



Top 10 alternatives that

outshine Couchbase



Introduction

[Couchbase](#) is recognized for its performance as a distributed document NoSQL database with SQL-like querying capability. While it was initially attractive for organizations seeking real-time processing capabilities and flexible data modeling, Couchbase has found its place in large-scale consumer-facing applications and back-end enterprise applications. However, as businesses grow and workloads become more complex, limitations in Couchbase's architecture — such as escalating costs, operational complexities, and constraints in certain performance scenarios — become apparent.

For instance, Couchbase's reliance on RAM to achieve high performance can drive up hardware costs as data volumes increase. Because it requires manual data rebalancing of clusters during scaling or node failures, it's more work for administrators. At the same time, because it is primarily

JSON-based and operations on binary data are limited (let alone for graph data structures), it can be harder to use with artificial intelligence (AI) driven applications. These challenges have led enterprises to explore scalable, modern alternatives that are more flexible and efficient.

This e-book details the top 10 database alternatives to Couchbase. Each alternative has been selected based on its ability to address Couchbase's problems, whether you need to reduce the total cost of ownership (TCO), lower latency for real-time workloads, or prepare for future AI and machine learning (ML) requirements.

[A 2024 study](#) cited a lack of AI skills and high costs/budget limitations as the top two obstacles for organizations delivering on their AI strategy. The alternatives presented here have been evaluated not only for their features but also for their ability to continue working into the future. By understanding these options, organizations can reduce infrastructure costs while meeting evolving demands.



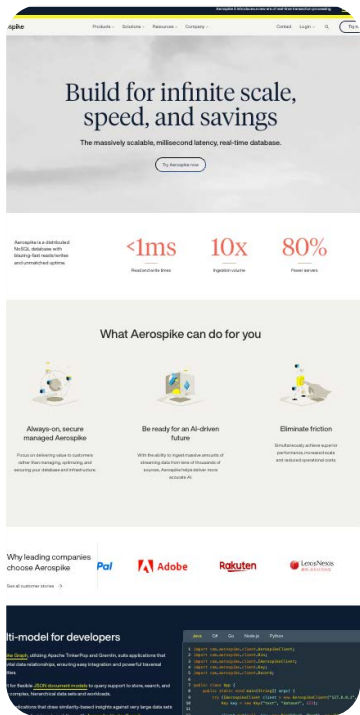


Aerospike

[Aerospike](#) is a distributed, multi-model NoSQL database that handles real-time big data applications with low latency and high reliability. Using its patented Hybrid Memory Architecture (HMA), Aerospike combines DRAM and SSDs to deliver high performance for less money. As of the release of Aerospike 8, it also processes distributed ACID transactions with high performance.

The [“Five signs you’ve outgrown Couchbase”](#) white paper shows reduced node requirements between 66% and 85% compared with Couchbase architectures. Companies moving to Aerospike have seen a 300% improvement in throughput while slashing infrastructure costs by millions annually.

Strengths



Predictable low latency

Offers sub-millisecond response times, making it suitable for mission-critical, mixed workload applications.



Efficient scalability

Multithreaded architecture lets organizations handle up to petabytes of data with 80% fewer nodes than Couchbase.



Advanced data models

Supports key-value, document, [graph](#), and vector data models for more developer flexibility.



Global availability

Supports synchronous active-active replication for consistent worldwide performance.



<1ms

Read and write times

10x

Ingestion volume

80%

Fewer servers

Challenges

- A steeper learning curve for new developers.
- Requires planning to take advantage of its full potential for hybrid memory configurations.

Use cases

- **AI/ML:** Build AI/ML apps ranging from retrieval augmented generation, semantic search, recommendations, and feature stores.
- **AdTech:** Powering real-time bidding platforms and user profile stores.
- **Financial services:** Enabling [fraud detection](#), real-time analytics, and digital payments processing.

Competitive edge

Aerospike's ability to reduce hardware requirements while maintaining high throughput makes it particularly suitable for enterprises focused on sustainability and a low TCO.

“

Aerospike has done the unthinkable: they cut our server footprint by a factor of six while boosting our performance 300%. The resulting total cost of ownership is saving us over a million dollars a year.

