

Aerospike® Database 6

PRODUCT BRIEF

Real-time, multi-model database delivers millions of transactions per second with strong consistency at unlimited scale

Changing customer expectations and industry disruptors are driving significant investments in digital transformation initiatives across multiple industries. To compete in today's real-time digital economy, companies must invest in the right gigabyte-to-petabyte scale data platform to enable the right decisions in the moments that matter. This is only possible by utilizing modern data architectures which support predictable performance at scale. A modern data platform consists of the following components: System of Engagement (SOE) databases that capture real-time data from edge and near-edge devices; System of Record (SOR) databases which store both the real-time data from the SOE database and historical data which acts as the single source of truth; globally distributed databases which can span across geographical regions and can process billions of records daily with no data loss; and analytic, AI, and ML databases that leverage the combined data. Overall, high-speed data transfer between these components is required as they enable real-time decisioning in the microsecond moments that matter.

New in Aerospike Database 6.0/6.1

- **Document data model** – comprehensive JSON document database capabilities.
- **Batch processing** – for all workloads - write heavy, read heavy, and mixed.
- **Secondary indexes for nested elements** – of document data for maximum performance and scale.
- **Massively parallel secondary indexes** – for real-time complex SQL queries at GB/PB scale.
- **FIP 140-2 compliance** – for public sector clients and high security use cases.
- **Warm start for secondary indexes** – enables developers to place indexes in shared memory for fast restart.

Highlights

Lowest latency

Sub-millisecond performance for the day's first transaction to the billionth, and beyond.

Unlimited scale

Predictable performance from gigabytes to petabytes of data. Build your real-time, mission-critical applications with peace of mind and adherence to the strictest of SLAs.

Industry-leading uptime

Five-nines uptime with globally distributed, strongly consistent data. Aerospike has customers with uptime for years on end. "Aerospike just works."

Reduce costs, grow business

Decrease energy consumption and carbon footprint with up to 80 percent server or cloud instance reduction, even as your data grows.

Aerospike Real-Time Data Platform

Aerospike provides predictable performance up to petabyte scale, with five-nines uptime with always accurate, global, consistent data for all components of the end-to-end platform (Figure 1). The platform consists of these elements:

- **Multi-model database** – Functions as an unlimited scale Key-Value, Document, Graph or Time Series database.
- **Edge database (SOE)** – Used for real-time decisioning based on local streaming and transactional data plus historical data pulled dynamically from the SOR.
- **Synchronous Active-Active replication** – Using cross-datacenter replication and strong consistency, Aerospike can be deployed across multiple geographically separated data centers with high resiliency, automated failovers, and no loss of data.
- **Real-time system of record (SOR)** – Stores transactional and historical data and pushes data as needed to the SOEs also powering ML and AI-based applications.
- **Aerospike Cross Datacenter Replication (XDR)** – Enables multiple geographically dispersed data centers to stay in sync through high performance replication.
- **No-code integrations** – Enables companies to directly integrate the Aerospike Database with Spark, Presto-Trino, Pulsar, JMS and more.

