Aerospike

Infinitely possible.™

Aerospike in action

Five case studies across key industries

Introduction

Harnessing the full potential of data is essential for any organization seeking to innovate, enhance customer experiences, and achieve operational excellence. The right technology partner can make all the difference, providing the tools to manage real-time data, scale effectively, and optimize costs. In this e-book, we explore how five leading organizations have transformed their businesses using Aerospike's powerful database capabilities. Each story is not just a testament to the versatility of Aerospike's technology, but also a roadmap for how businesses in any sector can achieve similar results.

Why these stories?

We've carefully selected these five case studies to represent a wide spectrum of use cases and industry verticals. Whether it's Adobe revolutionizing ad targeting in AdTech, PayPal preventing fraud in real time, Sony Interactive Entertainment seamlessly managing millions of user profiles, Barclays enhancing financial services, or Myntra optimizing e-commerce operations, each example provides actionable insights into how Aerospike empowers companies to innovate and scale.

What you'll discover

Across these pages, you'll learn how:

- **Real-time data access:** Enables instant decision-making and enhances customer experiences.
- **Scalability:** Supports growth, from millions of daily transactions to managing petabytes of data, without compromising performance.
- **Cost efficiency:** Reduces operational costs while maintaining or enhancing service quality.

Why it matters to you

No matter what industry your organization operates in, the challenges and opportunities presented in these case studies will likely resonate with your own business objectives. If your goal is to innovate like Adobe, secure transactions like PayPal, engage users like Sony, build trust like Barclays, or streamline operations like Myntra, then this e-book will provide the inspiration and insights you need to move forward.

Dive into these real-world stories to see how Aerospike is enabling businesses across the globe to redefine what's possible with their data. Discover how you can leverage the same technology to achieve your own strategic goals — because the next great success story could be yours.

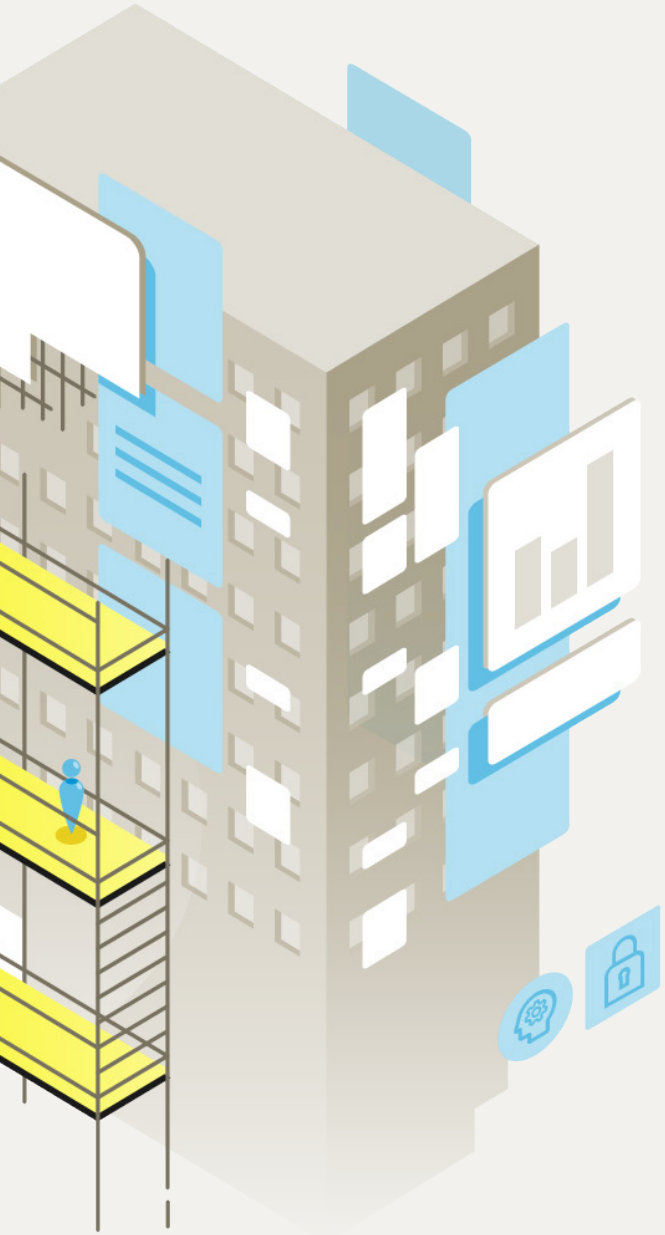

Distributed, real-time personalization on the edge

Adobe elevates digital experiences

Adobe is a multinational computer software company known for their innovation and creative products, and offers digital marketing management software via Adobe Experience Cloud. Adobe Experience Cloud supercharges digital experiences with generative AI, real-time insights, and end-to-end digital marketing tools built on the only platform designed for personalization at scale. The company has experienced an annual growth rate of 32% over the past decade and is valued above $200 billion.