—Guy Almog | Head of IT Engineering | Playtika

”





Amazon DynamoDB

[Amazon DynamoDB](#) is a fully managed, serverless NoSQL database service known for its high performance. As part of the AWS ecosystem, it is particularly appealing to organizations already using AWS.

DynamoDB runs more than 1 million active applications globally, with industries such as e-commerce and finance using its low-latency operations to satisfy its customers. They cite having hundreds of customers running at ½ million+ transactions per second.

Strengths



Serverless architecture

Eliminates infrastructure management, allowing developers to focus on application logic.



Global tables

Multi-region replication with strong consistency provides high availability and disaster recovery capabilities.



Integration with the AWS ecosystem

It integrates with AWS services such as Lambda, Kinesis, and S3, making it suitable for AWS shops.

The screenshot shows the Amazon DynamoDB website with the following content:

- Amazon DynamoDB**: Serverless, NoSQL, fully managed database with single-digit millisecond performance at any scale.
- Why DynamoDB?**: A central graphic with 'DynamoDB' in the center, surrounded by icons for performance, scale, and integration.
- Benefits of DynamoDB**:
 - Performance at scale**: Deliver applications with consistent high performance and low latency at any scale.
 - Designed for 99.999% SLA**: Offers the highest availability and uptime in the industry.
 - Scale-to-zero serverless**: Deliver applications with no capacity reservation and pay only for what you use.
 - Easy AWS integration**: Seamlessly integrates with other AWS services like Lambda, S3, and Kinesis.
- Scale and value of DynamoDB**: A section at the bottom of the page.



1M+

Customers

1/2M+

Requests/sec delivered for
100s of customers

200TB+

Table size used by
100s of customers

Challenges

- A complicated pricing structure, including charges for read/write units and additional features like DynamoDB Accelerator (DAX), makes predicting cost at scale difficult.
- Performance for mixed workloads isn't as good as that of databases optimized for high-throughput applications.

Use cases

- **Software application development:** Build user data caches for internet applications demanding high concurrency.
- **E-commerce:** Managing shopping carts, order processing, and customer profiles.
- **Gaming:** Player data, session history, and leaderboards.

Competitive edge

DynamoDB's serverless model and AWS integration make it suitable for AWS shops seeking a fully managed [cloud-native product](#).

“

When building a storage system, you have to think about components like replication, backups, and capacity management. By using Amazon DynamoDB and Amazon S3, we simplify these problems because AWS handles many of the complex tasks.

—Jonathan Lee | Alki Team Tech Lead | Dropbox

”





Apache Cassandra

[Apache Cassandra](#) is a scalable, open-source NoSQL database designed for fault tolerance and linear scalability. Known for its ability to handle large volumes of data, Cassandra is best suited for write-intensive applications.

Cassandra processes billions of requests daily for organizations like Netflix, with 900+ clusters and over 22,000 nodes. Its wide-column storage model efficiently handles structured and semi-structured data.

Strengths



Scalability

Built for horizontal scale; readily able to add nodes.



Fault tolerance and high availability

Data replicated automatically across nodes with support for multi-region clusters and availability even with node failures.



Distributed architecture

Provides high availability and disaster recovery.



800K ops/sec

90% write, 10% reads

600K ops/sec

50% write, 50% reads

Challenges

- It is complicated to run and maintain, especially for large clusters.
- Performance isn't as good for read-heavy or mixed workloads.

Use cases

- **Financial trading systems:** Tracking real-time market data and transactions with high availability.
- **Messaging Systems:** Deliver highly personalized communications within applications such as gaming.
- **Recommendation engines:** These are based on user profiles and preferences to drive engagement.

Competitive edge

Cassandra's decentralized architecture and scaling capabilities make it a good choice for enterprises with large real-time data demands and high write loads.

“

Since 2016, application engineers at Bloomberg have turned to Cassandra because it's easy to use, easy to scale, and always available. We serve up more than 20 billion requests per day on a nearly 1 PB dataset across a fleet of 1,700+ Cassandra nodes.

—Isaac Reath | Software Engineering Team Lead,
NoSQL Infrastructure | Bloomberg

”





Microsoft Azure Cosmos DB

[Microsoft Azure Cosmos](#) DB is a globally distributed database service designed for high availability, scalability, and low latency performance. As a fully managed NoSQL and relational database, Cosmos DB supports multiple data models, including document, key-value, graph, and vector. With industry-leading SLAs, multi-region writes, and tunable consistency models, it provides enterprises with high availability and data sovereignty while maintaining optimal performance for mission-critical applications.

