

White Paper

AEROSPIKE

Aerospike Database 5 on AWS

Build faster, scale-up, and save more in the Cloud

Contents

Overview	3
Aerospike Cloud Managed Service	4
Aerospike Quick Start for AWS	4
Aerospike AMI on the AWS Marketplace	5
Aerospike Kubernetes Operator	5
Key features of the Aerospike Data Platform	6
Why Aerospike for AWS	8
Business benefits	9
Top use cases	10
Summary	11

Overview

According to [research by Gartner](#), by 2022, 75 percent of all databases will be deployed or migrated to a cloud platform. Aerospike's AWS strategy is predicated on helping customers avoid the complexity of cloud deployments and minimizing its cost while providing a platform that can accelerate application development.

The Aerospike Real-time Data Platform enables organizations to act instantly across billions of transactions while reducing cloud instances by up to 80 percent. The Aerospike data 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 users.

Aerospike offers multiple deployment options to fit your organization's needs, whether looking into digital transformation, database modernization, agile development, or re-architecting applications. Aerospike can help you accelerate time-to-value, scale with predictable performance, and become more productive, reducing your cloud expense.

Customers such as AppsFlyer, Experian, Dream11, DBS, Beeswax, and Factor Eleven have deployed Aerospike on AWS as their foundation for the future, integrating into their DevOps cycles, connecting with AWS data pipelines, or providing observability for data governance through one of the following methods:

Highlights

- **Performance**
Evaluated and recommended Amazon EC2 instances for each particular workload growing from terabytes to petabytes with linear performance.
- **Connectivity**
Aerospike Connect for Amazon MSK, KMS, AWS Lambda, Amazon Kinesis Data Analytics, Amazon EMR, and more.
- **Cost optimization**
Save up to 80% with fewer instances using the hardware efficient Aerospike Hybrid Memory Architecture on AWS.
- **Flexible Deployment**
Select from Aerospike Cloud Managed Service, the simplified Aerospike Quick Start, automated Aerospike Kubernetes Operator, or self-managed cluster through the AWS Marketplace.

Aerospike Cloud Managed Service

The Aerospike Cloud Managed Service offers comprehensive administration and management of the real-time data platform by a team of experts for consistent deployment of highly-available Aerospike clusters. It provides multiple levels of SLAs that fit any particular need of the application lifecycle. Development teams can begin with a small Aerospike cluster to prototype, test functionality, or develop the application, then upgrade and promote those clusters into staging, tuning, and integration testing. Once those phases are complete, customers can promote the clusters into production, updating the SLAs at each stage for the appropriate level of availability, monitoring, and alerting for a 24x7x365 operation.

The Aerospike Cloud Managed Service team helps you define the right starting deployment size and characteristics for your cluster; determine the need for HA (High Availability) or DR (Disaster Recovery) options should a catastrophic infrastructure failure occur on the primary cluster. The Aerospike team also grants access to the Aerospike Cloud Console, the central location for monitoring, change requests, and ticket management for the entire Cloud Managed Service.

