

Aerospike® Connect for Elasticsearch

PRODUCT BRIEF

Enterprise search on real-time data using Aerospike Connect for Elasticsearch

Overview

Aerospike Connect for Elasticsearch enables the development of enterprise text search use cases for data residing in the Aerospike Real-time Data Platform. It enables users to easily and seamlessly integrate Elasticsearch - the most powerful open-source search engine - with Aerospike's high-performance, scalable NoSQL database.

Connect for Elasticsearch enables developers to create search applications on data in Aerospike, making it easier and faster to analyze and gain insights from their real-time data. The connector uses Aerospike XDR to asynchronously transport data from Aerospike into Elasticsearch clusters using change notifications.

Key Features

Aerospike Connect for Elasticsearch provides enterprise grade search capabilities to the Aerospike database using the world's most popular and powerful search engine

- Powerful full-text search enables developers to employ the most widely adopted open-source full-text search platform in Elasticsearch
- Efficiency through the use of Aerospike change notifications which minimizes the amount of data that needs to be loaded into Elasticsearch clusters
- Low latency data distribution to the Elasticsearch cluster using Aerospike's Cross Datacenter Replication (XDR)
- Search on text, metadata, JSON, and more - Developers have the flexibility to search text data and documents stored in Aerospike clusters

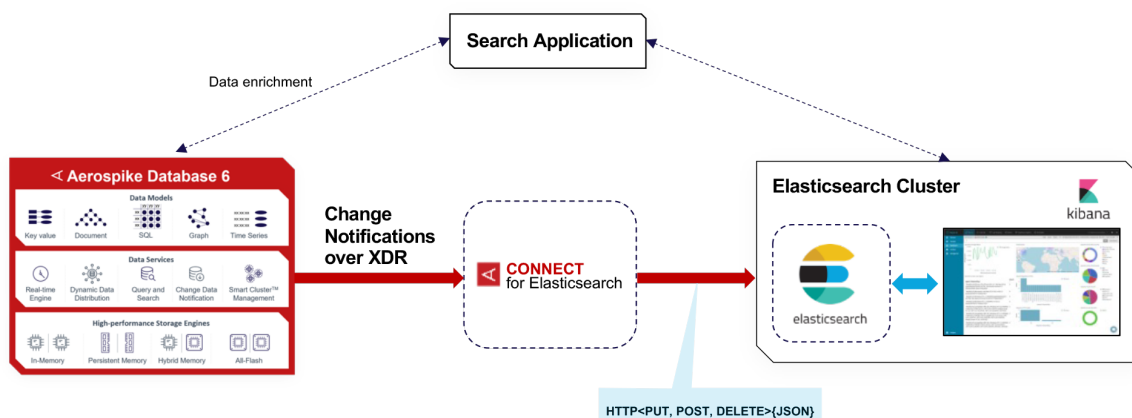


Figure 1: Connect for Elasticsearch: best of breed real-time data platform and search

How Aerospike Connect for Elasticsearch Works

Aerospike Connect for Elasticsearch is a connector between the two technologies using Aerospike's Cross Datacenter Replication (XDR) and Change Notification services. It provides a seamless and low-latency integration between Aerospike and Elasticsearch, enabling customers to use Elasticsearch for full-text search and other advanced querying capabilities on data in Aerospike. The image above illustrates the deployment architecture for the connector.

The connector subscribes to Change Notifications for mutations to data and asynchronously replicates data in Aerospike and makes it available in an Elasticsearch index in near real time. Additionally, the fine replication granularity of XDR allows developers to choose the namespace, set or bin to replicate in Elasticsearch depending on requirements around data availability, governance or resource usage. For instance, by configuring only the searchable bins in a record to be shipped to an Elasticsearch index you can efficiently use memory resources for your Elasticsearch cluster.

Our detailed technical documentation covers all this and more on the usage and configurability of the connector.

Use Case

Ecommerce and Retail

Retailers increasingly want to search the data they gather about their customers including social data, buying patterns, comments and reviews

Personalization and relevance

Customer and public service functions increasingly rely on analytics and search to improve the accuracy and relevance of their products and services.

User behavior and experience

Industries from ecommerce to banking increasingly need to search customer or user behavior based on documents and other textual data

ML model training

Model training requires fast and efficient data retrieval and indexing for large datasets, enabling more accurate predictions and improved performance.

Observability and log analytics

Observability solutions and industrial systems generate an ever increasing amount of JSON document and textual log data that can be searched to find anomalies

AEROSPIKE

The Aerospike Real-time Data Platform enables organizations to act instantly across billions of transactions while reducing server footprint by up to 80 percent. The Aerospike multi-cloud platform powers real-time applications with predictable sub-millisecond performance from terabytes to petabytes of data with five nines uptime with globally distributed, strongly consistent data. Applications built on the Aerospike Real-time Data Platform fight fraud, provide recommendations that dramatically increase shopping cart size, enable global digital payments, and deliver hyper-personalized user experiences to tens of millions of customers. Customers such as Airtel, Experian, Nielsen, PayPal, Snap, Wayfair and Yahoo rely on Aerospike as their data foundation for the future. Headquartered in Mountain View, California, the company also has offices in London, Bangalore and Tel Aviv.

©2023 Aerospike, Inc. All rights reserved. Aerospike and the Aerospike logo are trademarks or registered trademarks of Aerospike. All other names and trademarks are for identification purposes and are the property of their respective owners.