	term	37.4827	ms
100th percentile latency	term	37.6907	ms
50th percentile service time	term	4.61812	ms
90th percentile service time	term	5.10619	ms
99th percentile service time	term	6.8135	ms

Metric	Task	Value	Unit
99.9th percentile service time	term	22.1183	ms
100th percentile service time	term	25.0033	ms
error rate	term	0	%
Min Throughput	phrase	199.61	ops/s
Median Throughput	phrase	200.04	ops/s
Max Throughput	phrase	200.8	ops/s
50th percentile latency	phrase	3.86572	ms
90th percentile latency	phrase	4.96583	ms
99th percentile latency	phrase	22.5681	ms
99.9th percentile latency	phrase	33.5684	ms
100th percentile latency	phrase	28.2658	ms
50th percentile service time	phrase	3.5689	ms
90th percentile service time	phrase	4.2535	ms
99th percentile service time	phrase	8.6957	ms
99.9th percentile service time	phrase	24.5685	ms
100th percentile service time	phrase	27.6584	ms
error rate	phrase	0	%
Min Throughput	country_agg_uncached	4.99	ops/s
Median Throughput	country_agg_uncached	5	ops/s
Max Throughput	country_agg_uncached	5	ops/s
50th percentile latency	country_agg_uncached	182.291	ms
90th percentile latency	country_agg_uncached	201.585	ms
99th percentile latency	country_agg_uncached	257.343	ms
100th percentile latency	country_agg_uncached	267.904	ms

Metric	Task	Value	Unit
50th percentile service time	country_agg_uncached	181.161	ms
90th percentile service time	country_agg_uncached	196.189	ms
99th percentile service time	country_agg_uncached	216.762	ms
100th percentile service time	country_agg_uncached	267.778	ms
error rate	country_agg_uncached	0	%
Min Throughput	country_agg_cached	99.95	ops/s
Median Throughput	country_agg_cached	100.05	ops/s
Max Throughput	country_agg_cached	100.07	ops/s
50th percentile latency	country_agg_cached	5.57249	ms
90th percentile latency	country_agg_cached	6.47982	ms
99th percentile latency	country_agg_cached	9.33674	ms
99.9th percentile latency	country_agg_cached	27.5319	ms
100th percentile latency	country_agg_cached	32.0567	ms
50th percentile service time	country_agg_cached	5.4601	ms
90th percentile service time	country_agg_cached	6.25153	ms
99th percentile service time	country_agg_cached	7.83564	ms
99.9th percentile service time	country_agg_cached	13.6439	ms
100th percentile service time	country_agg_cached	31.9487	ms
error rate	country_agg_cached	0	%
Min Throughput	scroll	25.01	pages/s
Median Throughput	scroll	25.03	pages/s

Strengths

Last updated : 2020-08-03 11:24:48

ES is hosted in the cloud, making it easy for you to create and manage Elasticsearch clusters and ensure high availability in production environments. Its core benefits are detailed below:

Ease of Deployment and Management

An ES cluster can be created in a few minutes without deploying software and hardware. Additionally, ES comes with a cluster management tool called Kibana, which assist cluster and alert systems to facilitate daily operation and management.

Elastic Scaling

ES has various node specifications and storage media for different business needs. Clusters can be scaled up or down to meet your current business needs and control costs.

Elasticsearch X-Pack

ES integrates Elasticsearch X-Pack, which has advanced features such as security, SQL, and machine learning to improve the efficiency of security management, usage, and OPS of Elasticsearch clusters.

High Availability

ES can be deployed in multiple availability zones, guaranteeing service continuity in the event of force majeure such as network or power failure in one single availability zone. A COS data backup policy can periodically back up the data to ensure rapid restoration in case of data loss due to unexpected conditions. In addition, ES boasts specially created policies such as kernel optimization that help comprehensively ensure data security and service stability.

Security Reinforcement

ES can be deployed in a logically isolated VPC, giving you full control over your environment configuration and the ability to customize network access control lists and security groups. It features a blocklisting/allowlisting mechanism for Kibana and IP access requests, and the security capability of X-Pack enables access control at the field level, helping ensure the security of your resources in the cloud.

Openness and Service Integration

ES supports the complete system of ELK products and is compatible with standard open-source RESTful APIs and ecosystem components. It can be integrated with Tencent Cloud products such as

COS, FL, CMQ, and TencentDB to implement data transfer and backup to meet your needs in different business scenarios.

