Petabyte Scale in AWS with Aerospike® Utilizing Intel® Technology

Petabyte Scale in AWS with Aerospike® Utilizing Intel® Technology

Intel-enabled Aerospike servers on AWS dramatically reduce database TCO while delivering < 1 ms latencies.[1]

Contents

Introduction

The persistent trend of data growth

Business leaders today are looking toward technology-driven capabilities to help accelerate their business growth. Cloud computing, AI, and database solutions are essential ingredients in driving faster go to market for new services and products. The infrastructure to support remote workforces is granting employees more flexibility while helping make businesses more adaptable to global disruptions. And as ever, there is a persistent need to accomplish more with fewer resources amidst growing budgetary constraints while supporting the varied needs of more-specialized business units.

Massive database scalability without massive server footprints

A recent petabyte benchmark test by Amazon Web Services (AWS) and Intel shows how enterprises can use Aerospike®, a distributed NoSQL database management system (DBMS) platform, in combination with key innovations in Intel® Xeon® Scalable processors and Intel® Optane™ technology to save millions of dollars per cloud application. With this solution, enterprises can power their real-time applications with predictable, sub-millisecond performance up to petabyte scale.

Early implementations of the Aerospike and Intel deployment on AWS Cloud instances show remarkable results.

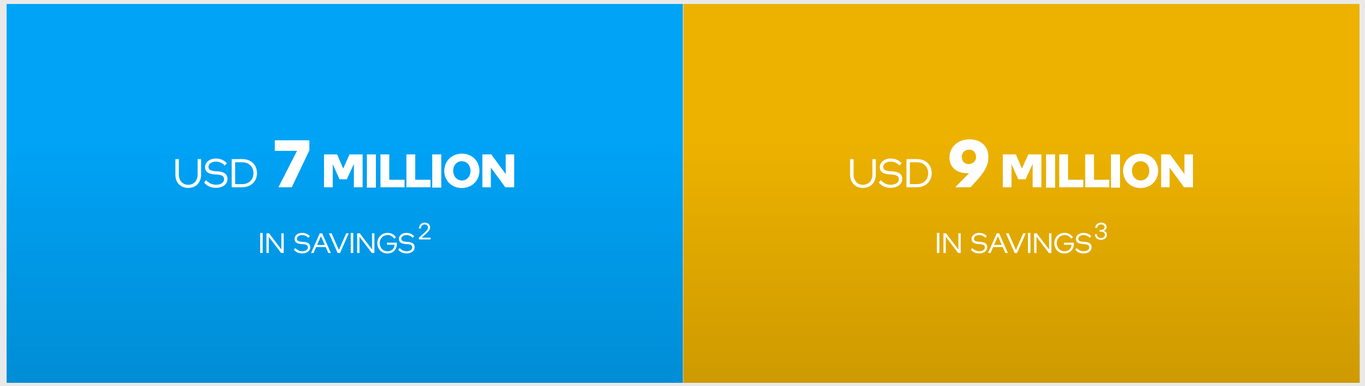

Signal, a TransUnion company, expects to save USD 7 million from a reduced server count from 450 to 60.

PayPal expects to save USD 9 million, with projected reductions in total cost of ownership (TCO) by 80 percent over five years.

> "Prior to Aerospike, we were using another in-memory data store and … running into challenges in terms of the cost of scaling… We moved to Aerospike for its hybrid memory architecture, to leverage next-generation memory and SSDs to their fullest advantage."
>
> — **Sai Devabhaktuni**
> *Sr. Director of Engineering*
> **PayPal ›**

Results based on real-world impact

To help determine a real-world impact, Aerospike® and Intel launched the petabyte benchmark test that simulated AdTech workloads, which offer a unique combination of high volume, high complexity, and high transactional speed challenges. The engineering team also set a high bar for success, establishing several key objectives to help ensure the solution's value to potential enterprise customers.