Challenge

A common, unified platform for intelligent personalization

Adobe needed to unify technology from well over 10 acquisitions, all of whom brought their own technology, infrastructure, and requirements. Their original solution with Cassandra, HBase, and an in-memory PostgreSQL database was costly to maintain, siloed in functionality, and impossible to unify into a common platform.

For Adobe Experience Cloud to function as an intelligent digital marketing platform, Adobe needed a solution that could quickly combine vast amounts of profile data, run intelligent services, and deliver real-time targeting, segmentation, personalization, and more. It also had to have low latency and high throughput to ingest, process, and move vast amounts of user data. Providing personalized digital experiences where customers were located — the edge — was a must-have.

Common edge compute system

Reduce operational load while achieving high performance and versatility in multiple use cases.

Enrich experiences with AI/ML services

Low latency response times when combining profile and interface data with intelligence capabilities to build personalized experiences.

Centralize disconnected information

Ingest, process, and combine vast amounts of fragmented data from many edge sources to deliver a 360 view of an enterprise customer.

Solution

Aerospike as the backbone for edge computation

With Aerospike, Adobe has a unified edge computing data platform that intelligently develops in-depth customer profiles and delivers personalized experiences with sub-millisecond response times. Aerospike handles several million requests every second, returns lookups within milliseconds, and is stable even in 95th percentile situations. As a key component of Adobe's edge computations story, Aerospike is becoming the data storage replication and distribution backbone for all edge computation across Adobe's infrastructure.

Single system

Real-time system of record and edge data store to process vast amounts of data, delivering on personalization and segmentation.

High performance

2ms latency at the 95th percentile and 5M transactions / sec high throughput while providing highly stable, predictable characteristics.

Zero downtime migration

Move off of original Cassandra, Hbase, and in-memory PostgreSQL with no downtime or SLA impact.

Lower costs by 3x

Replaces multiple solutions across different use cases and eliminates the need to manage different, complex and costly infrastructures.

Results

3x cost reduction, 50% smaller server footprint

Adobe was able to scale, increase customer satisfaction, and improve the developer experience, all while minimizing their storage footprint and costs.

> We started doing an evaluation agaisnt Cassandra, HBase. Those were the main two ones that we looked at. Postgres pretty much didn't get consideration because it was getting too expensive. And Aerospike was definitely coming out head-and-shoulders above and with respect to two important parameters ... the latency at the 95th percentile ... and the throughput that it can sustain.
>
> —Sandeep Nawathe | Sr. Director of Engineering | Adobe Experience Platform, Adobe

Real-time insights

Increased customer satisfaction and overall system throughput by 7x with 2ms latency @95th percentile and 4-5million TPS at the edge

Data at petabyte scale

20 terabytes of data on the edges with 10 petabytes of data in the Hub across data centers.

3x cost reduction

Reduced costs by 3x and saved up to 50% on storage footprint

1 system to manage, provision, deploy

Unified infrastructure provides a common programmatic API for developers, a single system to manage, provision, and deploy.

PayPal

PayPal puts data at the heart of its fraud strategy with Aerospike

The world's largest online payment system

PayPal is the world's largest online money transfer, billing and payments system. It owns the PayPal, Venmo, iZettle, Xoom, Braintree, and Paydiant brands. By leveraging technology to make financial services and commerce more convenient, affordable, and secure, the PayPal platform empowers over 325 million consumers and merchants in over 200 markets to join and thrive in the global economy.

Challenge

Scaling a fraud prevention solution with 32% annual data growth

While PayPal's fraud rate is impressively lower than the industry average – 0.17% of revenue compared to 1.86% – that small fraction represents over $1B in annual losses. In response, PayPal built a decision platform to identify emerging fraud patterns in real time while minimizing end-user friction. The platform processed, analyzed, and merged big data with fast data, such as customer enrollment, payment, invoicing preferences, and profile data. In 2015, PayPal realized their legacy NoSQL solution was unable to keep up performance while handling massive amounts of data it was collecting and processing. PayPal began a search for a cost-effective database that could both meet their short term and long term needs to scale horizontally, ensure high performance, and uptime.