Scenarios

Last updated : 2020-04-13 20:32:08

Log Analysis

Devices such as website servers, mobile devices, and IOT sensors can generate a high number of logs in various types that are stored on scattered nodes, which poses a great challenge to services relying on log search like troubleshooting and business analysis. ES provides a flexible, scalable, real-time, and centralized storage scheme and a full-text search feature to facilitate unified management and query of logs, helping you quickly identify and locate problems and improve the troubleshooting efficiency.

Full-Text Search

In-site search service for massive amounts of data such as search for ecommerce items, mobile apps, and organizational information is a necessary way to get the desired information efficiently. ES has a full-text search feature that can retrieve structured and unstructured data with ease. It also provides simple and convenient RESTful APIs and clients in various programming languages to help you quickly build a stable search service and integrate it into your existing business framework.

Business Intelligence (BI)

While data-driven operation is gaining popularity in today's industry landscape, businesses such as ecommerce, mobile app, advertising, and media need to rely on data analysis and data mining as assistance in business decision-making; however, large-scale business data poses great challenges to statistics collection and analysis. To cope with this problem, ES provides a structured query feature and supports complex filtering and aggregated statistics, which helps you efficiently perform statistical analysis on massive amounts of data, discover problems and opportunities, assist in business decision-making, and tap into the value of data.

Elastic Stack (X-Pack)

Last updated : 2020-09-11 09:50:40

Overview

X-Pack features are Elasticsearch's official commercial features, including security, SQL, machine learning, and monitoring. It facilitates the application development and OPS management of Elasticsearch services. ES offers editions that come with such features, which you can select when purchasing and creating a cluster. The features in different editions are detailed below.

Purchase Guide

The screenshot shows the purchase configuration interface. Under 'Elasticsearch Version', '5.6.4' and '6.4.3' are visible, with '6.4.3' selected. Under 'Advanced Features', 'Basic edition', 'Platinum edition', and 'Open source edition' are shown. 'Platinum edition' is selected and has a 'HOT' badge above it. A help icon (?) is present next to both sections. Below the options, a note states: 'Advanced features (X-Pack) include security, SQL, alarm, machine learning and many more. For details, see [Feature Introduction to Each Edition](#)'.

As shown in the figure above, there are options for the X-Pack features on the ES purchase page. ES offers three editions that have different X-Pack features as follows:

Item	Basic	Platinum	Open Source
X-Pack included	✓	✓	×
X-Pack completeness	Partial	All	None

Purchase recommendation

In order to be able to use more advanced features in ES, we recommend that you choose the **Platinum Edition** when you create a cluster. The specific features and differences of each edition are detailed below. For pricing information, please see [Product Pricing](#).