Key objectives for the Aerospike and Intel petabyte benchmark test:

- Act instantly across billions of transactions
- Reduce server footprint up to 80%
- Achieve sub-millisecond performance on real-time applications
- Maintain five-nines uptime with globally distributed, strongly consistent data
- Have the ability to deploy quickly on-premises, in multicloud, or in hybrid configurations

Which would you rather manage?

Just because businesses are growing their data footprint to petabyte scale doesn't mean they have to expand their server footprint. Aerospike and Intel ran the test against a control group workload featuring a traditional open source NoSQL database management tool. The petabyte benchmark test proved successful across all objectives while showcasing a dramatic USD 11 million in savings over a hypothetical three-year period with a data growth rate of 15 percent per year.[1]

Here are the results up front:[1]

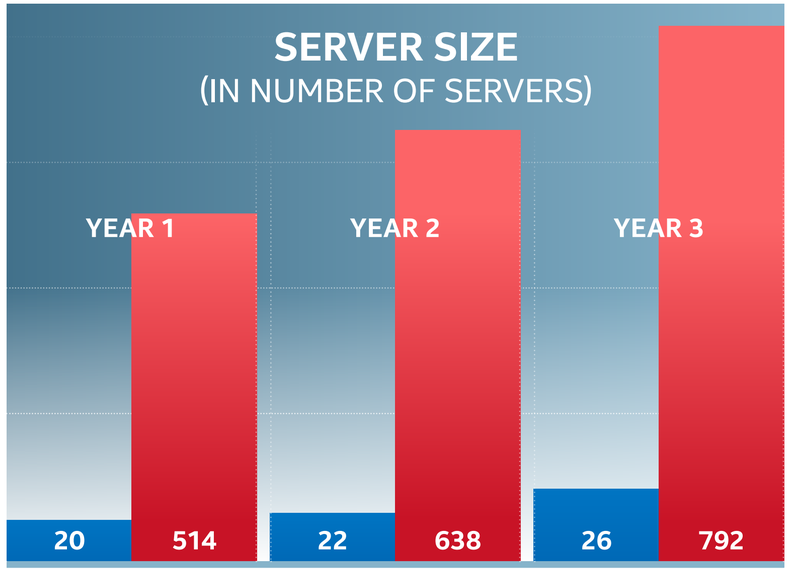

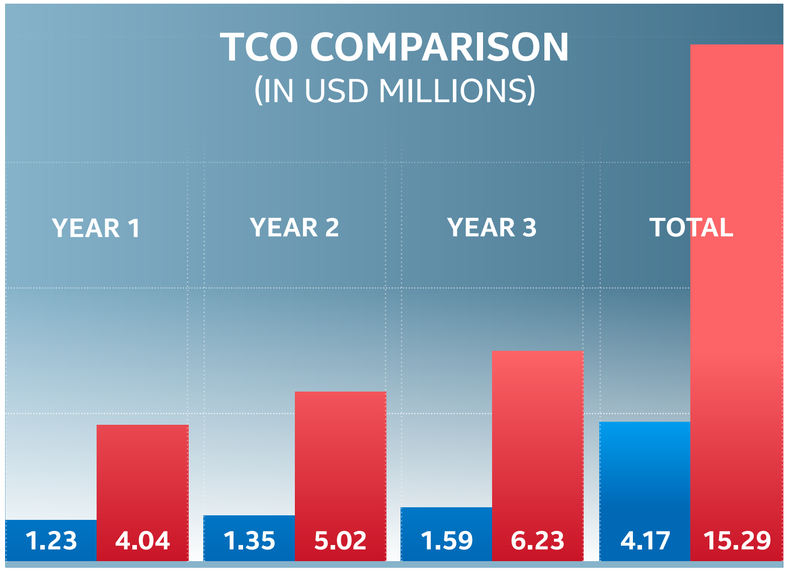

Aerospike vs. **Open source alternative**

Smaller TCO numbers are better.