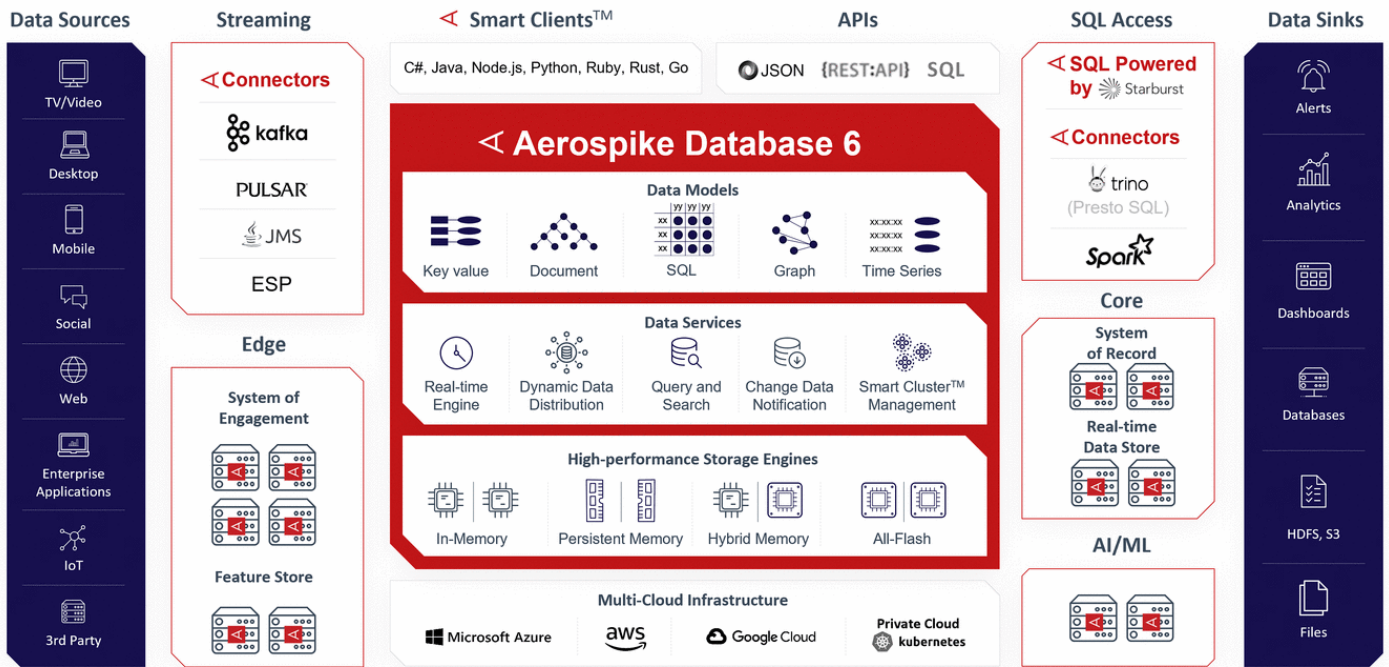


Figure 1. The Aerospike Real-time Data Platform

Meeting gigabyte-to-petabyte scale enterprise data management requirements

The Aerospike Real-time Data Platform, powered by the Aerospike Database 6 delivers millisecond response times at scales of gigabytes-to-petabytes for the strictest enterprise requirements:

Predictable performance

Low latency and high throughput- Multi-threaded parallel processing at the CPU combined with our patented Hybrid Memory Architecture (Figure 2) designed for flash storage devices uniquely deliver predictable high performance at scale.

Support for next-generation memory - Aerospike offers storage options for optimal density with high performance, starting with our patented Hybrid Memory Architecture (HMA). HMA stores indexes on DRAM and data on Flash. Aerospike HMA treats SSDs like DRAM, accessing them as raw devices at DRAM speeds with great parallelism.

Our storage engine approach has been extended to and optimized for All-Flash and persistent memory configurations, and we offer an in-memory option as well. See figure 2.

Smart Client™ architecture - Aerospike's Smart Client™ architecture ensures parallel access to multiple servers in a cluster for the highest possible performance.

Real-time data transfer between edge and core - Aerospike XDR enables multiple geographically dispersed data centers to stay in sync through high performance replication.

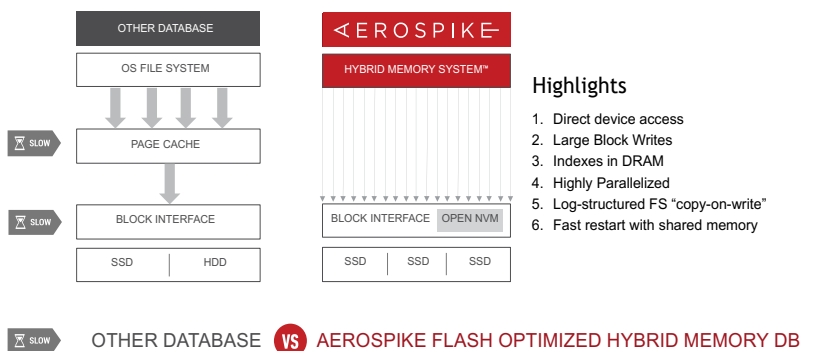


Figure 2: Aerospike's patented Hybrid Memory Architecture™

Multi-model engine

Document database - New in Aerospike Database 6, Aerospike Document Data Services provides comprehensive JSON document database capabilities including the storage, processing, and querying of JSON document data in real-time and at scales that have previously been unavailable in other databases and data platforms.

Scalability - The Aerospike Hybrid Memory Architecture with All Flash and Hybrid Flash options, coupled with Dynamic Cluster Management, allows the Aerospike database to scale to petabytes and store transactional/streaming/real-time data as well as historical data.

Compression - Aerospike's storage compression feature provides lossless compression of records written to persistent storage. Additionally, the communication between the Clients and the database is also compressed.

Rich data types - Scalar data types: Integer, Double, String, BLOB, Bytes; Collection data types: List (Ordered and nonordered), Map (Key-Ordered, Key Value-Ordered, Unordered); Probabilistic data types: HyperLogLog, HyperMinHash; Geospatial: GeoJSON data type.

Change notification framework - Allows Aerospike servers to stream changes to other systems via Aerospike Connect components. Allows for building easy yet reliable and scalable systems for complex event processing (CEP).

Query and Search

Secondary indexes - New in Aerospike Database 6 - secondary indexes are now available at the partition level for massive parallelism and improved performance and scalability.