Strengths



Seamless integration with Microsoft ecosystem

Deeply integrated with the broader Microsoft suite, including Azure AI, Power BI, and Microsoft 365, enabling organizations already in the Azure ecosystem to leverage unified data workflows and analytics.



Global active-active replication

Cosmos DB supports multi-region writes, ensuring low-latency access and automatic failover across distributed applications. Thus, it is a strong choice for globally distributed workloads.



Tunable consistency levels

It offers five consistency models—strong, bounded staleness, session, consistent prefix, and eventual—allowing developers to balance performance and data accuracy based on workload needs.



Challenges

- **Complex pricing model:** Cosmos DB's pricing is based on provisioned throughput (RU/s), which can be difficult to predict and may lead to unexpected costs if not carefully managed. Microsoft provides guidance on estimating RU/s, but workloads with fluctuating demand can experience unexpected charges.
- **Performance optimization for write-heavy workloads:** Cosmos DB requires careful selection of partition keys and indexing configurations to prevent hot partitions and maximize throughput efficiency in write-intensive scenarios.

Use cases

- **IoT and telematics:** Processes and stores high-velocity [Internet of Things \(IoT\)](#) data streams in real-time with low-latency reads and writes.
- **Retail and e-commerce:** Manages product catalogs, customer data, and personalized recommendations with global consistency and availability.
- **Gaming:** Supports real-time game state management, player leaderboards, and in-game transactions with millisecond latency. Cosmos DB automatically scales to handle traffic spikes during game launches and updates, ensuring seamless performance for millions of concurrent players.

Competitive edge

Azure Cosmos DB surpasses Couchbase with its globally distributed, multi-active replication architecture, providing

seamless, low-latency access to data across multiple regions. Its support for multiple data models—document, key-value, graph, and column-family—offers greater flexibility than Couchbase's document-focused approach. Also, its five tunable consistency levels provide more granular control over performance and data consistency, while Couchbase primarily operates with eventual consistency.

<10ms

Read/write time latency SLA

99.999%

Availability

“

Because we run on Azure, our software grows with our client. The growth is very predictable, so we know how much we need to provision ahead of time.

—Yoni Nijs | Cofounder, CTO, and Product Lead | Zero Friction

”





Apache CouchDB + MemCached

[Apache CouchDB](#) is a NoSQL document database designed for high availability, distributed scalability, and ease of replication. It uses a multi-master architecture and a RESTful API, making it an excellent choice for web and mobile applications. [MemCached](#), a high-performance, distributed caching system, is often paired with CouchDB to accelerate data retrieval by reducing read latency and offloading query loads. Together, they offer a powerful combination of durable storage and in-memory caching for optimized performance.

Strengths



Efficient caching for high-speed data retrieval

MemCached enhances CouchDB's performance by reducing database reads, improving response times, and increasing throughput for high-traffic applications. It supports horizontal scaling, allowing multiple nodes to share the cache workload efficiently.



Multi-master replication for distributed scalability

CouchDB's built-in replication ensures fault tolerance and data consistency across distributed environments, making it well-suited for globally distributed applications.



Lightweight and simple deployment

MemCached's in-memory architecture is designed for minimal resource consumption, making it an effective caching layer without introducing significant operational complexity.



Challenges

- **Limited query capabilities:** CouchDB's querying capabilities are more limited than Couchbase, requiring Mango queries or MapReduce for complex queries.
- **Cache volatility and data loss risks:** MemCached does not provide data persistence, meaning cached data is lost in the event of a server restart or failure, which can impact performance when reloading from CouchDB.

Use cases

- **Content management systems (CMS):** CouchDB's document model and MemCached's fast caching enable dynamic content delivery with reduced database load.
- **E-commerce platforms:** This combination supports real-time product catalog updates and session management, ensuring a smooth shopping experience for users.
- **Mobile and offline-first applications:** CouchDB's synchronization capabilities and MemCached's rapid data access provide an efficient solution for apps that require offline functionality with periodic syncing.

Competitive edge

The CouchDB + MemCached combination offers a flexible and scalable alternative to Couchbase by leveraging document storage with high-speed caching. While Couchbase provides an integrated caching layer and native indexing, CouchDB with MemCached allows developers to scale caching independently, optimizing for workloads with fluctuating demand.

