

Aerospike Connect for Kafka

PRODUCT BRIEF

Real-time processing of Kafka data streams

Overview

Kafka has become the default streaming data transport technology for data in motion across virtually all modern data pipelines. As low latency connectivity between these diverse systems is critical in enterprises adopting digital transformation, real-time data movement from the edge to the core and back is becoming crucial. For this reason, a strong integration of Kafka data streams and the Aerospike Real-time Data Platform delivers a perfect marriage of a reliable real-time data transport with a high-performance data persistence and data processing engine.

Connecting Aerospike with enterprise systems through Kafka

Aerospike Connect for Kafka leverages the cross-datacenter replication (XDR) protocol in Aerospike Database 6 to achieve sub-millisecond latencies from gigabytes to petabyte scale in Kafka-based streaming data pipelines. Kafka provides a transport that enables applications to publish and subscribe data through named topics. This integration enables Aerospike Database 6 to pull data from Kafka and to submit Aerospike change notifications via the Aerospike XDR protocol to Kafka. In addition, other applications and data stores can subscribe and publish to Aerospike Database clusters via Aerospike Connect for Kafka.

Highlights

- Powerful Change Data Notifications informs external systems and CDC applications when records are modified, inserted, or deleted from a data stream.
- Provides a high-performance persistence engine for applications to stream data in and out of Aerospike clusters.
- Seamless integration between Aerospike, Kafka, and other enterprise systems, including Spark, Pulsar and Trino.
- Enables event-driven and microservices architectures.

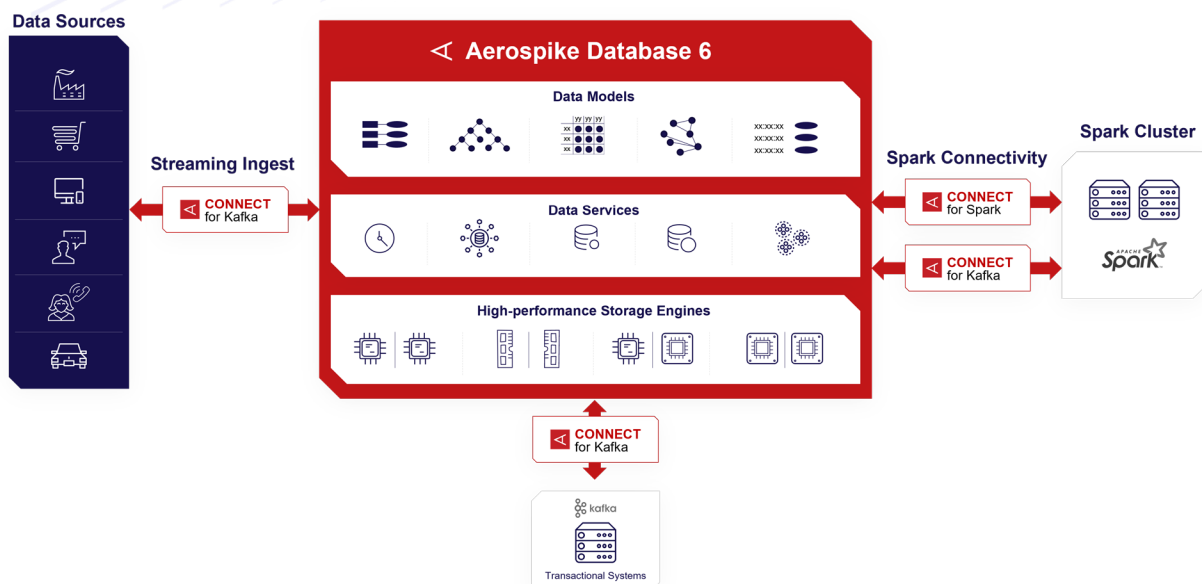


Figure 1: Aerospike Connect for Kafka

Aerospike Connect for Kafka

Aerospike Connect for Kafka supports bi-directional movement of data through messaging with Kafka enabling enterprises to deploy and integrate Aerospike seamlessly with other systems.

Key Features

- Stream change notifications from Aerospike using the XDR protocol to streaming data applications and analytics.
- Supports popular message formats such as Apache Avro, binary (MessagePack) and flat JSON, making application integration easier.
- Ability to route messages based on namespaces, sets, bin values, and static topic names from Aerospike to the downstream application via Kafka.
- Transform messages going into and out of Aerospike to downstream applications via Kafka.
- Rapid, parallelized ingest of data into Aerospike from streaming applications without worrying about data models compatibility between Aerospike and source systems.

Benefits

- Enterprise wide connectivity of Aerospike utilizing the Kafka as a message bus
- Enables event driven architecture & microservices architectures
- Ability to integrate Aerospike with Kafka sinks and sources that are not directly supported by Aerospike
- Make Aerospike a part of your high throughput and low latency streaming pipeline
- Lower TCO by leveraging the cost efficiencies of Aerospike and Kafka
- Deploy Anywhere

Typical Use Cases

Financial Services and FinTech

Bridge legacy systems such as mainframes to modern services such as fraud detection, risk management, and customer 360 using best of breed Kafka streaming and the Aerospike Real-time Data Platform

e-Commerce and Retail / CPG

Faster and more reliable use of behavioral, social, recommendation data as well as clickstream integration with product data.

Industrial Internet

Edge and device data synchronization with back-end systems.

Telecommunications

Stream data from Aerospike (to be used for high velocity, high volume ingest at an IoT edge system) to back-end core systems. Persist the data, train AI/ML models, and perform real-time analysis.

Data Lake / EDW Integration for real-time Analytics

Integrating operational data with data lakes and data warehouses for the broadest set of real-time analytics use cases.

AdTech

Real-time clickstream data synchronization and integration.

AEROSPIKE

Aerospike unleashes the power of real-time data to meet the demands of The Right Now Economy. Global innovators and builders choose the Aerospike real-time, multi-model, NoSQL data platform for its predictable sub-millisecond performance at unlimited scale with dramatically reduced infrastructure costs. With support for strong consistency and globally distributed, multi-cloud environments, Aerospike is an essential part of the modern data stack for Adobe, Airtel, Criteo, DBS Bank, Experian, PayPal, Snap, Sony Interactive Entertainment, The Trade Desk, and Wayfair. A global company, Aerospike is headquartered in Mountain View, California, with offices in London, Bangalore, and Tel Aviv.

©2022 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.