Real-time decisioning
Needed to quickly process and analyze data to identify emerging fraud patterns in under 200 milliseconds

Faster data processing
Requirement to grow from 3.5 transactions per second to 7 million transactions per second

Scaling for huge amounts of data
Data scale challenges included hundreds of petabytes

Real-time decisioning

Needed to quickly process and analyze data to identify emerging fraud patterns in under 200 milliseconds

Faster data processing

Requirement to grow from 3.5 transactions per second to 7 million transactions per second

Scaling for huge amounts of data

Data scale challenges included hundreds of petabytes

Solution

Requirements met with Aerospike

When their legacy NoSQL database couldn't deliver results without a deterioration in performance, PayPal turned to Aerospike. With its high-performance NoSQL database, Aerospike is built explicitly to run on flash and Persistent Memory (PMem), not DRAM, while consistently delivering speed at scale.

> Prior to Aerospike, we were using another in-memory data store and we were running into challenges in terms of the cost of scaling...
>
> —Sai Devabhaktuni | Sr. Director of Engineering, PayPal

Meet SLAs 99.95%, reducing fraud exposure by 30x

Complete machine learning based fraud detection calculations in the prescribed time window to avoid allowing a potential fraudulent transaction

Improve query performance and data consistency

Strongly consistent data simplifies application logic and thus query performance simultaneously

Slash data load times

While providing reliable access to fresh data — resulting in 12x reduction in re-indexing time.

Cost effectively scale up and out

Handle extreme 32% annual growth in data and objects while keeping total cost of operations low

99.99% uptime

Maintain consistent high availability in a 24/7 environment necessary for digital payments

8 million TPS

Enable more than eight million executions per second across the database environments, which includes RDBMS and NoSQL systems

Results

Scalability, growth, and cost savings — right away and over time

PayPal's partnership with Aerospike delivered huge initial successes while providing the company a growth path to take advantage of the latest hardware and software advances. PayPal's initial deployment with Aerospike's patented Hybrid Memory Architecture™ (HMA) solution provided massive cost savings and huge improvements in fraud calculation, all while maintaining high performance and frictionless customer experience.

After 2015, PayPal's fraud decision platform ballooned to over 4,000 database instances, with 100 petabytes of data, and added 32% data storage each year. In the meantime, a multiyear collaboration between Intel and Aerospike enabled optimizations for higher performance than SSDs, lower cost per GB than DRAM, higher node densities, and data indexes which persist over system restarts. PayPal took advantage of these advances and evaluated 2nd Gen Intel® Xeon® Scalable processors and Intel® Optane™ PMem to increase overall storage density and reduce DRAM costs with Aerospike. With these improvements, PayPal was able to further lower TCO and gain additional technical benefits over the initial solution.

30x reduction fraud exposure

Improved SLA adherence for completing fraud prevention calculations to 99.95% up from 98.5%

8x server footprint reduction

From 1024 servers down to 120.

3x cost reduction

$9m hardware cost savings projected from $12.5m down to $3.5m.

5x throughput improvement

From 200K to 1M transactions per second

Sony Interactive Entertainment

Enabling real-time personalization via machine learning

Sony Interactive Entertainment

Sony Interactive Entertainment is a multinational video game and digital entertainment company owned by Sony Group Corporation with headquarters in the U.S. and Japan. PlayStation is Sony's popular video gaming brand consisting of game consoles, handhelds, a media center, smartphone, online services, and magazines. The company has sold over 545 million Playstation consoles globally. Its online service, the PlayStation Network, includes a virtual market to purchase and download games and multimedia, a subscription-based online service called PlayStation Plus, and a social gaming networking service called PlayStation Home.

Challenge

Personalized decisions for 100s of millions of users

In 2016, after massive success with PlayStation 4, Sony decided to become a data-driven company. With huge amounts of data gathered from 103 million active users, 38.8 million Playstation Plus subscribers, and 5 million virtual headsets, Sony was well positioned to create a machine learning platform that their development teams could use to make models for personalization, enterprise reporting for business decisions, fraud detection, and more.

The challenge, however, was making the data accessible to the teams and data scientists. Sony needed a solution to handle hundreds of millions of users and several terabytes of data in a useful location, determine how to use the data to better understand their users, and then make decisions to personalize the customer experience.

High availability, low latencies at scale

Needed to reliably handle hundreds of microservices, delivering millions of requests, more than 100 billion data events per day, throughout many database clusters and multiple regions.

Data integration and accessibility

Needed to bring together multiple data islands, formats, and create a common data dictionary that all teams could use to implement sophisticated use cases, like machine learning models.

Reasonable total cost of ownership

Wanted to avoid the expenses associated with vertical scale and be able to plan grow while managing costs.