“

The replication in CouchDB is really the killer feature that sets it apart from other databases.

—Patrick Wolf | Hoptree

”

“

I like that MemCached allows for storage in RAM. It is key-value, but storing a JSON string is easy and quick. Especially so when using for its intended purpose as a cache without worrying about persistence. It's written in C, very close to the metal. Additionally, development on it is very active - with the most recent stable release a few months ago. I especially love that it can distribute its hash table across multiple machines, thus improving speed and reducing overall workload.

—G2.com review

”





DataStax Enterprise

[DataStax Enterprise \(DSE\)](#) is a scalable, cloud-native NoSQL database built upon Apache Cassandra. It is optimized for deploying critical workloads both on-premises and on [Kubernetes](#). It provides native [vector search](#) capabilities tailored for modern applications involving generative AI (GenAI). DSE combines operational simplicity, global scalability, and powerful performance, helping businesses efficiently manage data-intensive applications across distributed environments.

Strengths



Increased developer flexibility and productivity

Empowers developers by offering familiar APIs, extensive tooling, and seamless integrations with modern application stacks, accelerating application development and reducing complexity.



Cloud-native optimization

Designed explicitly for deployment in cloud and Kubernetes environments, simplifying operations and enhancing deployment flexibility.



Built-in GenAI capabilities

This feature provides native vector search features, ideal for modern AI-driven workloads requiring real-time data retrieval and integration.



11M

Rows of data,
consolidated from 450M

21K

Transactions per second

99%

Uptime requirement achieved

Challenges

- **Operational complexity:** DSE inherits Cassandra's operational overhead, requiring careful tuning and configuration, which can increase complexity and management efforts.
- **High infrastructure requirements:** Requires significant hardware resources, including storage and memory, leading to higher total cost of ownership at large scales.

Use cases

- **Healthcare:** Enables real-time data integration and analytics for improved patient care, personalized treatment plans, and streamlined operational workflows.
- **Gaming:** Powers real-time interactions and immersive experiences by handling high-throughput data and instantaneously responding to player actions.
- **Supply chain and logistics:** Provides real-time visibility and analytics, optimizing operational efficiency, reducing latency, and enhancing responsiveness across complex supply chain networks.

Competitive edge

DataStax Enterprise provides distinct advantages over Couchbase regarding real-time AI support, cloud-native architecture, and robust change data capture (CDC) for Cassandra, enhancing data integration capabilities.

“

We pulled in one year of data from Oracle and once we got it into DataStax Enterprise...it was a smooth ride and it was processing at a very high rate.

—Mukram Aziz | Senior Manager of Data Services | Capital One

”





Google Cloud Firestore

[Google Cloud Firestore](#) is a fully managed, serverless NoSQL document database designed for web, mobile, and server-side applications. Built on Google Cloud infrastructure, Firestore provides seamless scalability, automatic synchronization, and real-time data access. It is deeply integrated with Firebase and Google Cloud services, making it a popular choice for developers building cloud-native applications.

Strengths



Seamless real-time synchronization

Firestore has built-in, live data synchronization across devices and clients, making it ideal for collaborative applications with support for event-driven architectures via Cloud Functions for Firebase.



Fully managed serverless architecture

With automatic scaling, Firestore removes the need for database management, reducing operational overhead.



Deep integration with Google Cloud services

Firestore natively integrates with Firebase, BigQuery, and Google AI/ML tools, streamlining development for data-driven applications.



Challenges

- **Unpredictable cost scaling:** Firestore's pricing model is based on operations (reads, writes, and deletes), which can lead to unexpectedly high billing costs due to improper indexing or query optimization during usage spikes.
- **Latency and performance constraints at scale:** Firestore is optimized for ease of use rather than raw performance and may experience latency issues at high scale, making it less suitable for high-throughput, low-latency workloads compared to alternatives like Aerospike.

Use cases

- **Mobile and web applications:** Firestore's real-time syncing and offline capabilities make it a strong choice for mobile-first applications.
- **Gaming and leaderboards:** Firestore's ability to sync real-time player data makes it well-suited for multiplayer games and dynamic leaderboards.
- **Data analytics:** Firestore integrates seamlessly with BigQuery, enabling real-time data ingestion for analytics and analysis for large-scale analytics workloads, business intelligence, and machine learning applications with BigQuery.

