

White Paper

AEROSPIKE

---

# Achieving resiliency with Aerospike's Real-time Data Platform

## Contents

Executive overview	2
Database challenges in an “always-on” enterprise	2
Common business objectives	3
Popular technologies and architectures	4
The Aerospike approach	5
Core server architecture	6
Multi-site clustering	10
Cross Datacenter Replication (XDR)	13
Archive	15
Technology summary	16
Customer profiles	17
AppNexus/Xandr	17
PayPal	17
European financial institution	18
Nielsen	18
Verizon Media/Yahoo	19
Summary	19
About Aerospike	21

## Executive overview

Downtime: it's the bane of global enterprises. Lose immediate access to your operational data, and you lose business. It's that simple.

But providing ultra-fast, predictable access to real-time transactional data around the clock is no simple task. Software upgrades, hardware failures, network glitches, natural disasters, and usage spikes are common problems that threaten to compromise increasingly aggressive service-level agreements (SLAs). Layer those concerns against a backdrop of massively growing data volumes, globally distributed business environments that require flexible deployment options (on-premises, in the cloud, or hybrid configurations), and business imperatives to minimize the total cost of ownership (TCO). You have an incredible challenge that traditional database architectures just can't handle.

That's where Aerospike comes in. A distributed real-time data platform unlike any other, Aerospike provides 99.999% uptime (five nines), offers synchronous and asynchronous data replication for business continuity across multiple locations, automates routine maintenance and failover operations, and enforces strong, immediate data consistency in single-site and multi-site deployments. Beyond that, Aerospike scales readily from 100s of gigabytes to petabytes without massive server footprints (or associated costs) while delivering predictable, ultra-fast data access speeds for mixed workloads, even during periods of peak demand.

Indeed, leading firms in finance, retail, telecommunications, advertising technology (Ad Tech), and other industries have turned to Aerospike to power their "always-on" business models. In this paper, you'll learn about Aerospike's key technologies that provide firms with a high-speed, highly scalable platform that services critical applications around the clock with unparalleled cost-efficiency. You'll also explore how Aerospike customers use these technologies to power their businesses 24x7, often saving \$1 - \$10 million per application compared with other approaches.

## Database challenges in an "always-on" enterprise

As firms in many industries face business mandates to operate without downtime or lags in service, IT architects are re-evaluating their data platforms. After all, accessing, managing, and securing critical operational data is vital to business resiliency and continuity. Rendering such data inaccessible in the target timeframe can damage a firm's reputation, diminish customer loyalty, and cause substantial revenue losses.

Not surprisingly, IT architects routinely strive to build infrastructures that address different business continuity goals. Recovery time objectives (RTO) and recovery point objectives (RPO) are vital considerations. RTO identifies the tolerable downtime (if any) resulting from an outage, while RPO identifies the acceptable data loss (if any) resulting from an outage. Specific applications cannot tolerate downtime nor data loss, even in a natural disaster, such as a significant climatic event taking an entire data center offline. Thus, you would need multiple active

data centers at different locations with fully synchronous data replication and automated failover from a data platform perspective. By contrast, applications that can tolerate data loss (perhaps several seconds or minutes worth of transactional data) but little or no downtime can be well served by replicating data asynchronously across multiple systems at different locations.

## Common business objectives

Most enterprises employ multiple technologies to satisfy their varied business continuity and disaster recovery strategies. Table 1 identifies several common business objectives and popular technologies that support each.

	RTO	RPO	Data platform capability
<b>Core system resiliency:</b> Exceptional availability	99.999% uptime  RTO = 0	No data loss  RPO = 0	<ul style="list-style-type: none"> <li>● Self-managing, self-healing</li> <li>● Data redundancy</li> <li>● Balanced data distribution</li> <li>● Rolling upgrades</li> <li>● Strong consistency</li> </ul>
<b>Global high availability:</b> Protect mission-critical apps; multi-site	99.999% uptime  RTO = 0	No data loss  RPO = 0	<ul style="list-style-type: none"> <li>● Multiple data centers with synchronous replication</li> <li>● Active/active sites</li> <li>● Built-in load balancer for transparently routing traffic</li> </ul>
<b>Global business continuity:</b> Protect non-critical apps	High uptime (seconds - minutes)  RTO > 0	Data loss  RPO > 0	<ul style="list-style-type: none"> <li>● Multiple data centers with asynchronous replication</li> <li>● Active/active or active/passive sites</li> <li>● Granular replication options</li> </ul>
<b>Archive:</b> Provide long term retention; support compliance	Variable uptime (days or more)  RTO > 1 day	Data loss  RPO > 1 day	<ul style="list-style-type: none"> <li>● Backup / recovery tools</li> <li>● Support for low-cost archive targets (e.g., cloud, NFS)</li> </ul>

**Table 1: Reliability and availability levels commonly required for enterprise applications**

A deeper look at resiliency requirements for a real-time data management platform can lead to several critical questions, such as:

- To what extent can routine maintenance and upgrades be handled without downtime or impact on runtime performance?
- How does the system address common failures, e.g., a corrupt storage device, an unresponsive server node, etc.? How much operator involvement is required?

- In what ways can the data platform be deployed to gracefully cope with unpredictable scenarios, such as a natural disaster? How much operator involvement is required?
- How does the platform guard against spikes in data access latencies during routine operations and failure scenarios? What mechanisms enable the system to meet target SLAs during peak loads and seasonal spikes in demand?
- What types of data replication technologies are supported, and in what configurations? For example, are active/active and active/passive architectures supported across multiple cloud service providers, on-premises deployments, and hybrid configurations?
- How does the platform minimize runtime performance degradation for replicated data scenarios?
- Are data consistency guarantees available? How do these map to differing application requirements?

In a moment, you'll learn how Aerospike answers these and other critical questions. But first, let's review several popular technologies and architectural approaches.

## Popular technologies and architectures

### Clustered computing

Many modern data platforms support clustered computing environments with shared-nothing architectures for scalability and performance. Maintaining redundant copies of user data on a single system promotes high availability during many failure scenarios. However, some implementations are more cost-efficient than others. For example, a platform that recommends a replication factor of three for user data (one primary copy and two replicas) can deliver high data availability but incurs high TCO due to a large server footprint and corresponding operator overhead.

### Data distribution and rebalancing

Other architectural aspects also impact the extent to which data is available within target timeframes. For example, an application's ability to meet SLAs is affected by data distribution and data rebalancing techniques. The simplicity of accommodating upgrades to cluster hardware or software, the operational ease of adding or deleting nodes, and the ability to quickly route client requests to the correct server node are all of paramount importance. Multi-tiered platforms that employ a cache to the front-end of a relational or non-relational DBMS present additional availability considerations. Failure in one tier can slow or restrict operational access, potentially resulting in lengthy restart times, stale reads, or even data loss.