Aerospike Expressions - Execute operations and functions closer to the data for more efficient comparisons of values, even extending the ability to read, write, and execute complex computations as arguments in other operations.

Set indexes - Provide efficient access to a Set within an Aerospike Namespace. This

feature allows fast queries of records within a Set in a petabyte-scale database.

Durability/Consistency

Durability - Data can be replicated asynchronously across geographies and synchronously written to other nodes in the cluster and to flash storage without disruption for the highest durability.

Consistency - Aerospike provides strong consistency on primary key access that has been confirmed through Jepsen test results. Data held in Aerospike is always guaranteed to be correct in all scenarios. Aerospike Multi-site Clustering brings the Jepsen validated Strong Consistency to deployments across multiple sites.

High Availability/Replication

Uptime and high availability - Aerospike provides high availability and a demonstrated uptime of five 9s or more which is made possible by our unique cluster management and intelligent client technology.

Synchronous data replication - Aerospike Multi-Site Clustering supports always-on, strongly consistent, globally distributed transactions at scale. It provides a true real-time Active-Active solution with great resiliency on WANs.

Asynchronous data replication - XDR delivers expressions-based fine-grain control of asynchronous replication of data across geographically distributed clusters. It can be used to create a global data hub, allowing to route and augment data captured at the edge to other clusters.

Enterprise Security

Encryption - Aerospike supports full transport encryption, as well as in-database transparent data encryption.

Authentication - LDAP and Kerberos authentication mechanisms are supported. Rich sets of access control options are available including ACLs.

Authorization - Aerospike provides a sophisticated role-based access control (RBAC) system.

Auditing - Aerospike can be configured to generate audit log messages on a wide variety of security events.

Centralized Secret Management option - Aerospike allows the following security items to be managed by and stored within HashiCorp Vault's KV secrets engine:

- LDAP user credentials and TLS certificates
- XDR remote destination passwords
- Encryption-at-rest key
- Network TLS certificates and keys

Integrations

Integration with existing data stores and systems - for building modern data pipelines and powering highly scalable/low latency AI/ML applications. The Aerospike Connect product line is currently composed of:

SQL Access

- Aerospike SQL Powered by Starburst
- Aerospike Connect for Spark
- Aerospike Connect for Presto-Trino

Streaming Data Connectors

- Aerospike Connect for Kafka
- Aerospike Connect for Pulsar
- Aerospike Connect for JMS
- Aerospike Connect for Event Stream Processing (ESP)

Deployment Options

In data centers and private clouds

In public cloud - Google Compute Platform, Amazon Web Services, Microsoft Azure, Alibaba Cloud and others.

Aerospike Cloud Managed Service

Aerospike experts deliver and maintain an optimized deployment of the Aerospike database in cloud environments with white glove service.

Database Product Features

License

Aerospike Server License Type - Commercial License
 Aerospike Client License Type - Apache v2

Support

Binaries - Tested & Verified
 Enterprise Production Support
 Hot Patch Availability
[Community Support](#)

Architecture

[Multi-site Clustering](#)
[Strong Consistency](#)¹
[Rack Awareness](#)

Storage Engine

[All Flash](#)¹
[Persistent Memory support](#)

Performance

[Unlimited Transactions/Queries per Second](#)
[Namespaces 32 max](#)
[Objects per Namespace per Node 2 \(0.5 Trillion max\)](#)

Security

[FIPS 140-2](#)
[TLS Transport Encryption](#)
[Access Control Lists \(ACLs\)](#)
[Data Encryption-at-Rest](#)¹
[LDAP Authentication](#)¹

Operations

[Backup & Restore](#)
[Cross Datacenter Replication \(XDR\)](#)
[Rate Quotas](#)
[Change Notification](#)
[Rapid Rebalance](#)
[Uniform Balance](#)
[Delay Fill Migrations](#)
[Quiescence](#)
[Compression](#)¹
[Fast Restart](#)
[Durable Delete](#)
[Read Page Cache](#)
[IPv6](#)

Query and Search

[Aerospike SQL Powered by Starburst](#)¹
[Secondary indexes](#)
[Aerospike Expressions](#)
[Set Indexes](#)

Aerospike Connect

[Aerospike Connect for Spark](#)¹
[Aerospike Connect for Kafka](#)¹
[Aerospike Connect for JMS](#)¹
[Aerospike Connect for Pulsar](#)¹
[Aerospike Connect for Presto-Trino](#)¹
[Aerospike Connect for Event Stream Processing \(ESP\)](#)¹

¹ Available with additional licensing

² See [Known Limitations](#) for more information

For detailed discussion of Aerospike Database features and editions, please see:

<https://aerospike.com/products/features-and-editions/>



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.