Competitive edge

Google Cloud Firestore excels in real-time data synchronization and effortless scalability, making it a compelling choice for cloud-native applications.

99.999%

Availability SLA

“

Once we implemented our new statistics processing system, we were able to update our contributors' site metrics much faster... By providing this granular level of data to our contributors, we are helping them better optimize their content and deliver the best possible pieces to their readers.

—Benjamin Harrigan | Software Architect | Forbes

”





MongoDB

[MongoDB](#) is a flexible, high-performance NoSQL database known for its ease of use and developer-friendly features. It was designed to replace relational databases in scenarios where flexible schema, scalability, and document-oriented data models provide advantages over traditional table-based storage. MongoDB offers scalability and integration with over 100 technologies. As of 2024, the company reported \$1.7 billion in annual revenue, with 68% generated from its database-as-a-service (DBaaS) platform.

Strengths



Enterprise-grade adoption

Frequently shortlisted by enterprises for its broad applicability across industries.



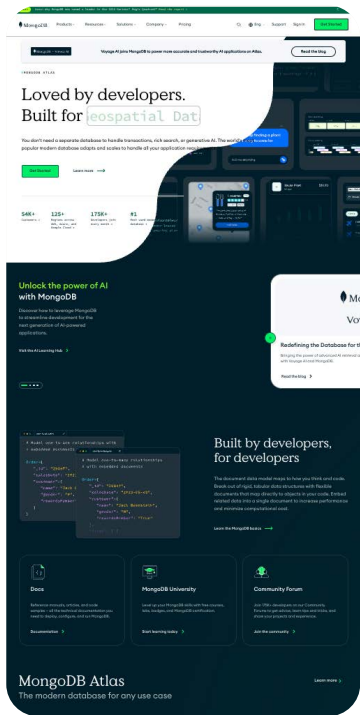
Extensive tech stack integration

Offers over 100 integrations with various platforms and tools.



Developer-friendly

Easy to learn, with abundant community resources and documentation.



99.995%

Uptime SLA

20%

Cost savings (Shutterfly)

14K

Requests per second (Beamable)

Challenges

- **Scaling beyond operations:** To scale horizontally, adding more shards is required. While this is primarily an operational task, the effectiveness of scaling depends heavily on initial data modeling and sharding key design. If these aren't optimized, it can lead to imbalanced clusters where some nodes handle disproportionately large data loads, causing performance bottlenecks and necessitating complex application changes.
- **High hardware requirements:** MongoDB recommends a replication factor of three (3) and advises allocating 50% of RAM as a guideline to keep frequently accessed data in memory. This leads to higher hardware demands, increasing infrastructure costs, and limiting scalability for resource-constrained deployments.

Use cases

- **Artificial Intelligence (AI):** Integrates vector search and real-time analytics for AI-driven applications.
- **Edge Computing:** Enables low-latency data processing for real-time decision-making at the edge.

- **Internet of Things (IoT):** Manages large-scale real-time sensor data efficiently.

Competitive edge

MongoDB offers flexibility and ease of use but is susceptible to performance degradation with scale due to the need for proper shard key selection and efficient indexing. It surpasses Couchbase with advanced analytics, mature cloud-native features, and broader market adoption for easier integration and developer support.

“

MongoDB Atlas gives us the ability to scale vertically and horizontally with relatively little effort and without huge operational overheads.

—Ali El Rhermoul | CTO | Beamable

”






Redis Enterprise

[Redis Enterprise](#) is a high-performance, in-memory NoSQL database designed for caching, real-time analytics, and message brokering. Built on the open-source Redis foundation, Redis Enterprise extends its capabilities with enhanced scalability, reliability, and persistence options for enterprise applications. As of 2024, Redis has raised \$356.8 million in funding across ten rounds, serving over 8,500 customers and reaching annual revenues of \$154 million.

Strengths



In-memory performance

Offers sub-millisecond latency and high throughput, making it ideal for real-time data processing and caching scenarios.



Rich modular ecosystem

Extends functionality through modules such as RedisJSON, RedisTimeSeries, and RedisAI, enhancing its versatility for diverse applications.



