Elasticsearch Service

Release Notes

Product Documentation

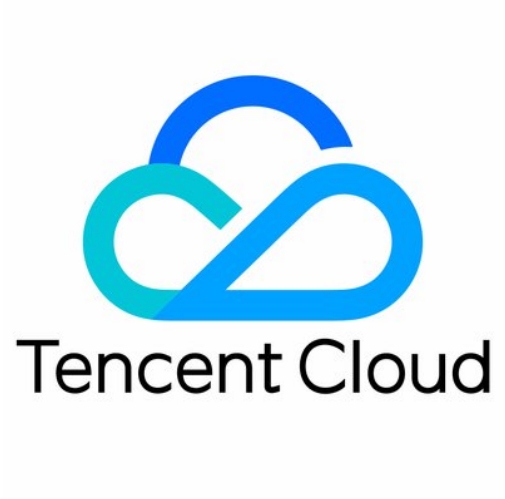

Release Notes

Last updated：2022-06-28 21:28:05

June 2021

Update	Description	Release Date	Documentation
Supported Kibana node specification customization	You can purchase Kibana node models with different specifications (nodes with 1 CPU core and 2 GB memory are free of charge) for use in scenarios with massive data analysis and export tasks that require a high Kibana performance. In addition, ES supports high Kibana availability across AZs.	2021-06-25	Creating Clusters
Launched the Enhanced SSD model	The Enhanced SSD is supported, which is available only with Star Lake servers and suitable for I/O-intensive scenarios with high latency requirements.	2021-06-25	Pricing
Supported monitoring dedicated master nodes	The performance of dedicated master nodes can be monitored, making it easy for you to stay up to date with their running status.	2021-06-25	Viewing Monitors
Enabled slow log by default	Starting from June 25, 2021, slow log is enabled for newly purchased clusters by default to capture slow search logs and slow index logs.	2021-06-25	Querying Cluster Logs

May 2021

Update	Description	Release Date	Documentation
Launched in Beijing Zone 6 and Guangzhou Zone 7	Clusters can be created and managed in Beijing Zone 6 and Guangzhou Zone 7.	2021-05-22	Creating Clusters
Supported the Korean analysis plugin	The nori analyzer provided by Elasticsearch for Korean is supported, which is suitable for full-text search and analysis of Korean documents.	2021-05-14	Plugin List

March 2021

Update	Description	Release Date	Documentation
Supported Elasticsearch 7.10	Elasticsearch 7.10 is supported.	2021-03-01	What's new in 7.10
Supported the visual cluster architecture view	The architecture view straightforwardly displays the cluster deployment status (node type, number of nodes, etc.) to provide an overview of the cluster and node running status (normal, offline, warning, etc.). It also offers monitoring and management capabilities.	2021-03-01	-
Launched the High I/O model	High I/O is a type of CVM instance with a large-capacity local SSD disk and high read/write performance. It is suitable for scenarios with high requirements for read/write performance and disk capacity.	2021-03-01	Pricing
Launched in the Thailand region	Clusters can be created and managed in the Thailand region.	2021-03-01	Creating Clusters

January 2021

Update	Description	Release Date	Documentation
Supported heteronyms in the pinyin plugin	The support for heteronyms is optimized for user-friendly input. For example, the original analysis results of a username are "cengmoumou" and "cmm", and the optimized results are "cengmoumou", "zengmoumou", "cmm", and "zmm".	2021-01-27	-

December 2020

Update	Description	Release Date	Documentation

Increased the storage upper limit of a single cloud disk to 30 TB	For clusters on v6.8 or above, the storage capacity of a single cloud disk can be increased to 30 TB, which helps reduce the number of nodes and the cluster creation costs while using the same storage disk specification.	2020-12-30	-
Supported switching the client request node type	In hot/warm mode, you can switch the read/write traffic of CLB to only hot nodes so as to improve the overall read/write performance of your clusters.	2020-12-30	-
Increased the storage upper limit of a single cloud disk to 30 TB	For clusters on v6.8 or above, the storage capacity of a single cloud disk can be increased to 30 TB, which helps reduce the number of nodes and the cluster creation costs while using the same storage disk specification.	2020-12-30	-
Supported switching the client request node type	In hot/warm mode, you can switch the read/write traffic of CLB to only hot nodes so as to improve the overall read/write performance of your clusters.	2020-12-30	-
Supported the Standard SA2 (Star Lake) model	Also known as "Star Lake", Standard SA2 is Tencent's proprietary server model. It is optimized for cloud-based scenarios, with its chips deeply customized and optimized to deliver an ultra-high performance and stability. The prices of Standard SA2 are 30% lower than Standard S1 with comparable performance.	2020-12-01	Pricing
Supported the Big Data model	The Big Data model is a type of CVM instance with large-capacity local disks mounted to it. It provides a larger local storage capacity at lower costs to help reduce the creation costs in various scenarios with high storage usage and low access frequency, such as logs.	2020-12-01	Pricing
Supported the MEM Optimized model	Featuring a CPU to memory ratio of 1:8, the MEM Optimized model is suitable for use cases involving high memory utilization, such as massive aggregate analysis.	2020-12-01	Pricing
Supported Chinese and configuration items such as `timeout` in Kibana	You can switch the UI language of Kibana (Chinese and English) and modify the `timeout` parameter.	2020-12-01	-