X-Pack Overview

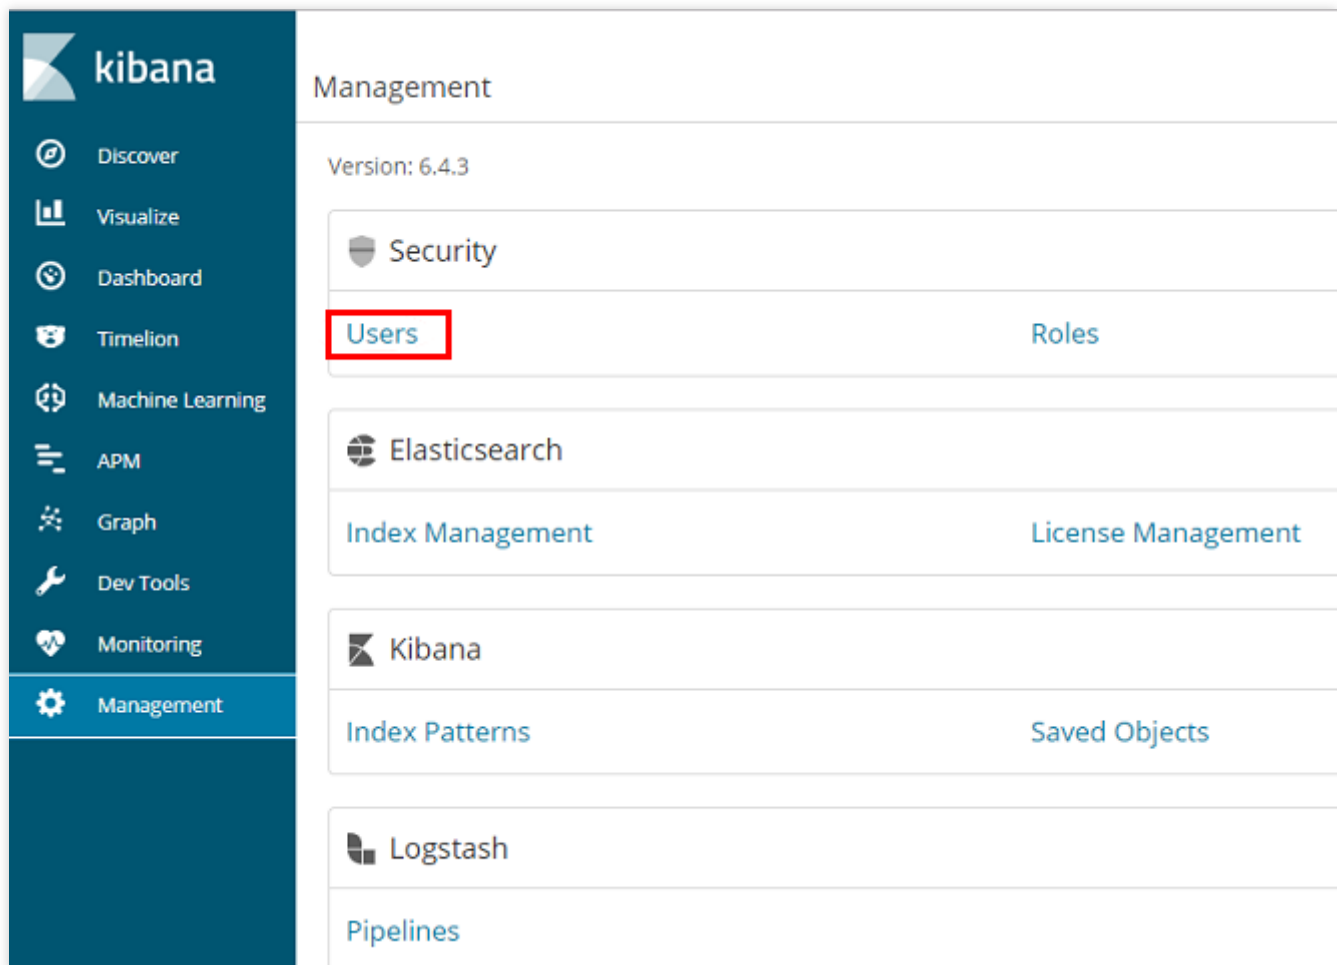
This document describes some of the commonly used X-Pack features. For more information, please see Elasticsearch's official [Elastic Stack subscriptions](#) and [API documentation](#).

Note :

- Some features vary by editions (Basic, Platinum, and Open Source).
- Some features are unavailable in earlier ES versions. For more information, please [submit a ticket](#).

Security

This feature supports refined read/write permission control at the index and field levels and effectively protects data security by enabling data security protection and business access isolation, granting access to the right people, and preventing malicious attacks and data leakage.

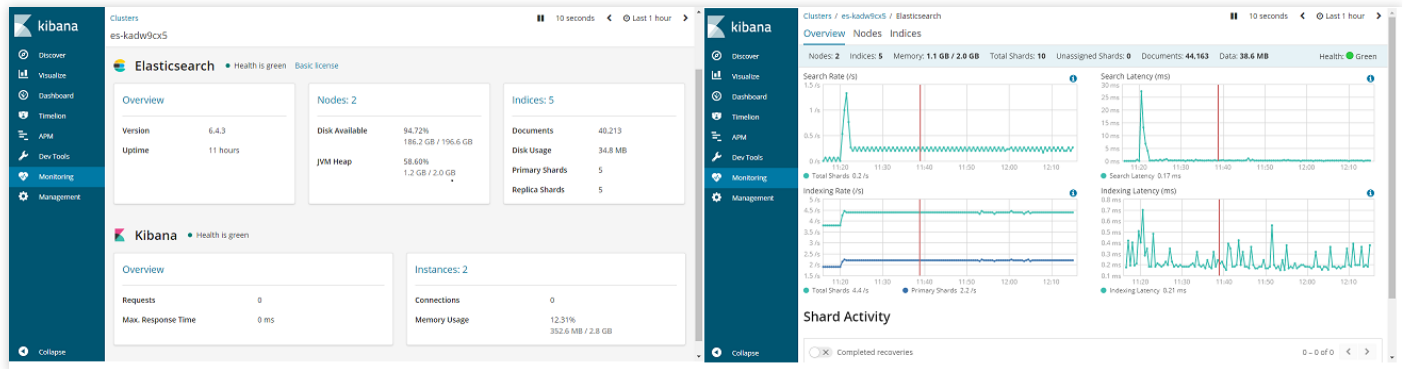
**Machine learning**

In the application scenario of custom data alerting, it is sometimes difficult to set rules and thresholds to define the changes. In this case, the trend in data changes and reasonable fluctuation range can be predicted by the unattended machine learning feature, and when the data deviates from the normal trend, alarms will be triggered and notifications sent.