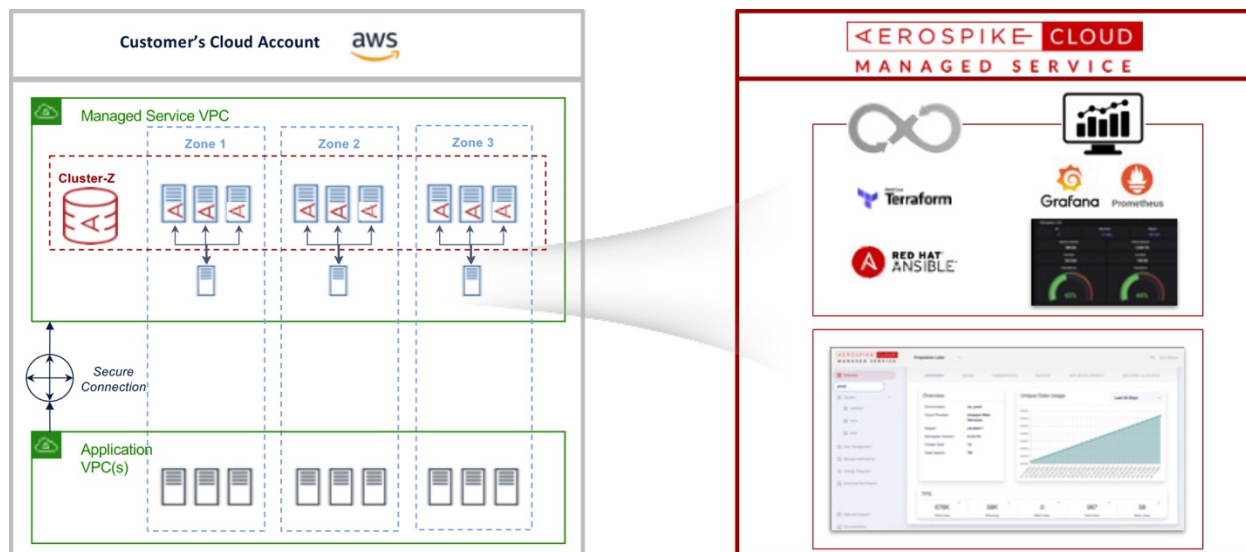


Fig. 1: Aerospike Cloud Managed Service architecture

Aerospike Quick Start for AWS

The Aerospike Quick Start is an automated deployment guide for rapid deployment of Aerospike clusters on AWS built as a collaboration between Aerospike and AWS.

The deployment guide uses an AWS CloudFormation template built following best security and resource utilization practices to help you deploy a working cluster on a new or existing VPC (Virtual Private Cloud) on your preferred Amazon EC2 (Elastic Cloud Compute) instance.

The automation of the template guides you through the few necessary steps to successfully deploy the cluster in under 15 minutes. An evaluation license for a single-node cluster¹ is provided by default to start your development right away, with the option to use a license supplied by Aerospike for multiple nodes or multi-cluster deployments (BYOL).²

Aerospike AMI on the AWS Marketplace

The Aerospike Real-time Data Platform is available as an Amazon Machine Image (AMI) in the AWS Marketplace, where customers can subscribe and deploy it as a self-managed system into their cloud infrastructure.

Aerospike provides a CloudFormation template that customers can use to deploy into a new or existing VPC; the template guides you through the process for configuring Availability Zones, security policies, cloud instance types, and memory and storage needs. The Aerospike team can provide sizing and configuration metrics support for a custom fit to the specific workload.

A single-node cluster license¹ for development and testing of the Aerospike Database 5 Enterprise Edition is available by default. To deploy a multi-node, multi-cluster, or use the XDR (Cross Datacenter Replication) capabilities, customers require additional licensing. Customers can either use an available license or obtain a new license by contacting Aerospike² using the following link: <https://aerospike.com/forms/contact-us/>

Aerospike Kubernetes Operator

The Aerospike Kubernetes Operator lets customers efficiently automate the management of everyday tasks such as the configuration, provisioning, scaling, and recovery of Aerospike clusters, thereby reducing the complexity of manual deployment and lifecycle management.

The Aerospike Kubernetes Operator enables organizations to deploy using best practices and easily coordinate a large number of clusters; it provides granular control of security and encryption to meet regulatory compliance, deploying multi-level enterprise security.

With features like rack awareness, persistent storage volumes, automatic scalability, and support for Amazon Elastic Kubernetes Service (EKS), the Aerospike Kubernetes Operator lets users automate operations and upgrades directly into AWS.

¹The evaluation license is a single-node Enterprise Edition license that includes the use of our Aerospike Connect collection for SPARK, Kafka, Presto, and more.