### Multi-site support

Real-time data platforms vary even more wildly when it comes to multi-site support. Data replication technologies, while common, often differ in certain areas:

- Data replication approach (synchronous or asynchronous).
- What data can be replicated (i.e., the degree of granularity supported)

- Where the data can be replicated (across cloud platforms, on-premises, and hybrid deployments).
- Replication process efficiency.
- Architectural topologies supported.

## **Active/active**

Quite commonly, replication technologies are deployed to support active/active or active/passive database environments. Active/active databases span two or more regions and service application requests at all locations, making each site “active.” Data records are replicated across regions so that reads can be processed anywhere. In some architectures, writes of a given data record are only handled at a single master location; other architectures allow such writes to occur at multiple locations. Each approach has its challenges involving availability, consistency, and performance.

## **Consistency**

For example, one consideration is, if writes for any given record are processed anywhere, how is the data kept synchronized and consistent? The two-phase commit protocol (introduced by distributed relational DBMSs in the 1990s) requires a global transaction coordinator to communicate repeatedly with each participant to ensure that the transaction is committed by all or none. The result is strong consistency is enforced but is costly and slow. Consequently, many non-relational database vendors offer eventual consistency, guaranteeing that eventually, all access to a given record will return the same value, assuming no other updates are made to that record before copies converge. But this comes at a cost. During normal operations – not just failure scenarios – readers may see stale data. Furthermore, suppose the same record is updated more than once before all copies converge. In that case, the database must resolve conflicts, which adds runtime overhead and results in data loss.

## **Active/passive**

Of course, not all enterprise applications require the high availability afforded by an active/active architecture. An active/passive approach consists of one active data center that processes all read/write requests. A passive database is maintained at a remote data center, standing by in case the active system fails. In such architectures, the passive database is typically updated asynchronously; failover may be automatic or require manual intervention.

## **The Aerospike approach**

Aerospike is a distributed multi-model NoSQL platform that provides fast and predictable read/write access to operational and transactional data sets. These data sets span billions of records in databases holding up to petabytes of data supporting systems of engagement or systems of record. Aerospike delivers exceptional availability and runtime performance with dramatically smaller server footprints. Aerospike is optimized to exploit modern hardware, including multi-core processors with non-uniform memory access (NUMA), non-volatile memory extended (NVMe) Flash drives, persistent memory (PMem), network application device queues (ADQ), and more. Such optimizations, coupled with Aerospike’s self-managing and self-healing

architecture, provide firms with distinct advantages for both on-premises and cloud deployments. Other features that distinguish Aerospike from many alternate platforms include its ability to automatically distribute data evenly across its shared-nothing clusters, dynamically rebalance workloads, and accommodate software upgrades and most cluster changes without downtime.

Aerospike’s high availability is one of the strategic benefits that draw users to its platform. Indeed, a presentation by [So-net Media Networks](#) attested to using Aerospike for “8 years of operation without a single stop.” Another presentation by [AppNexus/Xandr](#) on “What 100% uptime looks like” described how the firm relies on Aerospike to “be available every second of every day.” Similarly, a technology leader at [Cybereason](#) who migrated to Aerospike after using a different key-value platform asserted, “. . . what surprises us about Aerospike is how easy it is to work with and how easy it is to operate. Using Aerospike, we are able to create a big data solution, persistent and highly available.”

Let’s explore several critical Aerospike technologies in greater detail that enable the platform to deliver exceptional data availability without compromising performance or relying on large server footprints. We’ll start with Aerospike’s core server architecture and cover more specialized topics, such as its support for multi-site clusters and cross-datacenter replication.

## Core server architecture

The core architecture of any real-time data management platform is critical to its ability to support high levels of data availability and business resiliency. A white paper entitled [Introducing Aerospike's architecture](#) describes Aerospike’s internal design and essential features in detail, so we’ll just highlight a few architectural aspects shown in Fig. 1.

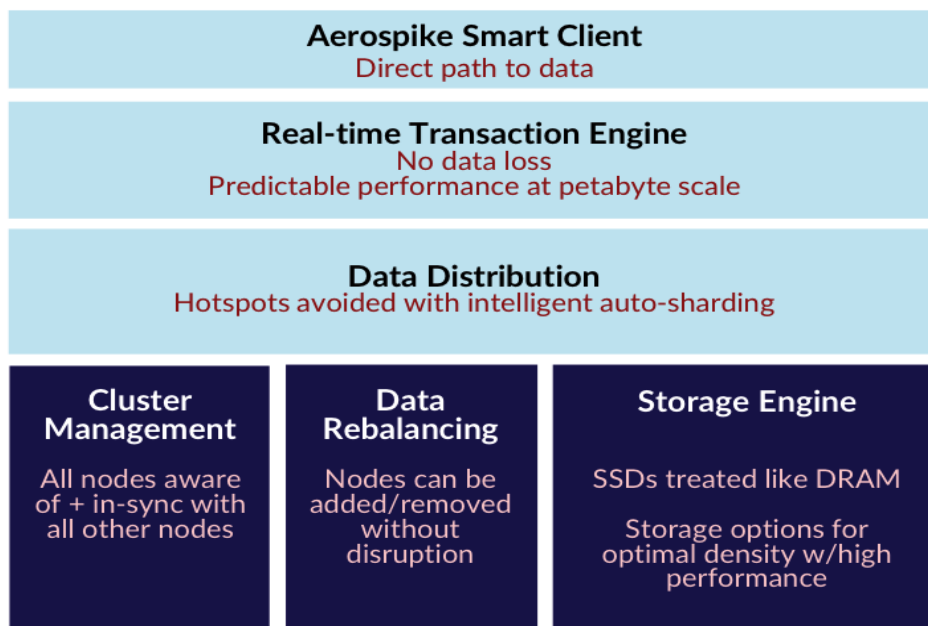
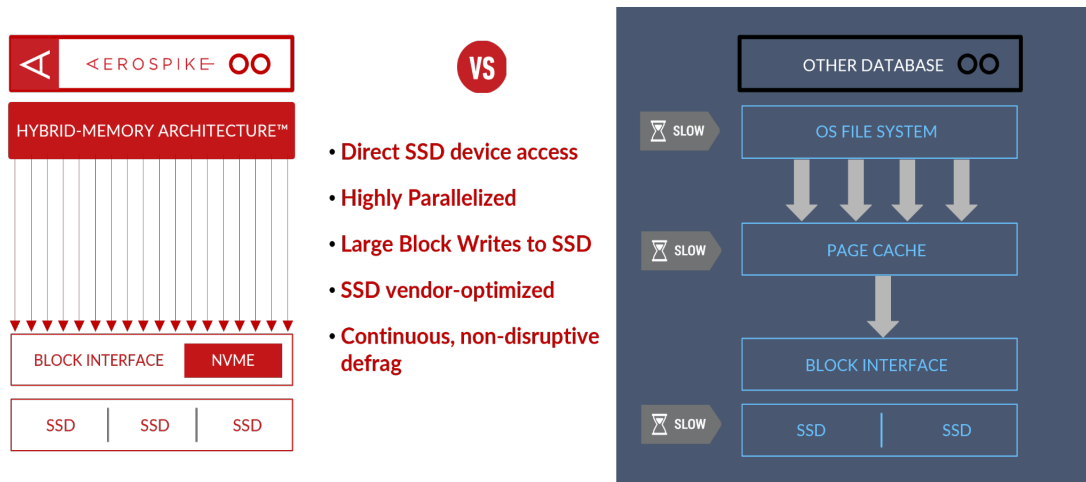