Monitoring

Monitoring information can be comprehensively collected at multiple levels such as cluster, node,

and index, helping you understand the cluster operations in real time and facilitating your application development and OPS.



• SQL

This feature makes full-text search and statistical analysis of Elasticsearch data possible through traditional database SQL tools. CLI and REST access methods are supported. **The Platinum Edition further supports JDBC connection.** This feature enables you to seamlessly connect ES with your existing business systems and thus reduces your learning costs for new technologies.

```

1 POST _xpack/sql?format=txt
2 {
3   "query": "select * from website"
4 }

```


id	title
1	My first blog post
2	My second blog post
3	My first blog post
4	My second blog post
5	My first blog post
6	My second blog post
7	My first blog post
8	My second blog post
9	My first blog post
10	My second blog post
11	My first blog post
12	My second blog post
13	My first blog post
14	My second blog post
15	My first blog post
16	My second blog post
17	My first blog post
18	My second blog post

In terms of SQL support, the Open Source Edition integrates with other SQL plugins. For more information, please see [elasticsearch-sql](#).

Detailed comparison among editions

This section mainly compares and highlights some key features of different Elasticsearch versions. As Elasticsearch is in a stage of rapid development, and the support for various features by different versions is constantly adjusted, we do not guarantee that the following content can stay in sync with the changes in the community.

For the latest and most accurate feature comparison, please see Elasticsearch's official [Elastic Stack subscriptions](#).



In the table below, , , and are used to indicate the feature completeness. : all; :

partial; : none.

Module	Feature	Open Source	Basic	Platinum
Elasticsearch	Scalability and resiliency			
	Query and analytics			
	Data enrichment			
	Management and tooling			
	Security			
	Machine Learning			
Kibana	Explore and visualize			
	Stack management and tooling			
	Stack monitoring			

	Share and collaborate			
	Security			
	Machine learning			
Beats	Data collection			
	Data shipping			
	Module			
	Monitoring and management			
Logstash	Data collection			
	Data enrichment			
	Data shipping			
	Module			
	Monitoring and management			

ELASTIC APM	APM server	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	APM agents	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	APM dashboards in Kibana	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	APM UI	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Distributed tracing	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Machine learning integration	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Elastic Logs	Log shipper (Filebeat)	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Dashboards for common data sources	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Logs UI	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Elastic Infrastructure	Metric shipper (Metricbeat)	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Dashboards for common data sources	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Infrastructure UI	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Elastic Uptime	Uptime monitor (Heartbeat)	●	●	●
	Uptime dashboards in Kibana	●	●	●
	Uptime UI	—	●	●

Detailed descriptions of certain Elasticsearch features:

In the table below, ✓ means the feature is available, - means not available.

Elasticsearch Feature Module	Item	Open Source	Basic	Platinum
Management and Tooling	REST APIs	✓	✓	✓
	Language clients	✓	✓	✓
	Snapshot/restore	✓	✓	✓
	_source only snapshot	-	✓	✓
	SQL interpreter CLI	-	✓	✓
	Data rollups	-	✓	✓
	Index lifecycle management	-	✓	✓
	Frozen indices	-	✓	✓
	Upgrade Assistant APIs	-	✓	✓
	JDBC client	-	-	✓
	ODBC client	-	-	✓
Security	Encrypted communications	-	✓	✓
	Role-based access control	-	✓	✓

	File and native authentication	-	✓	✓
	Audit logging	-	-	✓
	Attribute-based access control	-	-	✓
	Field- and document-level security	-	-	✓
Machine Learning	Anomaly detection on time series	-	-	✓
	Population/entity analysis	-	-	✓
	Log message categorization	-	-	✓
	Root cause indication	-	-	✓
	Alerting on anomalies	-	-	✓
	Forecasting on time series	-	-	✓

Capabilities and Restrictions

Last updated : 2020-08-12 09:15:23

ES is a cloud-based PaaS service developed based on open-source Elasticsearch. It enables you to quickly build an Elasticsearch cluster service to develop applications such as log analysis and data search. The following describes its capabilities and use limits.