Solution

A backend database for runtime decisions around 100M+ active users

Sony created a data ocean with federated data ownership for their internal teams, allowing each team that created or sourced the data to own and store that data as a data lake within their own cloud account. A centralized catalog allowed any team or data scientist to access the data to create machine learning models or reports to drive business decisions.

Personalization

Rapid customer identification and authentication with their behaviors and preferences, then customize the user experience in a high-performance environment.

Engaging social feeds

Communicating in-game via chat, messaging and voice. Find friends in-game or via connected social applications. Follows, comments, ratings.

Avoid fraudulent transactions across platforms and during surges in real time

Subsidize play-for-free games with in-app advertisements before transforming to pay-to-play.

Results

Achieving real-time personalization with a machine learning feature store using Aerospike

With Aerospike, Sony built a lightweight machine learning platform that allowed machine learning engineers to create, deploy and run their models, as well as manage workflows.

> Given the personalization scale and the low latenc requirements we had, we felt the Aerospike design — based on flash optimized in memory database technology — is really a good thing for us.
>
> —Suresh Bathini | VP of Software Engineering for PlayStation | Sony Interactive Entertainment

Slashed data load times

12x reduction in re-indexing time while providing reliable access to fresh data

Large scale

Handled 100B+ data events per day and 5TB+ data storage

8 million TPS

Enabled more than eight million executions per second across database environments, which includes RDBMS and NoSQL systems.

Automatic sharding

Made sharding operationally less painful

Low total cost of ownership

Small cluster was able to handle several terabytes of data

Low latency

Under 10 milliseconds

Machine learning at scale to fight credit card fraud

Barclays

Barclays, a British universal bank based in London, moves, lends, invests, and protects money for businesses and clients worldwide. Older than the UK itself, Barclays has a long history of offering innovative financial products, including personal and corporate banking, cards, insurance, mortgages, payments, ecommerce, wealth management, and more.

Challenge

Fragmented platforms caused inaccurate decisions and missed fraud activity

With the ever-growing number of devices, channels, and amount of available data online, opportunities for fraud are also on the rise. Fraud techniques and tools are becoming more sophisticated and accessible with easy access to low-cost computing. Customer needs and behaviors have become more complex, with demands for real-time seamless transactions, declines in brand loyalty, and increased comfort in sharing personal information over social media.

Barclays experienced an 83% increase in data from 2015 to 2020 and was projected to generate 463 exabytes daily by 2023. With over 20 million customers making over 30 million payment transactions a day, Barclays business units found it difficult to accurately analyze user profile information to detect and prevent credit card fraud. Problems included:

Inefficiencies and inconsistencies

With multiples of one product, business units couldn't share user profile data or fraud rules, resulting in false positives or false negatives.

Ballooning costs

Costs to manage and run multiple custom platforms were uncontrollable — hardware resourcing, technical staffing, operations.

Missed fraudulent activity

Vertical and horizontal scale-out problems led to an inability to store or manage the user profile data required to capture fraudulent activity.

Inability to meet SLAs

Performance at peak loads and unpredictable latencies caused delays in detecting fraud, resulting in failures to meet payment cycle SLAs.

Limited opportunities to grow and scale

Barclays found it increasingly difficult to evolve complex, bespoke engineering solutions to achieve goals and meet requirements.

Solution

A single, shared machine learning platform

The payment fraud team at Barclays implemented a machine learning fraud detection platform, powered by Aerospike, as a single platform across all business units. The platform accesses and analyzes user profile data for credit card fraud at sub-millisecond speeds. Characteristics of the solution include:

High throughput with low latencies

Scalability to process millions of transactions daily and handle 6x data growth

Strong consistency and security

Barclays has eliminated false positives and negatives, loss of user profile data, and is able to keep up with security requirements as they evolve.

Real-time data handling

Aerospike's Hybrid Memory Architecture allows Barclays to stores indexes in RAM and data on disk

Single platform

Simplified architecture and reduced number of platforms compared to the original structure

Results

Aerospike delivers performance and reliability while handling 6x data growth

With Aerospike, Barclays was able to fight financial fraud in real time with shared rules, user profile data, and information across all its business units. The machine learning platform is able to meet current and future projected performance requirements with a simplified architecture and unmatched reliability.