October 2020

Update	Description	Release Date	Documentation
Launched in the Russia region	Clusters can be created and managed in the Russia region.	2020-10-26	Creating Clusters

September 2020

Update	Description	Release Date	Documentation
Supported three-AZ cluster deployment	In addition to dual-AZ cluster deployment, three-AZ cluster deployment is also supported, which helps improve the disaster recovery capability of your clusters.	2020-09-28	Multi-AZ Cluster Deployment
Supported YML configuration customization	A YML customization editor is provided to flexibly configure more YML parameters.	2020-09-27	YML File Configuration
Supported selecting rolling mode and blue/green mode for cluster upgrade	You can select the upgrade mode based on your business conditions: • In rolling mode, nodes in the cluster are restarted one by one and quickly upgraded on a rolling basis without any interruption to the system service, but the online performance may be affected. • In blue/green mode, the same number of new nodes as the existing nodes are added to the original cluster with no cluster restart required. This upgrade process is seamless and smooth but time-consuming.	2020-09-27	Suggestions and Principles for Cluster Specification Adjustment
Supported single node restart	When a single node fails, you can restart it rather than the entire cluster to solve the problem.	2020-09-27	-
Optimized cluster monitoring	Cluster and node monitoring metrics are adjusted, metric grouping is standardized, and chart display is optimized based on actual OPS scenarios, making monitoring easier to use.	2020-09-27	Viewing Monitors

Supported selecting scenario-based templates	The default index templates for the common use cases of ES are provided, which help optimize the cluster and index configurations and reduce the cluster exceptions and performance issues caused by improper use.	2020-09-27	Scenario-based Cluster Template Configuration
Lowered disk pricing	Disk prices are lowered in certain regions.	2020-09-27	Elasticsearch Service Price Reduction Notice

July 2020

Update	Description	Release Date	Documentation
Supported presetting the plugin list	Over 10 open-source and proprietary mature plugins that provide rich features are supported, including IK Analyzer and Smart Chinese Analysis. You can install and uninstall them based on your business needs.	2020-07-17	Plugin List
Supported the QQ analysis plugin	Developed by Tencent's ES team in collaboration with the NLP team, the QQ analysis plugin is widely used in Tencent businesses such as QQ, WeChat, and QQ Browser. On the basis of traditional dictionary-based analysis, it supports named-entity recognition (NER) and custom dictionaries. It has become industry-leading on key metrics such as analysis accuracy and speed.	2020-07-17	QQ Analysis Plugin
Supported synonym configuration	You can upload synonym files, which helps simply synonym library configuration.	2020-07-17	Synonym Configuration
Supported selecting rolling mode and blue/green mode for cluster configuration adjustment	You can select the configuration adjustment mode based on your business conditions: • In rolling mode, nodes in the cluster are restarted one by one and quickly adjusted on a rolling basis without any interruption to the system service, but the online performance may be affected. • In blue/green mode, the same number of new nodes as the existing nodes are added to the original cluster with no cluster restart required. This	2020-07-17	Suggestions and Principles for Cluster Specification Adjustment

	configuration adjustment process is seamless and smooth but time-consuming.		

June 2020

Update	Description	Release Date	Documentation
Launched in Guangzhou Zone 6	Clusters can be created and managed in Guangzhou Zone 6.	2020-06-30	Creating Clusters
Launched in the Tokyo and Virginia regions	Clusters can be created and managed in the Tokyo and Virginia regions.	2020-06-09	Creating Clusters

May 2020

Update	Description	Release Date	Documentation
Integrated X-Pack features	New clusters support advanced X-Pack features such as alerting, LDAP authentication system, cross-cluster search (CCS), and cross-cluster replication (CCR) (to use such features, previously created clusters need to be upgraded).	2020-05-19	-
Supported editing tags	You can modify the cluster tags on the details page, which makes it easier for you to flexibly categorize and manage resources.	2020-05-19	-

April 2020

Update	Description	Release Date	Documentation
Lowered the product prices	ES prices are lowered overall.	2020-04-10	Pricing

Supported models with higher specifications	High-Specced models are supported, including Standard S1 with 24 CPU cores and 48 GB memory, Standard S1 with 24 CPU cores and 96 GB memory, Standard S1 with 32 CPU cores and 128 GB memory, and Standard S1 with 48 CPU cores and 96 GB memory.	2020-04-10	Pricing
Improved the disk upgrade efficiency	The disk capacity of CVM instances can be expanded directly, which significantly reduces the amount of time it takes to expand the disk capacity and improve the process reliability.	2020-04-10	-

February 2020

Update	Description	Release Date	Documentation
Supported v7.5	Elasticsearch 7.5 is supported.	2020-02-21	7.5.0 release highlights
Supported multi-AZ deployment	Multi-AZ deployment can guarantee service continuity in the event of force majeure such as network or power failure in one single AZ, thus improving the cross-data center disaster recovery capability of your business.	2020-02-21	Multi-AZ Cluster Deployment
Supported ES upgrade	ES can be upgraded from a low version to a high version and also to advanced features (Open Source - Basic - Platinum). You can upgrade your clusters based on your business needs to get the latest features while ensuring seamless business transition.	2020-02-21	Upgrading ES Clusters
ES clusters now support public network access	Clusters with user authentication enabled can be accessed over the public network, but an IP allowlist needs to be configured for security protection.	2020-02-21	ES Cluster