Figure 1: Aerospike’s core server architecture is designed for high data availability

As mentioned earlier, Aerospike is a shared-nothing, multi-threaded, NoSQL system designed to operate efficiently on a cluster of server nodes, exploiting modern hardware and network technologies to drive reliably fast performance, often at the sub-millisecond level. As referred to in the lower right of Fig. 1 and shown in greater detail in Fig. 2, Aerospike treats solid state drives (SSDs) as raw devices. Aerospike writes data in large blocks using a highly efficient custom file format that avoids wear-leveling issues common with other providers. Moreover, administrators can configure Aerospike’s storage usage to suit their cost, performance, availability, and scalability objectives. Users of Aerospike can keep indexes and data all in traditional memory (DRAM), all in PMem, all on SSDs, or in hybrid configurations (with indexes in DRAM or PMem and user data on SSDs). As of this writing, no other operational data platform offers such flexible options.




*Software written in C to natively talk to hardware, not an API layer*

**Figure 2. Aerospike exploits SSDs in ways other systems don't**

As referred to in the lower left of Fig. 1 and illustrated in greater detail in Fig. 3, Aerospike’s cluster management design ensures that all nodes are aware of the states of other nodes in the cluster through “heartbeats” and other mechanisms. This enables the system to automatically detect when a new node is added or when an existing node becomes inaccessible due to a hardware failure, a rolling software upgrade, a planned provisioning change, or other reason. A shift in cluster status triggers an automatic rebalancing of data to ensure high availability and strong runtime performance even during the rebalancing process itself. In addition, Aerospike’s built-in support for data redundancy is critical to ensuring that standard hardware or network failures won’t render essential data inaccessible or require manual intervention for recovery. Unlike some systems, Aerospike typically requires only 2 copies of user data (i.e., a replication factor of 2) to deliver exceptional uptime. Indeed, this is Aerospike’s default setting and what most of its production users deploy to achieve high availability on relatively small server footprints.

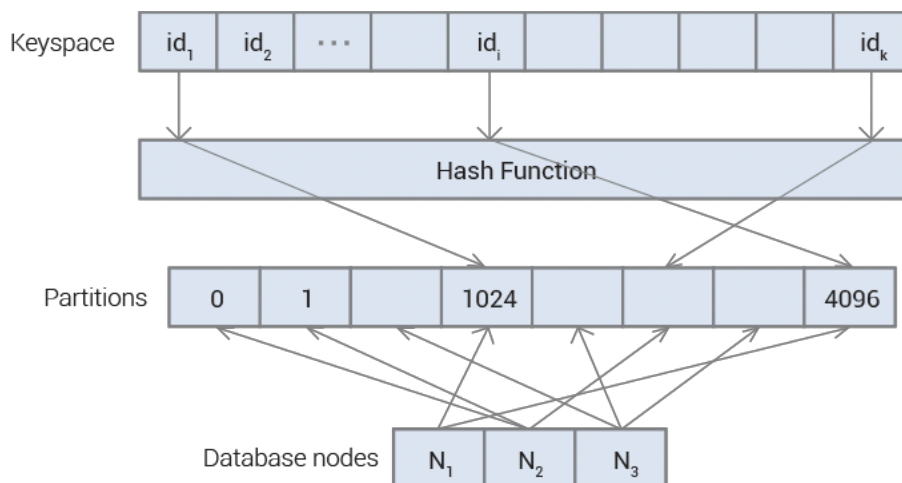


- Cluster discovers new node using Paxos Gossip protocol
- Partition table recalculated throughout cluster
- Migrations occur
- “Smart Client” updated with new partition map
- Remove a node and the opposite occurs!  
Replicas become primaries where needed + new replicas created to avoid downtime. Allows for rolling upgrades.



**Figure 3: Aerospike quickly and automatically responds to changes in cluster state**

Aerospike automatically employs a deterministic, uniform data distribution scheme to prevent data skewing and hot spots in the cluster. This scheme hashes each record’s primary key into a 160-bit digest using the RIPEMD algorithm and distributes the digest space into 4096 non-overlapping partitions (see Fig. 4.) This approach is very robust against collisions and promotes a well-balanced cluster. However, data access latencies become unpredictable without uniform data distribution, compromising the availability of data to specific applications and causing SLA violations. Furthermore, Aerospike’s data distribution technology supports “rack awareness,” which enables copies of user data to be stored on different hardware failure groups (i.e., different racks) located in different zones or data centers. Doing so further promotes high availability during common failure scenarios without operator intervention, as you’ll learn in a [subsequent section on multi-site clustering](#).