²Contact Aerospike for a multi-node license. <https://aerospike.com/forms/contact-us/>

Key features of the Aerospike Data Platform

- Hybrid-Memory Architecture**

Aerospike enables the use of flash storage (SSD, PCIe, NVMe) in parallel on one machine to perform reads at sub-millisecond latencies at very high throughput (100K to 1M) in the presence of a heavy write load. This use of SSD enables enormous vertical scaleup at a 5x lower total cost of ownership (TCO) than pure RAM.

Aerospike implements a hybrid memory architecture wherein the index is purely in-memory (not persisted), and data is stored only on persistent storage (SSD) and read directly from the disk. Disk I/O is not required to access the index, which enables predictable performance.

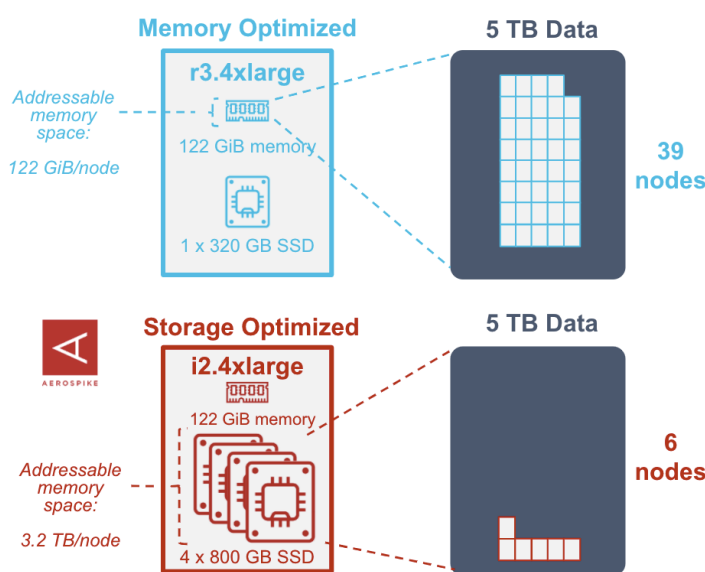


Fig. 2: Legacy architecture compared to Aerospike’s HMA.

- Cross Datacenter Replication - XDR**

Aerospike Database Cross Datacenter Replication (XDR) provides dynamic, fine-grained control for the replication of data across geographically separate clusters. It can be used to build globally distributed applications, with low latency reads and writes at each cloud region, accessing the data that is meaningful to each region while complying with data locality regulations.

The Aerospike XDR feature replicates data asynchronously between two or more clusters. A site can be a physical rack in a datacenter, an entire datacenter, an availability zone in a cloud region, or a cloud region. This flexible configuration enables different deployment models based on data availability or governance regulations.




Fig.3: Overview of Aerospike XDR

- **Query Optimization**

Aerospike utilizes multiple technologies such as set indexes, secondary indexes, and Aerospike Expressions to power query performance on data at the Edge, in a System of Record, and across data centers and clouds. These fast yet powerful queries can be utilized to analyze data to support real-time decisions, multi-tenancy, shared services, and compliance needs with great control and precision. Additionally, given the Aerospike real-time data platform’s ability to scale from terabytes to petabytes of data with linear performance, these queries can be run efficiently on large data sets.

- **Aerospike Connect**

The Aerospike Connect product line integrates Aerospike with popular frameworks, including Spark, Kafka, Pulsar, JMS, and Presto. It enables the design and development of new generations of data-intensive applications utilizing the speed and scale of the Aerospike real-time data platform massively parallel engine.

The Aerospike connectors allow seamless integration with AWS data pipelines and services like Amazon MSK (Managed Streaming for Apache Kafka), AWS KMS (Key Management Service), AWS Lambda, Amazon Kinesis Data Analytics, and Amazon EMR.