Transactions per second (TPS): Aerospike processed up to 5 million TPS at < 1 ms latency

TCO: Aerospike costs USD 4.17 million vs. open source TCO of USD 15.29 million over three years

Server footprint: Aerospike at 26 nodes vs. open source at 792 nodes by year three

More data, more challenges

Traditional database solutions continue to be challenged not only by the amount of data being generated, but also by the mixed types of data that need to be ingested and analyzed. Transactional, demographic, and alternative data types, including multimedia and social data, are adding to a pool of information that presents new opportunities for insight. This complexity places new demands on database infrastructures and IT organizations across a spectrum of needs: performance, operational ease, elasticity, availability, data consistency, enterprise integration, and cost efficiency.

Petabyte pressure points to look out for

As database operations grow, these pressure points can help indicate the need for a more comprehensive and strategic solution.

SCALE

Many open source and commercial solutions simply can't scale up to multipetabyte-level processing for mixed workloads, with a crippling drop in reliability or availability that renders the solution unusable.

PERFORMANCE

Relational database management systems (RDBMSs) provide strong data consistency but can't deliver ultrafast performance at scale with low TCO. Traditional caching systems often exhibit erratic latencies at 5+ TB.

COMPLEXITY

Open source and some commercial NoSQL systems suffer operational complexity, unpredictable performance, and sprawling server footprints. Businesses typically end up generating a lot of "weirdware" to compensate for shortcomings in open source software.

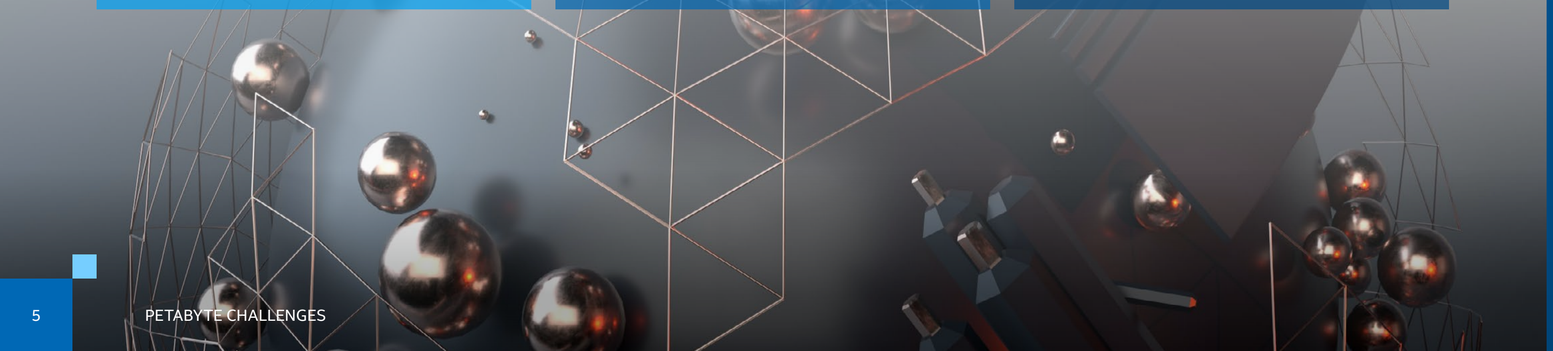

> "In that rare engineering trifecta, we've found that [the Aerospike® and Intel® solution is] faster, cheaper, and far more reliable."
>
> **— Jason Yanowitz**
> *Vice President of Engineering*
> **Signal ›**

What is weirdware?

A colloquial term, *weirdware* refers to deploying workarounds that connect disparate platforms or get existing platforms to perform a specialized function that's not officially supported. Weirdware is problematic because it leads to increased technical debt, which refers to deploying an easy solution now that will require rework later.

The Aerospike® and Intel difference

Aerospike and Intel deliver fast storage I/O at petabyte scale with strong consistency, without hundreds of servers to manage and maintain. Truly the best time for a business to migrate their workloads to Aerospike is before they start encountering challenges with manageability, scalability, and uptime. For businesses that expect their data or transactional workloads to increase in the next few years, now is the time to start planning a move to Aerospike.

Top three challenges/reasons to consider Aerospike

The enterprise server footprint and TCO keep expanding.