**Figure 4: Aerospike’s partitioning scheme distributes data evenly throughout the cluster**

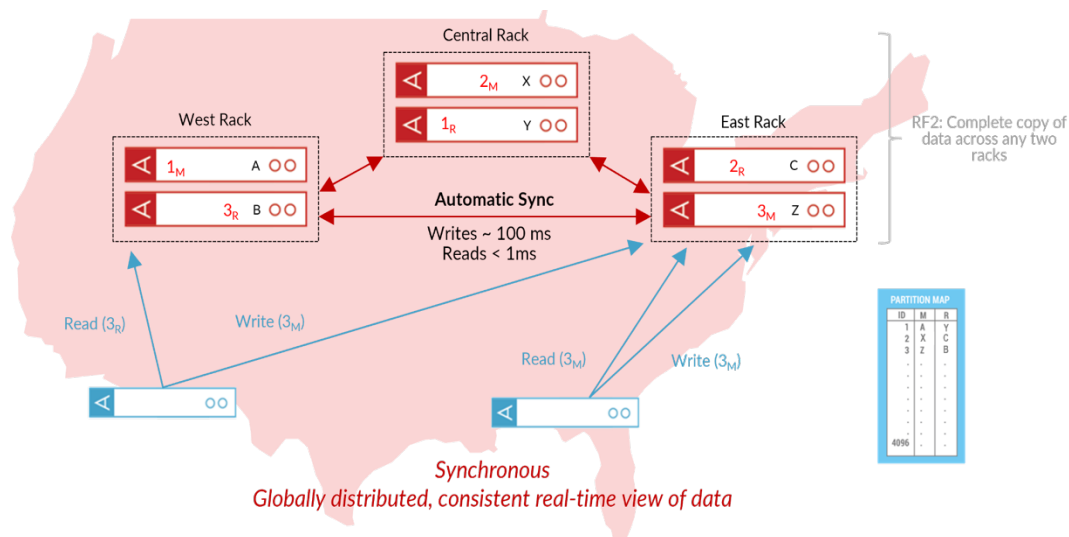
Aerospike's real-time transaction engine enforces immediate data consistency, even at petabyte scale. Administrators can configure namespaces (databases) to operate in strong consistency (SC) mode, preventing stale reads, dirty reads, and data loss while preserving ultra-fast data access. An earlier, slightly more relaxed data consistency mode (available / partition tolerant, or AP mode) is also supported for backward compatibility with early Aerospike releases. Regardless of consistency mode, any write to a record is applied synchronously to all replicas in a local cluster. Aerospike applies each record operation atomically. Restricting the transaction scope to a single record enables Aerospike to deliver strong consistency and exceptional runtime performance not seen with other approaches.

Aerospike's Smart Client™ layer includes specific optimizations that aren't common in other systems. For example, it maintains an in-memory structure that details how data is distributed in the active cluster. As a result, the Aerospike client layer can route an application request directly to the appropriate node, minimizing costly network "hops." By contrast, some systems use a proxy service to process and route each client request, creating performance bottlenecks and introducing a single point of failure. Still, other systems rely heavily on inter-node communications, with the node receiving the request re-routing it elsewhere. This, too, can impact the timeliness of data access.

The growing popularity of cloud deployments has prompted Aerospike to adopt various resiliency features as many users rely on Aerospike to meet stringent SLAs even in cloud environments, where virtualization can sometimes cause unpredictable behaviors. Many of these resiliency features are helpful in hybrid and on-premises deployments as well. For example, as described in [this blog post](#), Aerospike has integrated a series of start-up sanity checks to detect and flag configuration issues that can be particularly troublesome to resolve under heavy production workloads. Essentially, Aerospike logs warnings when it detects any known "anti-patterns"; if desired, administrators can use a configuration setting to enforce strict conformity to best practices. In addition, Aerospike has implemented various "circuit breakers" that temporarily suspend or throttle certain load-generating mechanisms when Aerospike is unexpectedly stressed, typically due to a transient resource problem (such as network congestion). As you might imagine, determining when it's best to throttle back rather than press ahead is no trivial task, so Aerospike employed its considerable field experience and engaged actively with customers to develop appropriate circuit breakers for specific situations.