Well-known cache

Redis is renowned for its role as a high-speed in-memory cache, widely adopted for reducing latency in web applications, accelerating data access, and supporting high-performance use cases in industries like finance, e-commerce, and gaming.



99.999%

Availability SLA

75%

Savings vs. relational DBs

350%

ROI over three years

Challenges

- **Costly for larger datasets:** Since Redis is employed largely as a cache with data in memory, scaling workloads beyond two terabytes can become prohibitively expensive.
- **Weak reliability and consistency:** Redis relies on eventual consistency and lacks strong built-in high-availability features, making it less suitable for mission-critical applications where data consistency and uptime are essential.

Use cases

- **Real-time leaderboards:** Processes large volumes of streaming data for online gaming, stock market trading, purchase trends, and ranking.
- **Caching for high-performance applications:** Reduces latency for frequently accessed data in web applications, enhancing user experience in sectors such as e-commerce and gaming.
- **Session management:** Stores and manages user session data efficiently, ensuring quick data retrieval and seamless experiences for online platforms.

Competitive edge

Redis Enterprise outperforms Couchbase in real-time data processing and sub-millisecond latency performance, making it a superior choice for caching and high-speed transactions. Its modular architecture, combined with enhanced scalability features and enterprise-grade support, offers broader flexibility than Couchbase.

“

Redis Enterprise is a game-changer and a wonderful partner. It is a fast, high-performance database that has improved our Apache Spark document processing pipeline and enabled more accurate hybrid semantic search for our document chunk embeddings.

—Taqi Jaffri | Co-founder and Head of Product | Docugami

”





ScyllaDB

[ScyllaDB](#) is a high-performance, distributed NoSQL database designed as a drop-in replacement for Apache Cassandra. Built on a C++ codebase with a fully asynchronous, shared-nothing architecture, ScyllaDB offers ultra-low latencies and high throughput, making it ideal for real-time, data-intensive applications. It is optimized to leverage modern hardware, including multi-core processors and SSDs, providing consistent performance across large-scale workloads.

Strengths



Flexible wide-column data model

ScyllaDB's distributed wide-column architecture enables efficient reads and writes for large-scale workloads for semi-structured or changing schemas.



Optimized for modern hardware

ScyllaDB delivers superior performance and resource efficiency compared to Cassandra through multi-core optimizations, SSD optimizations like IO scheduling, and advanced networking techniques such as multi-queue NIC support.



Compatibility with Cassandra and DynamoDB

ScyllaDB supports Cassandra Query Language (CQL) and DynamoDB APIs, simplifying migrations and ensuring smooth integration with existing applications.



1.7ms

Latency

180K

Transactions per second

Challenges

- **Tunable consistency model:** ScyllaDB primarily focuses on availability and partition tolerance (AP), making it less suitable for use cases requiring strong consistency without performance compromises.
- **Limited multi-model support:** ScyllaDB primarily supports wide-column data structures, offering key-value stores with limitations. It has no native support for document or graph data models—only via third-party integrations.

Use cases

- **Real-time analytics:** Processes high-velocity data streams for applications in AdTech, fraud detection, and customer behavior analysis, providing near-instant insights.
- **Time series and IoT data management:** Manages time-series data from millions of IoT devices with high-speed writes and Time-to-Live (TTL) for data expiration.
- **Recommendation engine:** Creates personalized recommendations for onsite users to increase conversions for commerce.

Competitive edge

ScyllaDB outperforms Couchbase in handling real-time, high-throughput workloads via efficient hardware utilization and auto-sharding. It has fewer nodes and better active-active support with its multi-site clustering with synchronous replication.

“

On just 3 nodes, we can scale up to 35 of the c5.xl instances and we'll process over 6500 records per second, while also running the real-time workloads.

—Dan Podhola | Principal Software Engineer | Zillow

”






Explore new alternatives to Couchbase

Choosing the right alternative to [Couchbase](#) depends on your organization's unique requirements, including performance needs, scalability demands, and budget considerations. Each alternative discussed in this e-book offers distinct advantages designed to address Couchbase's limitations, whether by enhancing real-time performance, reducing infrastructure costs, or simplifying operational complexity. By carefully evaluating these options, businesses can unlock greater efficiency, scalability, and future-readiness for modern data workloads.





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>.

©2025 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.

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