The enterprise is finding it harder and harder to meet service-level agreements (SLAs).

The enterprise is struggling to manage peak or seasonal workloads.

Top three ways Aerospike can solve these challenges

Aerospike exploits the latest technologies.

Multicore Intel® Xeon® Scalable processors, non-uniform memory access (NUMA), non-volatile memory express (NVMe) SSDs, Intel® Optane™ persistent memory (PMem), and network application device queues (ADQ) help enable unprecedented density and utilization, resulting in higher impact with fewer nodes.

Aerospike provides self-managing, self-healing features with high availability and data consistency.

Configurations can automatically balance workloads across virtualized resources to ensure uptime, without the need for administrator intervention.

Flexible deployment options allow customers to scale operational workloads easily.

Hybrid and multicloud environments allow Aerospike customers to quickly add or remove compute and storage nodes on demand.

> "Very soon we reached a point that Redis couldn't scale anymore. We moved to Aerospike. Aerospike is developer friendly. The scale problem was the first to be solved."
>
> — **Gil Shoshan**
> *Software Engineer*
> **ironSource ›**

Freeing up the server footprint, technology budget, and time for innovation

The real-world impact of Aerospike® can best be measured by the success of its customers. Not only does the Intel-enabled Aerospike deployment help cut costs and consolidate infrastructure, but it also frees up developer resources and time to focus on big-picture initiatives and new services development.

KEY BENEFIT:
More time to focus on new product development

"Before Aerospike, we were spending more and more of our time on the care and feeding of Cassandra and less and less time on the building of new product offerings. With Aerospike, we've now cleared the road map, and we're just focused on adding new functionality to our platform for our customers."

— **Jason Yanowitz**
Vice President of Engineering
Signal ›

KEY BENEFIT:
Responsive services and decision-making in high-risk environments

"With our previous database, we were operating at a request duration that would hurt our customers because they couldn't make a risk decision in the allocated time window … Aerospike helps us … to return in almost all cases with a very low latency so that all risk decisions can be made."

— **Matthias Baumhof**
VP Worldwide Engineering
ThreatMetrix ›

KEY BENEFIT:
Support millions of customers without a trade-off in user experience

"With our past database [MongoDB], whenever there was a search in concurrent price updates from many services, we saw degradation in the buyer experience. Now with Aerospike, we can push through huge price changes while maintaining the same response time experience on the buyer's side—even with millions of buyers."

— *Vice President of Engineering*
Snapdeal ›

KEY BENEFIT:
Save millions with a consolidated server footprint

"Aerospike has done the unthinkable: they cut our server footprint by a factor of six while boosting our performance 300 percent … saving us over $1 million a year."

— **Guy Almog**
Head of IT Engineering
Playtika ›

The petabyte benchmark test in detail

To help put the Aerospike® configuration through its rigors, engineers chose AdTech workloads because of the characteristics and challenges they present. AdTech is largely about serving up ads to specific end users across various web pages and applications based on users' profile characteristics. This process necessitates high volumes of demographic data that can account for several petabytes or more per profile.

On the back end, businesses bid for each ad placement, an automated process that needs to support 10 million or more bid requests per second, with bid responses completing in under 75 ms and auctions closing in under 100 ms per ad placement. For the test, Aerospike ran both the user profile store and campaign store operations in parallel. This test involved a complex front end, a complex back end, and was the perfect battleground to test Aerospike's capabilities.

AWS instances chosen

Note: The AWS instances used in the benchmark were not equipped with Intel® Optane™ persistent memory modules, which could potentially raise the performance levels even higher by accelerating storage I/O. As of the writing of this paper, Intel and AWS are collaborating on integrating Intel Optane persistent memory into key instance offerings.

Up to 5 million read-only transactions per second at 100% < 1 ms latency[1]

For AdTech and many other use cases, transaction latency and uptime are critical to services delivery. Even though many enterprises can achieve several nines (99 percent or more) of uptime, the tiniest outage or failed transaction can have a huge impact on business reputation.