## Multi-site clustering

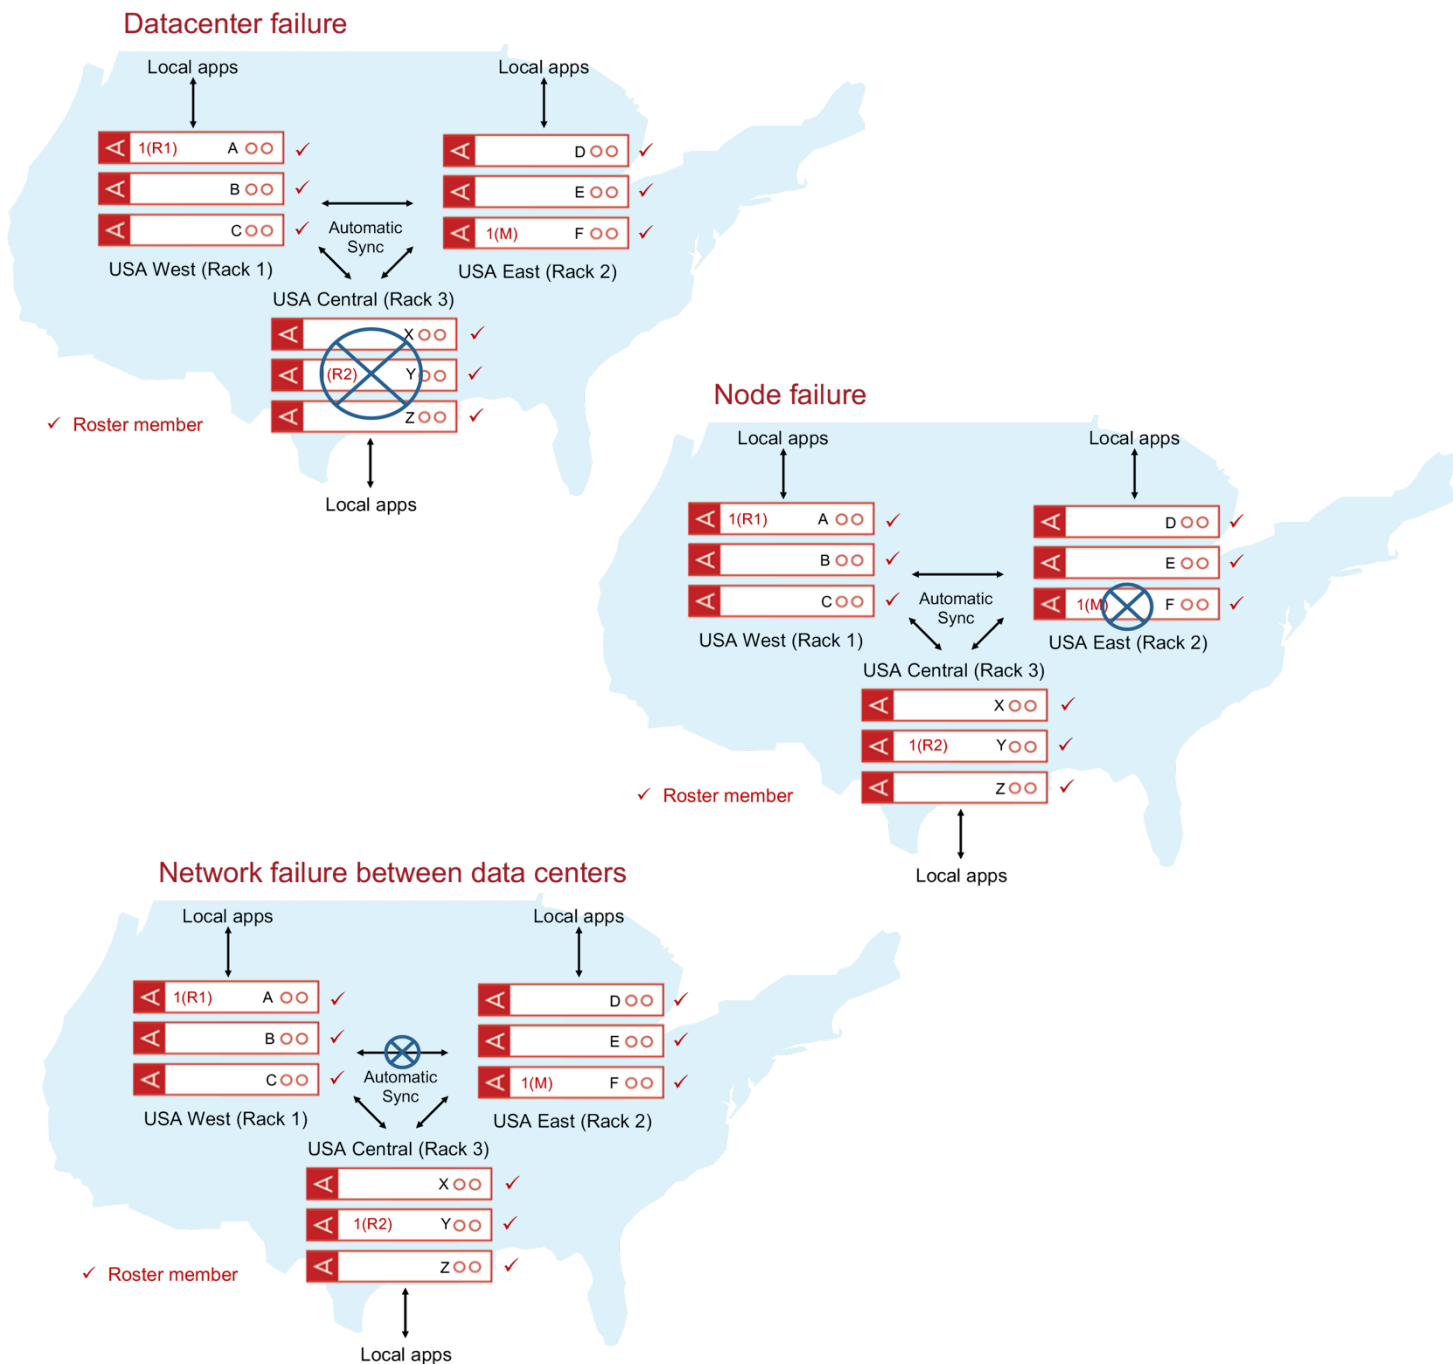
As the name implies, Aerospike's multi-clustering support enables a single cluster to span multiple geographies, as shown in Fig. 5. This allows users to deploy shared databases across distant data centers and cloud regions with no risk of data loss. Automated failovers, high resiliency, and strong performance are hallmarks of Aerospike's implementation. Two features underpin Aerospike multi-site clustering: rack awareness and strong consistency.



**Figure 5: A single Aerospike cluster can span multiple sites with no risk of data loss**

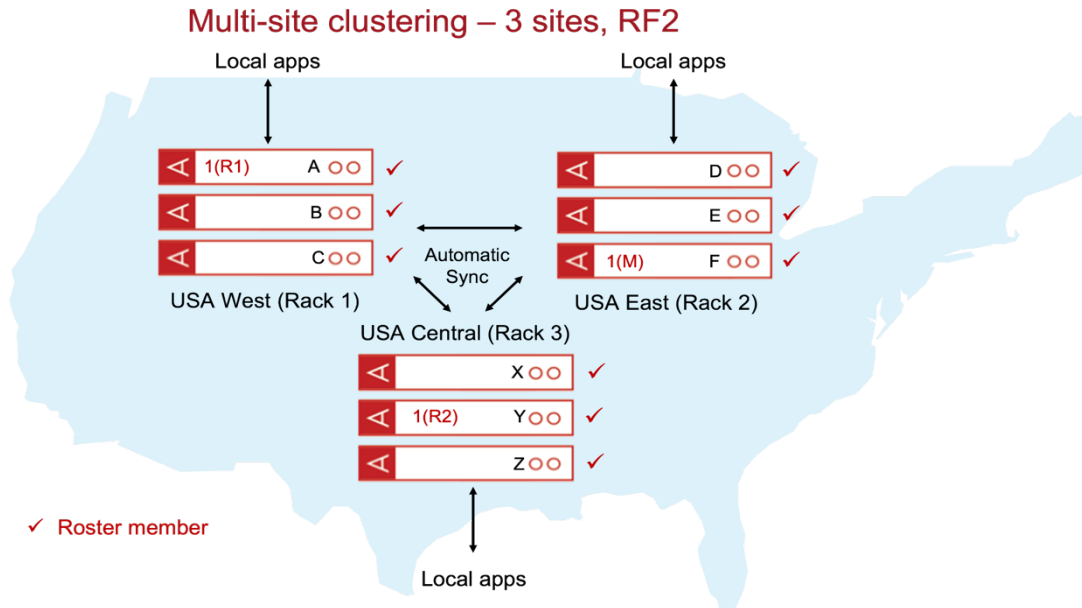
As discussed earlier, rack awareness allows replicas of data partitions to be stored on different hardware failure groups (different racks). As a result, administrators can configure each rack to store a full copy of all data to maximize data availability and local read performance. In addition, Aerospike will automatically and evenly distribute data across all nodes within each rack.

One node maintains a master copy of a given data partition at any time. Other nodes (typically on other racks) store replicas, which Aerospike automatically synchronizes with the master. Aerospike tracks the locations of masters and replicas using an internal roster; it also understands the racks and nodes of a healthy cluster. This enables Aerospike to detect automatically -- and quickly overcome -- various failures as shown in Fig. 6, such as the loss of a single node at a data center, a network failure between datacenters, or even the loss of an entire data center.



