

WHY IS AEROSPIKE'S HYBRID MEMORY ARCHITECTURE DIFFERENT?

Highlights

- Combines strong consistency with high performance
- Delivers elastic, predictable performance at internet scale
- SSD optimization yields massive performance benefits
- Massively parallel architecture allows linear scalability
- Eliminates complex, costly caching layers
- Applications become simpler, more maintainable and faster to market
- Self-healing clusters with identical nodes gives high availability

High Performance, Enterprise-Grade Key-Value Store Database

More and more enterprise applications are driven by fast decisions closer to data and hence are architected as edge or fog systems. For edge-based systems that require high throughput and extremely low latency over large volumes of data, traditional architectures call for a large DRAM cache in front of a persistent store. It is believed that anything else simply will not perform. However, Aerospike's Hybrid Memory Architecture™ (HMA) offers a fundamentally different approach. By storing information on fast SSD devices with only indexes stored in DRAM, many advantages are realized without compromising on the performance.

Aerospike's HMA provides lower latency, simpler architecture and higher throughput than cache-based architectures. Robust clustering and high availability make it ideal for mission critical use cases that require strict SLAs on latency. The already low TCO benefits Aerospike provides can be compounded by additional savings or extra business features that can be realized due to the ability to get more data, faster.

Strong consistency with high performance

Edge-based systems requiring strong consistency traditionally traded performance for data integrity. Aerospike's HMA allows strong consistency with virtually zero performance penalty, opening up realms of possibilities of new use cases across Finserv, Banking, Payments, Retail and Telco.

Performance is predictable


The cache is removed from the technology stack, so there are no cache misses. There is also no need to prime the cache. All data is accessed in the same way. Latency is nearly constant across all data accesses, from low data volumes to 100's of TB of data. This gives predictable performance across wide volumes of data.

SSD optimization yields massive performance benefits

SSDs are not simply fast hard drives. Their underlying characteristics are fundamentally different, requiring special considerations for optimal performance. Uniquely designed for SSDs, Aerospike's HMA takes full advantage of these characteristics, yielding performance gains far exceeding those resulting from porting existing storage technologies from rotational drives to SSDs.

Massively parallel architecture allows linear scalability

Scaling a database cluster should be as easy as adding nodes or adding hardware within existing nodes. This is realized in Aerospike's HMA architecture: the throughput and data storage capacity of the cluster will scale linearly both with the addition of nodes as well as



with additional hardware. Latency is unaffected as the cluster scales; thanks to Aerospike's Smart Partitions™ algorithm, database requests are sent directly to the node that holds the data with no need for intermediate processing.

Low Total Cost of Ownership (TCO)

Cached architectures rely on large amounts of DRAM for performance, often requiring 50-90% of the data to be stored in volatile memory that happens to be expensive. Aerospike's HMA requires significantly lower quantities of DRAM, always accessing the information from comparatively lower cost SSDs. The scalability and density of SSDs far exceeds that of DRAM and the cost is significantly lower, resulting in significantly lower TCO. Reductions of 4x-10x the number of nodes and similar TCO savings have been realized by Aerospike customers.

Applications become simpler, more maintainable and faster to market

There are inherent complexities in maintaining caches as well as data integrity issues that a high-speed, cache-less architecture like HMA avoids.

Self-healing clusters with identical nodes gives high availability

Aerospike automatically distributes data over all the nodes in the cluster so each node has approximately the same data volume. Losing a node causes the cluster to rebalance evenly preventing costly cascading failures and ensuring multiple copies of the data remain available. Traffic is served normally while rebalancing occurs allowing business to continue.

About Aerospike

Aerospike is the world's leading enterprise-grade, internet scale, key-value database whose patented Hybrid Memory Architecture™ enables digital transformation by powering real-time, mission critical applications and analysis. Only Aerospike delivers strong consistency, predictable high performance and low TCO with linear scalability. Serving the financial services, banking, telecommunications, technology, retail/e-commerce, adtech/martech and gaming industries, Aerospike has proven customer deployments with zero downtime for seven years running. Recognized by industry analysts as a visionary and leader, Aerospike customers include Nielsen, Williams Sonoma, Kayak, Neustar, Bharti Airtel, ThreatMetrix, InMobi, Applovin and AppNexus. Aerospike is based in Mountain View, CA, and is backed by New Enterprise Associates, Alsop Louie Partners, Eastward Capital Partners, CNTP and Silicon Valley Bank.

©2018 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.