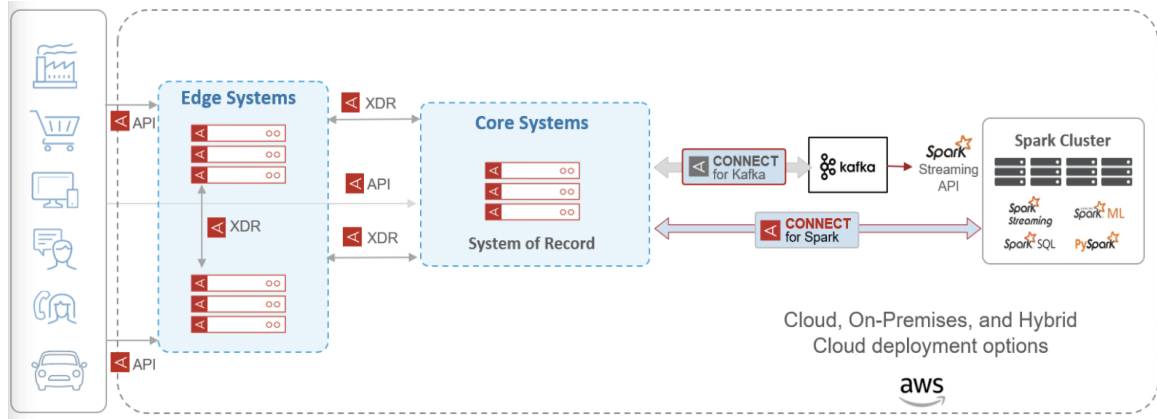


Fig. 4: Sample deployment with Aerospike Connect

- **Smart Client**

Aerospike’s Smart Client automatically distributes both data and traffic, load balancing to all the nodes in a cluster to improve the performance of data-oriented applications.

The Aerospike Smart Client is a first-class observer of the cluster and its nodes. Each node maintains a list of all the neighboring nodes, and the client uses this list to discover the entire cluster set. Once the client knows all the nodes and their roles, it caches a map of nodes, partitions, masters, and replicas for a highly efficient client-server interaction, then the administrator can define policies to perform reads from a local zone for a more efficient operation, then the administrator can define policies to perform reads from a local zone or specific replicas for a more efficient operation.

This feature is why, in Aerospike, there is single-hop access to data from the client; therefore, the linear scale-out ability of the Aerospike cluster is purely a function of the number of clients or server instances.

Why Aerospike for AWS

- **Optimized for AWS**

Aerospike in close collaboration with AWS built CloudFormation templates for rapid deployment of Aerospike clusters following cloud deployment best practices for security and performance.

- **Real-time database**

Aerospike Database 5 supports real-time data operations for latency-sensitive applications, even for large data sets up to petabyte-scale with the highest throughput and microsecond latency as demonstrated in the collaborative benchmark and white paper “Running operational workloads with Aerospike at petabyte scale in the cloud on 20 nodes.”

- Predictable performance**

Aerospike is optimized carefully defining a list of recommended Amazon EC2 instances for different workload profiles to use our Hybrid-Memory Architecture (HMA) and CPU parallel processing capabilities to provide consistent and predictable performance.