**Figure 6: Sample failure scenarios**

Let's look more closely at a sample multi-site cluster configuration. Fig. 7 illustrates a configuration in which each data center has one rack with three nodes, and each node has a copy of the roster. With a replication factor of three, this example shows the roster-master copy of the data in the partition on Node 3 of Rack 2; replicas exist on Node 1 of Rack 1 and Node 2 of Rack 3.



**Figure 7: Aerospike cluster spanning 3 racks at 3 different sites; each rack contains 3 nodes.**

Aerospike routes an application’s request to read a data record to its local data center’s appropriate rack/node. Referring to Fig. 7, an application in USA East seeking to read data will access the master copy in Rack 2 Node 3 with one network hop. A US Central application seeking the same record will enjoy “one-hop” access, as Aerospike will retrieve the data from the replica in Rack 3 Node 2. Sub-millisecond reads are common.

Write transactions are processed differently to ensure strong, immediate data consistency across the cluster. Like reads, writes can be initiated by any authorized application, but Aerospike routes each write transaction to the rack/node containing the current master of the data. The master node ensures that the write activity is reflected in its copy and all replica copies before the operation is committed.

Returning to Fig. 7, a write request for the data will be routed to Rack 2 Node 3 in USA East because it contains the master of the target record. Again, this will occur regardless of where the client application resides. As you might expect, the routing and data synchronization processes introduce some overhead. Writes won’t be as fast as reads, but they’ll still be fast: often a few hundred milliseconds or less, which is well within most target SLAs for distributed transactions.

Resiliency in the face of failure is a critical requirement of any multi-region operational database, which can experience a data center failure, a single node failure, a network failure between regions, etc. A [separate white paper](#) on Aerospike’s multi-site clustering technology walks you through several such failure scenarios. It explains how Aerospike automatically and seamlessly detects and reacts to cluster changes and how it forms a new sub-cluster within seconds to service read/write requests - even though a portion of the full cluster is unavailable. Under extreme situations, Aerospike may restrict certain operations to preserve data correctness; however, its

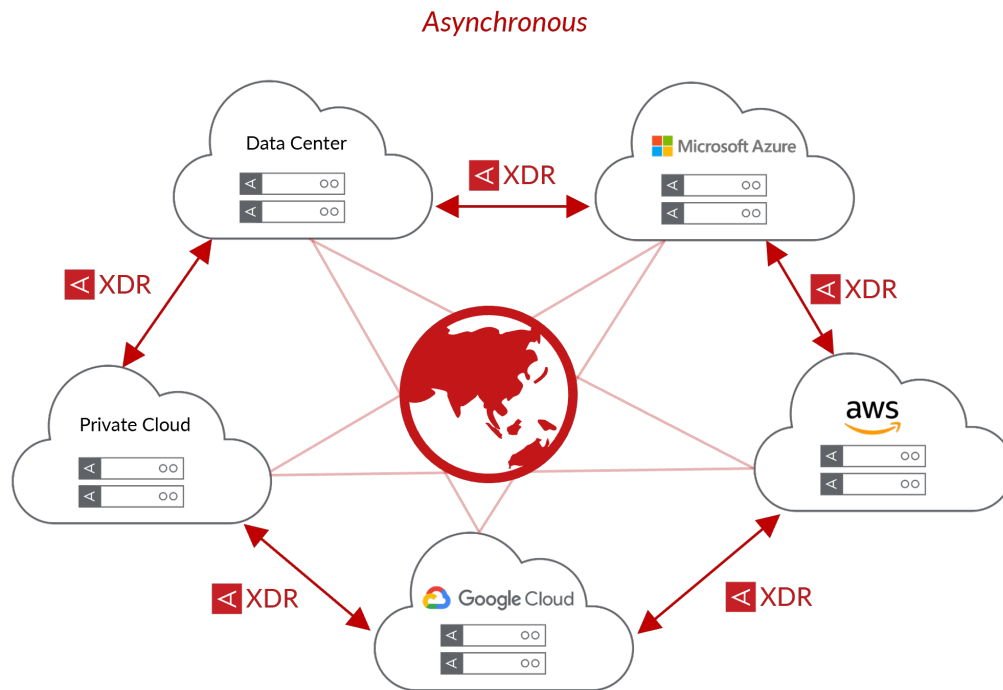
architecture is designed to afford the highest level of data availability while maintaining strong, immediate data consistency. While it's beyond the scope of this paper to explore specific multi-site cluster failure scenarios, we'll summarize Aerospike's general failover approach so you can understand the fundamentals.

When Aerospike detects a change in cluster status, it consults its roster to determine if the roster-master of the requested data is unavailable. If so, Aerospike designates a new master from the available replicas; this new master is usually on another rack. In addition, Aerospike automatically creates new replicas so that each data partition has the required number of replicas (based on the system's replication factor setting). This occurs in the background without impacting availability; high network bandwidth and pipelined data transfer speed this process. When the cluster is split into multiple sub-clusters, only one sub-cluster can accept requests for a given partition to ensure data consistency. Aerospike follows these rules to determine what the operational sub-cluster can do:

1. If a sub-cluster has both the master and all replicas for a partition, then the partition is available for both reads and writes.
2. If a sub-cluster has a strict majority of nodes and has either the master or a replica for the partition, the partition is available for both reads and writes.
3. If a sub-cluster has exactly half of the nodes and has the master, the partition is available for both reads and writes.

## **Cross Datacenter Replication (XDR)**

Cross-Datacenter Replication (XDR) transparently and asynchronously replicates data between two or more Aerospike clusters in different geographies. A site can be a physical rack in a data center, an entire data center, an availability zone in a cloud region, or a cloud region. Firms often use XDR to replicate data from Aerospike-based edge systems to a centralized Aerospike system. XDR also enables firms to support continuous operations during a crisis (such as a natural disaster) that takes down an entire cluster. As shown in Fig. 8, XDR operates across cloud platforms, on-premises data centers, and hybrid environments.



**Figure 8: XDR asynchronously replicates data across Aerospike platforms in the cloud, on-premises, and hybrid deployments**

To accommodate varying business requirements, XDR can replicate data at different levels of granularity. For example, firms can replicate one or more namespaces (databases), one or more sets (collections of records), as well as all or a subset of bins (fields) within records. Firms can even restrict the data to be shipped based on filtering expressions -- an important feature for compliance and location-specific privacy concerns. XDR's flexible data shipment granularity helps firms reduce network overhead, avoid overprovisioning CPU and storage at target destinations, and reduce the cost of moving cloud data. Furthermore, each Aerospike cluster using XDR can have its own dynamically defined list of destination clusters for data shipment, avoiding overhead and minimizing costs.