Product Composition

ES consists of ES cluster, its core component, and Kibana, its visual data analysis tool. For data collection and transfer to the ES cluster, you can deploy data collection tools such as Beats and Logstash or develop custom applications based on your own business needs.

Available Configurations

Node specification

Parameter Name	CPU Cores	Memory
ES.S1.SMALL2	1	2 GB
ES.S1.MEDIUM4	2	4 GB
ES.S1.MEDIUM8	2	8 GB
ES.S1.LARGE16	4	16 GB
ES.S1.2XLARGE32	8	32 GB
ES.S1.4XLARGE32	16	32 GB
ES.S1.4XLARGE64	16	64 GB

The node with 1 core and 2 GB of memory is intended for testing purpose only and is not recommended for production environments.

Storage

SSD cloud disks are used. The disk capacity is 100 GB-6 TB on a single node.

Number of nodes

The number of nodes is limited to 2-50. As an ES cluster typically consists of nodes deployed in a distributed manner, a master node is required to manage it. In order to prevent the risk of split brain caused by possible node failures, you are recommended to select at least three nodes for a cluster.

Configuration selection

See [Evaluation of Cluster Specification and Capacity Configuration](#).

Network Access

Private network access in VPC

In order to ensure data security, ES is built in your VPC, and you can only access an ES cluster from your VPC to write and query data. If you need to access a cluster over the public network for development and debugging purposes, you can connect your local IDC to the VPC using [VPN](#). In this case, please take effective measures to protect your data.

Kibana page

You can access the Kibana page over the public network. For the sake of data security, a password and access IP blocklist/allowlist need to be set for the Kibana page.

VPC network selection

Once an ES cluster is created, its VPC network cannot be changed; therefore, please make a good plan for your business deployment in advance when creating a cluster.

Related Concepts

Last updated : 2020-02-26 22:15:13

An Elasticsearch cluster is generally a distributed one consisting of multiple nodes. The nodes communicate and cooperate with one another to provide searching and indexing services (client requests can be forwarded to the optimal node among all the nodes). Different nodes play one or more different roles. There are many node roles in Elasticsearch, and the most important two ones are data nodes and master nodes.

Data Node

It is mainly responsible for operations related to the storing, processing, and manipulating of data and index shards, such as I/O-, memory-, and CPU-intensive operations like CRUD, search, and aggregation. During the use of cluster, you should closely monitor the resource utilization of the data nodes and ensure cluster stability by adding more nodes to scale the cluster up when the service is overloaded.

Master Node

It is responsible for making cluster-wide operations lightweight, such as creating or deleting indices, tracking which nodes are part of which clusters, and deciding which shards to assign to which nodes. It is important to have a stable master node for the cluster health.

Master-eligible Node

This refers to a node that is eligible to be selected as a master node. Any node that meets the requirements for a master node (all nodes by default) can be selected as a master node through the master selection process.

By default, all nodes are data nodes and eligible to be a master node, which is very convenient for small clusters. Because the requests for index processing and data searching are I/O-, memory-, and CPU-intensive for data nodes, they may cause pressure on the node resources. As the cluster grows, in order to ensure that the master nodes are stable and free from pressure and to ensure the cluster stability, the master nodes should be separated from the data nodes.

Dedicated Master Node

This is a node set to serve only as a master node in an Elasticsearch cluster.

Suggestions on Dedicated Master Nodes

Configuring dedicated master nodes is mainly to ensure the stability of the cluster as it grows. It is recommended to configure at least 3 dedicated master nodes:

- If the number of dedicated master nodes is 1, there is only one eligible master node. `discovery.zen.minimum_master_nodes` can only be set to 1, and there is no backup in case of network failure.
- If the number of dedicated master nodes is 2, there are 2 eligible master nodes. If `minimum_master_nodes` is set to 1, although there is a backup node, there may be a risk of split-brain (i.e., each eligible master node sets itself as the master node) when the master node is re-selected in case of network failure. If `minimum_master_nodes` is set to 2, as the number of eligible master nodes falls short, no master node can be selected in case of failure.
- If the number of dedicated master nodes is 3, there are 3 eligible master nodes. If `discovery.zen.minimum_master_nodes` is set to 2, even if one eligible master node is lost in case of network failure, there is still one master node that can be re-selected.

For more information, see [ES Node Description](#).