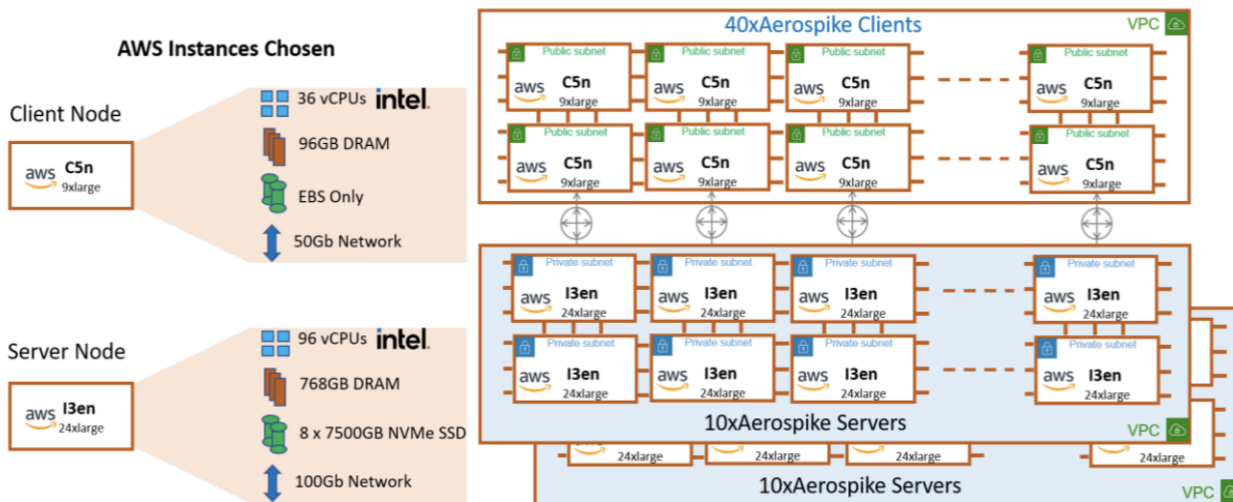


Fig. 5: Overview of Aerospike’s petabyte benchmark architecture

- Ease of deployment**

Multiple deployment options to choose from and combine to fit your needs, from a fully cloud-managed service to self-managed AWS Marketplace to automated Kubernetes operation, or rapid deployment through a Quick Start guide, Aerospike has a frictionless integration model for any IT operations model.

Business benefits

- Accelerate time-to-value**

Built on years of experience architecting and managing the Aerospike real-time data platform both on-premises and in the cloud, Aerospike's best practices and pre-defined cluster templates for infrastructure and database configurations accelerate the deployment cycle.

- Increase Productivity**

Focus on aligning your organizational resources with your goals instead of using valuable resources and time to design and maintain your Aerospike deployment. Rapidly deploy new clusters using the CloudFormation templates for self-manage service or use the Aerospike Cloud Console to request changes in the managed service.

- **Global data availability**

Aerospike provides options to effectively and efficiently manage geographically distributed data across multiple cloud regions and availability zones. Aerospike supports two distinct active-active deployment models, multi-site clustering (MSC) and Cross Datacenter Replication (XDR) for building a global data hub.

- **Cost reduction**

Stop over-deploying in the cloud and instead use the hardware efficient Aerospike Hybrid Memory Architecture on AWS. Deploy the resources that you need when you need them without massive upfront commitments. Generate an environment with predictable performance from a few terabytes up to petabytes of data with five-nines of uptime while reducing the number of cloud instances.