XDR replication is log-based. Applications write to their "local" cluster, and XDR logs minimal information about the change (not the full record). For efficiency, XDR batches change data and ships only the latest version of a record. Thus, multiple local writes for one record generate only one remote write -- an essential feature for "hot" data. XDR also keeps a pool of open connections for each remote cluster, pipelining replication requests in a round-robin fashion.

XDR avoids any single point of failure. Indeed, it ships changes as long as at least one node survives in the source and destination clusters. It also adjusts easily to new node additions in any cluster and can utilize all available resources equally.

Firms can define active/active or active/passive architectures using XDR. However, XDR replication is always asynchronous. With an XDR-based active/active configuration, applications can write data at any location. Conflicting writes can occur if the same data is updated concurrently at multiple sites, resulting in different values at different locations or loss of some data. Some applications can tolerate this relaxed form of data consistency. For those that cannot, firms may choose to restrict write access to specific data to a given location, thereby avoiding the potential for conflicting concurrent updates. Note that the XDR's asynchronous nature means that the most recent updates from the failed site may not be available at the other sites when a site failure occurs.

XDR is best suited to situations where high data availability with low data access latencies are mandatory. Its architecture trades strong data consistency between sites for low-latency reads and writes at all sites. A [solution brief](#) provides further details on XDR.

## Archive

Quite commonly, firms seek to offload inactive or cold data from their operational platforms to long-term storage for compliance or other reasons. Indeed, government regulations and corporate policies are driving firms to retain more data for longer periods of time and to employ an infrastructure that allows for rapid retrieval of archived data on demand.

To fulfill such requirements, Aerospike's backup facility enables firms to efficiently archive data to a network file system (NFS) for low-cost storage. Its built-in `asbackup` command enables administrators to create full-image archives of desired namespaces or sets, or incremental backups based on timestamps. Furthermore, administrators can restrict backups to data subsets matching a provided expression. Command options also enable administrators to compress data using `zstd` compression and to encrypt data using AES 128/AES 256 in Counter (CTR) mode with an accompanying private key. Restoring archived data created through Aerospike's backup tool is readily accomplished via Aerospike's `asrestore` command.

For added efficiency, Aerospike offers throttling options for both backup and restore operations. For example, a command option can be invoked to instruct write operations to backup files to not exceed the bandwidth specified (in MiB/s). This effectively throttles the scan on the server side, too, as the Aerospike's backup facility will not accept more data than it can write. A similar command option is offered for restore to throttle read operations from backup files so that they will not exceed the specified I/O bandwidth (in MiB/s) and its database write operations won't exceed the specified number of transactions per second.

Finally, administrators can selectively restore data from backups, limiting the scope of what will be restored according to business needs and avoiding any unnecessary overhead or data transfers. For example, restore options can limit the sets and bins to be restored and omit restoration of certain objects altogether, such as user-defined functions and secondary indexes.



## Technology summary

As you've seen, Aerospike provides a wealth of technologies to help firms achieve their varied data availability and business continuity goals. Table 2 revisits the common business goals described in an [earlier section](#) and outlines corresponding Aerospike technologies that enable firms to achieve these goals.

	Aerospike Technology
<b>Core system resiliency:</b> Exceptional availability, single site	<b>Aerospike Database 5</b> <ul style="list-style-type: none"> <li>● Self-managing, self-healing</li> <li>● 99.999% uptime with replication factor of 2</li> <li>● Balanced data distribution</li> <li>● Rolling upgrades</li> <li>● Node addition/removal without disruption</li> <li>● Immediate data consistency: SC or AP modes</li> <li>● Cloud or on premise</li> <li>● Smart Client™</li> <li>● Deep hardware exploitation; flexible storage options</li> <li>● High efficiency</li> <li>● Automated detection of “anti-patterns”</li> <li>● Ultra-fast, predictable performance</li> </ul>
<b>Global high availability:</b> Protect mission critical apps	<b>Aerospike Multi-site Clustering</b> <ul style="list-style-type: none"> <li>● Based on Aerospike Database 5</li> <li>● Single system spans multiple sites</li> <li>● Active/active design</li> <li>● Automated failovers</li> <li>● Synchronous data replication</li> <li>● Strong, immediate data consistency</li> <li>● Cloud, on premises, or hybrid</li> </ul>
<b>Global business continuity:</b> Protect key apps	<b>Cross-Datacenter Replication</b> <ul style="list-style-type: none"> <li>● Based on Aerospike Database 5</li> <li>● Active/active or active/passive design</li> <li>● Asynchronous data replication</li> <li>● Granular replication levels</li> <li>● Smart data shipment algorithms</li> <li>● Flexible data consistency options</li> <li>● Cloud, on premise, or hybrid</li> </ul>
<b>Archive:</b> Protect least critical apps, provide long term retention, support compliance	<b>Aerospike Database 5</b> <ul style="list-style-type: none"> <li>● Optional backup to NFS</li> <li>● Support for low-cost data storage</li> <li>● Encryption, compression options</li> <li>● Filtering, throttling options for high performance</li> </ul>

**Table 2: Aerospike technologies for achieving various business objectives**

So how do these and other Aerospike technologies fare in production? Let's explore some customer use cases.

## Customer profiles

Aerospike is deployed globally by firms across industries, including finance, retail, technology, telecommunications, ad tech and others. Dozens of case studies and customer testimonials are available on [Aerospike's website](#). This section will highlight a few firms that were drawn to Aerospike in part by its exceptional data availability and overall reliability.

### AppNexus/Xandr

A technology company that powers the sale and purchase of digital advertising, AppNexus (acquired by AT&T in mid-2018 and renamed Xandr) processed more than 10 billion transactions daily in 2018 -- more than double the combined average daily transactions of the New York Stock Exchange, Nasdaq and Visa at that time. Availability is critical to its success, as the executive responsible for its global infrastructure and operations noted in an [Aerospike User Summit conference presentation](#):

*AppNexus does not close. Ever! AppNexus has no natural or scheduled maintenance windows where ad serving can be unavailable.*

So how do they fulfill such extreme demands? They rely on Aerospike. Indeed, AppNexus turned to Aerospike to manage its 100+ TB of real-time data even before XDR and other features discussed in this paper were generally available.