> With Aerospike, we were able to dramatically reduce stand-in processing (STIP), data consistency issues, as well as false positives and false negatives for future transactions.
>
> —Dheeraj Mudgil | VP, Enterprise Fraud Architect | Barclays

80% reduction in latency

Vastly improved latencies compared to the original system

4x greater throughput

Handles over 10M transactions daily

<100 ms response time

Less than 100ms response time for 99.99 percentile of transactions

6x data growth

Scale to handle growth from 3TB to 30TB+ over 3 years

Real-time personalized homepages for improved customer experiences

Myntra

Myntra began in 2007 as a shop that personalized gifts and grew into a major Indian fashion e-commerce company focused on making the world a more stylish, colorful, and happier place. With $425M in revenue and 11,000 employees, they use technology to help people look great as well as democratize fashion and lifestyle. Myntra provides best in class customer experiences with one of the richest fashion product lineups in the country, including in-house and exclusive brands such as Roadster, HRX, and House of Pataudi.

Challenge

Bottlenecked feature store lookup process was hurting the customer experience

Myntra develops a deep understanding of their customer to provide a real time, personalized shopping experience that recommends sizing, styles, brands, and products specifically tailored to individual preferences. When a customer opens the Myntra homepage, Myntra uses a machine learning (ML) ranking model created from its feature store to populate the screen with widgets aimed at improving the customer experience, resulting in high click-through rates and revenue.

Myntra's model relied on feature lookup processes built on Redis to accurately build the experience. With hundreds of millions of users, 20K orders per minute at peak, a million active users per day, and over 500K concurrent users during peak events like sales, the feature lookup process was taking too long to execute, resulting in poor user experience. Myntra needed a different approach to accommodate the scale and performance they required to serve their growing customer base. Their needs included:

Speed and scale

Deliver homepage personalizations at throughputs of 100K-400K requests per minute and latencies under 40ms at the 99th percentile.

Eliminate bottlenecks in the feature store

Fetch features and embeddings in real time with low latency.

Reduce costs

Needed to rein in spending, reduce cloud resource footprint, and lower TCO.

Ecosystem integrations

3rd party integration with the machine learning data platform, including Spark, Kafka, Trino, and other connectors.

Consolidate the tech stack

Simplify the Azure technology stack that included petabytes of data stored across multiple SQL, NoSQL, and cache systems powering ~500 services, data science, machine learning, and observability.

Solution

Improved user experience, CTR, and revenue

After discovering through evaluation that Aerospike's enterprise solution would help them overcome Myntra's limitations in namespace, node expansions, scaling, and infrastructure costs, Myntra migrated more than 25 clusters and hundreds of services onto the Aerospike Database Enterprise Edition. The migration was seamless with zero downtime and kept everything in sync.

As a result, Aerospike significantly improved the performance and scalability of their feature lookup, resulting in improved customer experience, higher click-through rates (CTR), and stronger revenue.

P99 response times under 5ms

Aerospike excelled in performance compared to other data stores and provides P99 response times of less than 5ms

Efficient, consolidated footprint

Reduced to 8TB, 25 production clusters, 150 different nodes and overall reduction in cloud resource footprint

Cross datacenter replication and strong consistency

Hybrid Memory Architecture™ (HMA) allowed data to be kept on SSD and primary index in memory

Strong partnership

Exceptional support and engineering involvement with quick response and resolution times; willingness to help with new use cases

Results

Reliable performance in business-critical services

Today, Myntra views Aerospike as a trusted partner with reliable database technology for their most critical workloads. Myntra plans to expand their use of Aerospike, continuing to consolidate clusters, expand their use of XDR to improve disaster recovery preparedness, leverage enterprise security and encryption capabilities, and potentially upgrade to a container based architecture.

> Any time you search on the Myntra app, Myntra hits an Aerospike cluster on the back end. [It handles] millions of operations per second at a P99 read as well as write. Reads are submillisecond [latencies].
>
> —Nikhit Nair | Associated Architect | Myntra

Millions of search operations per second

Search is a business-critical service that requires millions of operations per second and P99 of under 1ms reads.

Millions of transactions per second

Builds ML models from features and is spread over 8 clusters with tens of millions of transactions per second

Billions of notifications daily

Sends out a billion notifications every day with users receiving, on average, 5 notifications a day, all processed within a 2-3 minute window

500K personalization operations per second

Supports half a million operations per second at sub millisecond response times for homepage personalization

About Aerospike

Aerospike is the real-time database built for infinite scale, speed, and savings. Our customers are ready for what's next with the lowest latency and the highest throughput data platform. Cloud and AI-forward, we empower leading organizations like Adobe, Airtel, Criteo, DBS Bank, Experian, PayPal, Snap, and Sony Interactive Entertainment. Headquartered in Mountain View, California, our offices include London, Bangalore, and Tel Aviv.

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

2440 W. El Camino Real, Suite 100, Mountain View, CA 94040 | (408) 462-2376