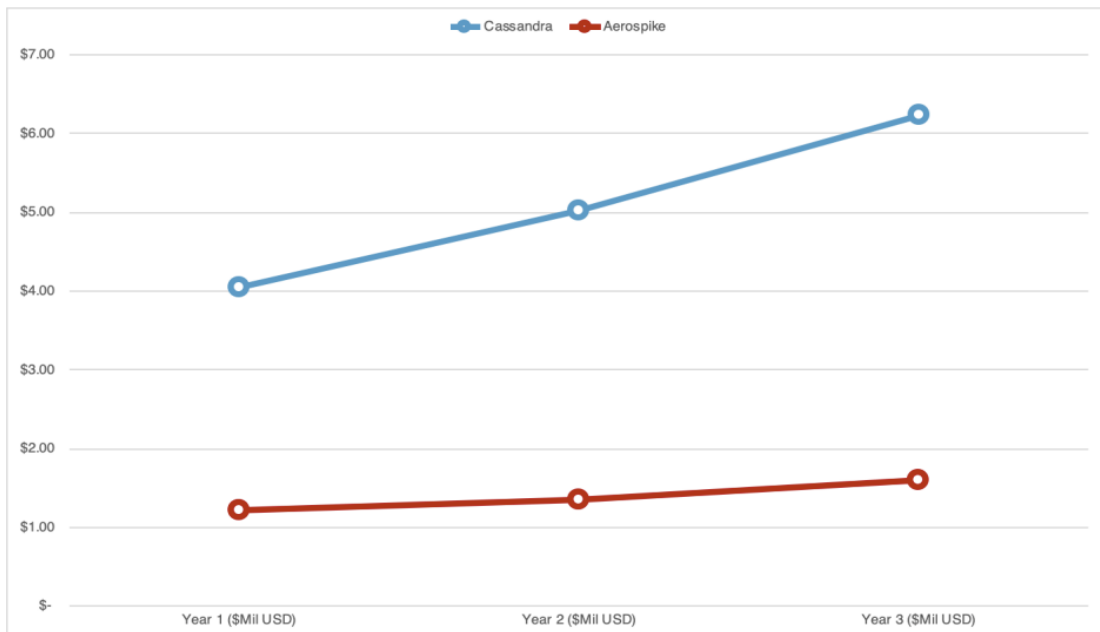


Fig. 6: Total operational expenses by year for Cassandra and Aerospike

Top use cases

- **Fraud Prevention**

Enable passive identity validation and detect fraud by analyzing millions of transactions, billions of data points, and petabytes of historical information in milliseconds. Reduce fraud exposure by 30x and de-risk ML model testing more frequently with more data and reduce false positives employing ten times the number of attributes in your same calculation window. Visit the *Aerospike Financial Services solutions page*.

<https://aerospike.com/solutions/industry/financial-services-payments/>

- **Recommendation Engine**

To create a truly friction-free digital, omnichannel commerce shopping experience a recommendation engine combines innovative math with domain-specific knowledge to increase online engagement. Given low online switching costs and fickle online behavior, providers must offer a great experience to attract and keep customers. These decisions need to happen in milliseconds, supporting diverse applications while processing all the available data from customers, warehouses, in-store, and supply chains. *Read how snapdeal provides dynamic pricing, high availability, and responsiveness for 20,000 sellers.*

[http://pages.aerospike.com/rs/229-XUE-318/images/Aerospike Case Study_Snapdeal.pdf](http://pages.aerospike.com/rs/229-XUE-318/images/Aerospike_Case_Study_Snapdeal.pdf)

- **Customer360**

Increase revenue, provide better customer experiences and improve operations with a real-time customer engagement platform. There is a massive amount of streaming data coming in from mobile, 5G, and IoT sensor applications. The challenge is to manage, leverage, and analyze this data across silos, optimizing new investments in modern infrastructure and applications, and provide the best customer experience – all at real-time speeds.

Learn more about Aerospike in the Telecommunications industry on our solutions page. <https://aerospike.com/solutions/industry/telco/>

- **Ad Tech**

With millions of requests per second, real-time bidding execution requires vast volumes of data processing at unprecedented rates, and firms face tremendous challenges to deliver new faster applications and real-time analytics from terabytes to petabytes of data.

Aerospike's Hybrid memory architecture and built-in parallelism enable sub-millisecond latencies and millions of operations per second throughput on large data sets. *Learn more on our AdTech web page.* <https://aerospike.com/solutions/industry/adtech/>

Summary

Aerospike offers multiple deployment options to fit your organization's needs, whether looking into digital transformation, database modernization, agile development, or re-architecting applications. Aerospike can help you accelerate time-to-value, scale with predictable performance, and become more productive, reducing your cloud expense.

- The Aerospike Cloud Managed Service reduces the friction and complexity of running a highly sophisticated data platform.
- The Aerospike Quick Start is an automated reference deployment guide to help you quickly build a new AWS environment.

- Deploying Aerospike directly from the AWS Marketplace gives you the freedom of a self-managed service while reducing your cloud footprint.
- The Aerospike Kubernetes Operator allows you to automate the deployment and management of the Aerospike Cluster on Amazon EKS reducing the complexity of a manual deployment.

For more information on the different deployment options visit the [Aerospike on AWS](#) page, and for a list of all our functionality is available in our [product matrix](#).

About Aerospike

The Aerospike Real-time Data Platform enables organizations to act instantly across billions of transactions while reducing cloud infrastructure up to 80%. The Aerospike data platform powers real-time applications with predictable sub-millisecond performance up to petabyte scale 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.

For more information, please visit <https://www.aerospike.com>.