Aerospike® and Intel ran two tests, each for four hours and each running the user profile and campaign operations in parallel across all keys in the user profile and campaign databases. Test one focused on read-only operations against the user profile database, while test two introduced a 20 percent write pressure against the user profile database. Here are the results:

TEST #	DATA	WORKLOAD	TPS	LATENCY < 1 MS
1	User profile (1 PB unique, uncompressed)	Read only	5,009,980 reads	100%
1	Campaign (1.5 TB unique, uncompressed)	50/50 read/write	95,420 reads 95,420 writes	100% 100%
2	User profile (1 PB unique, uncompressed)	80/20 read/write	3,017,340 reads 754,160 writes	100% 99%
2	Campaign (1.5 TB unique, uncompressed)	50/50 read/write	95,800 reads 95,800 writes	100% 99%

Table 1: Benchmark results. User profile and campaign workloads ran concurrently.

Dollars and sense: USD ~11 million in TCO savings by year three[1]

For many firms, TCO is the biggest inhibitor to petabyte-scale deployments, and the biggest headaches are infrastructure, maintenance, and operational costs. To help address these concerns, Aerospike and Intel extrapolated capital expenditure (CapEx) and operational expenditure (OpEx) figures across three years for two hypothetical deployments—one featuring Aerospike, the other featuring an open source alternative.

For this test, Aerospike and Intel kept the same petabyte parameters in the TPS test but added in an assumed data growth rate of 15 percent per year. As a result, each configuration would need to add more server nodes and more administrative overhead each year. IT staffing costs were estimated at USD 180,000 per person/year per 200 server nodes. Here are the results:

Aerospike

	YEAR 1	YEAR 2	YEAR 3	TOTAL
Cluster size (Total servers)	20	22	26	
Cost per server (USD)	$60,409			
Infrastructure (USD)	$1,208,179	$1,328,997	$1,570,633	**$4,107,809**
Fully burdened maintenance & support (USD)	$18,000	$19,800	$23,400	**$61,200**
TCO (USD in millions)	$1.23	$1.35	$1.59	**$4.17**

Open Source Alternative

	YEAR 1	YEAR 2	YEAR 3	TOTAL
Cluster size (Total servers)	514	638	792	
Cost per server (USD)	$6,964			
Infrastructure (USD)	$3,579,599	$4,443,160	$5,515,646	**$13,538,405**
Fully burdened maintenance & support (USD)	$462,600	$574,200	$712,800	**$1,749,600**
TCO (USD in millions)	$4.04	$5.02	$6.23	**$15.29**

While the open source alternative presents a lower cost per server due to its free licensing, the overall configuration requires hundreds more servers (and ballooning CapEx/OpEx) in order to maintain comparable TPS and latency with the Aerospike configuration.

The breakdown

The results for the first test show that Aerospike running on Intel-enabled AWS nodes was able to achieve 5 million read-only user profile TPS with < 1 ms latency while also processing almost 200,000 campaign transactions, also at < 1 ms latency. The second test shows that even with 20 percent write pressure for user profile transactions, Aerospike achieved a total of 3.7 million TPS for user profiles, nearly all with sub-millisecond latency.

Interested in other workloads?

See more comparative benchmarks ›

Interested in benchmarking your own workloads?

Contact sales@aerospike.com ›

The bottom line

Scaling up to petabyte databases while supporting millions of transactions at sub-millisecond latencies doesn't have to be a huge headache. Aerospike® takes advantage of the latest innovations in Intel® hardware to help maximize density and operational ease while minimizing cost when compared to open source alternatives that can't scale properly.

Results at a real-life AdTech firm

The Trade Desk transitioned its database clusters to Aerospike and now manages real-time bidding with up to 11 million queries per second (QPS) and up to 800 billion queries per day. They manage this level of data processing with just 17 Aerospike clusters across 10 data centers.