*Aerospike has been in a critical path at AppNexus for 7+ years [as of 2018]. Aerospike must be available every second of every day . . . We perform upgrades/restarts at all hours. Middle of the day? Yes, indeed! During peak request times? Absolutely, positively yes. Why? Because we can! The impact is very, very minimal . . . Use Aerospike. Do not be afraid.*

### PayPal

A worldwide leader in online money transfers, billing and payments, PayPal partnered with Aerospike initially in 2015 to help the firm rapidly detect attempted fraud and minimize monetary losses. To manage risk, PayPal demanded ultra-fast performance and high scalability from its data platform. Its availability requirements were aggressive: at least 99.99% uptime every hour of every day without exception.

PayPal found its legacy database struggling to manage the massive amounts of data that the firm was collecting. It needed a cost-effective strategy to scale horizontally without compromising performance or uptime. The firm forecast future needs to involve 100s of petabytes of data and a doubling of transaction volumes (from 3.5 to 7 million transactions a second).

PayPal found Aerospike uniquely qualified to power the global demands of its business. After replacing its legacy in-memory system with Aerospike, PayPal slashed the number of missed fraudulent transactions by 30x and projected a \$9 million savings in hardware costs due to Aerospike's significantly smaller server footprint. And it didn't have to compromise on speed or availability, either. Indeed, PayPal's positive experience with Aerospike prompted the firm to expand its deployment. In 2020, PayPal was running more than 2000 Aerospike servers. As its [lead data architect for risk and compliance](#) said,

*PayPal is innovating deep analytics to rapidly respond to emerging fraud patterns . . . to accelerate detection, reduce losses, and achieve near-continuous availability.*

## European financial institution

A major European financial institution relies on Aerospike's multi-site clustering technology to enable money transfers between member banks within seconds regardless of the time of day. Therefore, strong, immediate data consistency and high availability were critical database requirements for this payments application, known as the TARGET Instant Payment Settlement (TIPS) service.

The bank deployed a single Aerospike cluster across two data centers to track payment state, each with three nodes. This Aerospike infrastructure readily met the bank's target of processing 2000 transactions per second and up to 43 million transactions per day with round-the-clock availability. It also supported the bank's mandate that costs be within €0.0020 per payment. However, other solutions didn't meet the bank's objectives for resiliency (100% uptime), consistency (no data loss and no dirty or stale reads), and low transaction cost. As a result, the bank plans to add a third data center to the cluster to increase capacity and further enhance resiliency.

## Nielsen

A leading provider of data to digital marketers, Nielsen Cloud Marketing can't afford to be offline. Its global business demands quick access to data around the clock without service interruptions. As of 2019, Nielsen deployed Aerospike in data centers across the USA, Europe, and Asia/Pacific to keep its business running. As the VP of infrastructure noted in a [presentation at one of the Aerospike User Summits](#),

*Cross-Datacenter Replication (XDR) is vital to our organization. We need to have redundant data centers. We need our user objects to be available in multiple facilities. Aerospike does this with ease. It's able to take advantage of a 20-gigabit link without a problem and ship data continuously.*

The firm stores device information and event history in real time for billions of users in Aerospike. Low latency transactions on user objects are essential to Nielsen's competitive edge. The firm conducts real-time modeling and analysis over the data, returning the information to the user in

200 milliseconds or less. Reliable replication across multiple locations is vital. To date, the firm has effectively weathered a data center outage (caused by a hurricane impacting New York City), a complete data center hardware refresh, and the daily operational challenges of a 24x7 global business using Aerospike.

*Aerospike has been a great partner over the years. The ability to replicate data across regions is something that Aerospike provides that very (few) other NoSQL databases do with ease. Without Aerospike, we'd be looking for a new NoSQL data store that performs as well as Aerospike does, and I haven't seen one out there.*

## Verizon Media/Yahoo

Faced with a disparate, aging data management infrastructure, Verizon Media (spun-out of parent Verizon in September 2021 and rebranded Yahoo) sought to consolidate and streamline its approach to managing mission-critical real-time data. After extensive evaluations, the firm decided to base its global modernization effort on Aerospike. Applications targeted for the new Aerospike infrastructure require low data access latencies, petabyte-level scalability, ease of operations, strong data consistency, global deployments spanning multiple data centers, and a low server footprint.

Verizon Media is deploying Aerospike in several worldwide data centers using cross-data center replication (XDR) as needed. With clusters of 100 nodes or less, many based on PMem, Verizon Media expects to achieve the scalability and performance it needs for 100+ applications. By replacing multiple open-source platforms (RocksDB, HBase) and some homegrown technology with Aerospike, Verizon Media expects to simplify its infrastructure, support emerging business requirements, prepare for future growth, and keep within its target budget. As one architect put it,

*"Aerospike is the only product to beat our own technology in 12 years, and we've tested everything."*

For details, see [Verizon Media's presentation](#) at the Aerospike Digital Summit 2021.

## Summary

While few firms would deny their need for highly available operational data, 99.99% uptime (or more) is really only part of the requirement. Affordability, scalability, low operational overhead, flexible deployment options, and ultra-fast, predictable performance are vital to real-time operational systems. And that's where the challenge arises.

Unlike other platforms, Aerospike doesn't force its customers to compromise. To achieve such goals, you don't need massive server footprints, hard-to-find IT specialists, or a complex collection of open source and commercial technologies. You just need Aerospike.

Sound hard to believe? Leading firms in finance, retail, telecommunications, technology, and other industries have turned to Aerospike to power their "always-on" business models, often after other

systems failed to live up to expectations. As this paper discussed, some users have been running Aerospike for years without interruption or downtime. And many have saved \$1 - \$10 million per application when compared with other approaches.

Aerospike provides 99.999% uptime (5 nines) during common failure scenarios and system upgrades and is available for cloud, on-premises, or hybrid deployments. Synchronous and asynchronous data replication technologies offer global firms business continuity across multiple locations. The platform's self-managing and self-healing nature automates routine maintenance tasks and handles many failover operations without any manual intervention. Furthermore, Aerospike scales readily from 100s of gigabytes to petabytes while maintaining relatively small server footprints and delivering predictable, ultra-fast data access speeds even during periods of peak demand.

Perhaps you're skeptical that any one platform could do all that. Why not find out for yourself? Visit [Aerospike's](https://aerospike.com) website to learn more, or contact [sales@aerospike.com](mailto:sales@aerospike.com) to arrange a briefing.

## About Aerospike

The Aerospike Real-time Data Platform enables organizations to act instantly across billions of transactions while reducing server footprint up to 80%. The Aerospike multi-cloud 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>.