Footprint, overhead, and TCO results by third year[1]

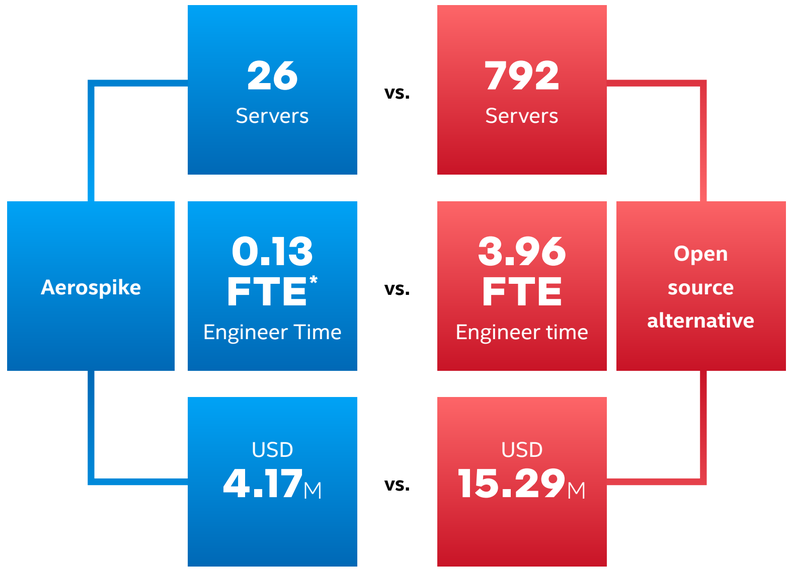

*Full-time equivalent

> "[We] run some of the largest [Aerospike] clusters in the world. Aerospike is the best NoSQL solution to support our millions of QPS workloads. [It's] amazingly simple to operate; dynamic cluster management and scalability are key."
>
> — **Matt Cochran**
> *Director of Engineering*
> **The Trade Desk ›**

Modernize and prepare for petabyte scale

Together, Aerospike® and Intel can be a cornerstone to how enterprises approach petabyte scale, as well as part of a larger modernization initiative. Data, customer bases, and customer expectations will only continue to grow.

Get started today

Dive deeper with the Aerospike white paper

Get started with Aerospike in the cloud today

About 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, Verizon Media and Wayfair 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.

> "Aerospike is the only product to beat our own technology in 12 years, and we've tested everything."
>
> — *Verizon Media at the Aerospike Digital Summit 2021* ›

> "We plan to store 500 TB of … readily accessible data [in Aerospike]. Surprisingly … we are able to come up with a pretty small footprint to store that huge amount of data while we are getting all the benefits of low latency and [are] able to keep up with all the SLAs."
>
> — *Vice President of Technology at a global brokerage firm*

intel + AEROSPIKE

1. “Running Operational Workloads with Aerospike at Petabyte Scale in the Cloud on 20 Nodes,” jointly developed by Aerospike, Intel, and AWS, June 2021. aerospike.com/lp/running-operational-workloads/

   Server configuration details: Aerospike Enterprise Server 5.4.0.2-1 on Red Hat 7; 20 nodes of AWS EC2 i3en.24xlarge instances with up to 3.1 GHz Intel® Xeon® Scalable (Skylake) processors with new Intel® Advanced Vector Extensions 512 (Intel® AVX-512) instruction set; 768 Gb of memory; 8x 7500 NVME SSDs; 96 vCPU; 100 Gbps network; OS: Amazon Linux 2.

   Client configuration details: 40 nodes of AWS EC2 c5n.9xlarge instances with 3.0 GHz Intel® Xeon® Platinum processors with Intel Advanced Vector Extensions 512 (Intel AVX-512) instruction set; sustained all-core Turbo frequency of up to 3.4 GHz, and single core turbo frequency of up to 3.5 GHz; 96 Gb of memory; 36 vCPU; 50 Gbps network.

2. “Debunking the Free Open Source Myth,” Aerospike website, date of access: October 2021. aerospike.com/debunking-the-free-open-source-myth/

3. “PayPal puts data at the heart of its fraud strategy with Aerospike,” Aerospike website, date of access: October 2021. aerospike.com/customers/paypal/

Notices and disclaimers

0122/SMR/CMD/PDF