

WHITE PAPER

5 signs you've outgrown DynamoDB



Snapshot

Why replace DynamoDB?



You're worried about TCO



You want better efficiency and price performance



You need greater flexibility



You want to simplify your data management infrastructure



You don't want to be locked into AWS

Introduction

Many organizations find that DynamoDB offers an easy onramp to NoSQL data management. Available only on AWS, DynamoDB readily supports key-value and JSON data, scales to support large data volumes, and integrates well with other AWS services. But, initial enthusiasm for DynamoDB can fade as business demands evolve. Concerns about TCO (total cost of ownership), price/performance at scale, and deployment limitations often emerge.

What are 5 signs that your organization may have outgrown DynamoDB?

1

You're worried about TCO

- Are you concerned about cost overruns as your business needs grow?
- Is it difficult to understand DynamoDB's pricing models and accurately predict expenses?

2

You want better efficiency and price performance

- Do you wish you could achieve ultra-fast data access speeds on a smaller server footprint?
- Would you like to reduce data storage requirements without compromising data availability?

3

You need greater flexibility

- Do you want a database that supports a broader range of data models?
- Would you prefer more choice in how your data is managed and stored to meet your price/ performance goals?
- Can you benefit from global (multi-site) transaction processing with strong, immediate consistency?
- Do you want the option to scale both vertically and horizontally, depending on your needs?

4

You want to simplify your data management infrastructure

- Would you prefer to use one data management system for a wider variety of applications, especially those that rely on multiple data models?
- Would you like to use the same system as a cache or persistent store?
- Do you need a cache that can front-end various legacy and non-AWS data stores?

5

You don't want to be locked into AWS

- Do you want more choices for supported cloud platforms?
- Would you like the option to deploy your data management system on-premises, in a cloud environment, or in a hybrid environment?

1. You're worried about TCO

Maintaining low total cost of ownership (TCO) for a data management system can be tough, particularly as data volumes and transaction rates grow. Pricing structures and technical features often play a big role in helping you accurately forecast operational expenses as your business needs evolve.

Aerospike and DynamoDB employ different pricing models, and many firms find that Aerospike's approach makes it easier to predict costs. By contrast, [DynamoDB's pricing](#) is more complex, requiring detailed knowledge of expected data access patterns and throughput to optimize costs. Users are charged for reading, writing, and storing data. [Optional services](#) (e.g., backup/restore, data import/export, and others) incur additional fees. For basic services, DynamoDB offers two different pricing models to appeal to users with predictable or unpredictable throughput rates.

With on-demand capacity mode, charges are based on read and write request units (RRUs and WRUs, respectively) consumed by the user workload. AWS automatically ramps up or down as needed, and users do not specify throughput rates. With provisioned capacity, users specify the number of RRUs and WRUs per second they expect to need and are charged accordingly. Auto-scaling can be used in provisioned mode to address performance concerns. With either mode, a single read or write command may require multiple RRUs or WRUs, depending on item size, data consistency level, transactional aspects, and other factors. For the lowest data access latencies, firms often employ the [DynamoDB Accelerator \(DAX\)](#) caching service, which incurs additional fees. The bottom line for DynamoDB users is that a thorough understanding of which operations incur additional "request units" and which optional services are needed to support business requirements is essential for expense planning.

Aerospike, on the other hand, charges for data volumes under management. There are no separate costs for data access operations, so no additional expenses are incurred as more applications rely on Aerospike-managed data to satisfy business needs. Aerospike does not charge additional fees for backup/restore, non-production (e.g., dev/test), or caching configurations – tasks that can incur additional fees for DynamoDB users. Low TCO at scale is one of the most common reasons customers choose Aerospike.

2. You want better efficiency and price performance

[Cost-efficient](#) use of hardware resources enables your data management platform to deliver the performance you need at an attractive price point. While both Aerospike and DynamoDB offer high scalability, only Aerospike enjoys a decade-long reputation for delivering ultra-fast, predictable data access speeds on relatively small server footprints.

Among the design points of Aerospike that contribute to its speed and cost efficiency is its [Smart Client layer](#). Aerospike maintains a client-side dynamic partition map that identifies the primary node for each partition, enabling the client layer to route read or write requests directly to the correct nodes without any additional network hops. This reduces data access latency and leads to more predictable data access speeds.

Meanwhile, DynamoDB maintains multiple [request router](#) instances that perform user authentication and authorization as well as consult a partition metadata component to identify all partition replicas. Depending on the nature of the request (read or write) and the desired data consistency level, the router will forward the request to the replica, serving as the partition leader or to a follower.

In addition, Aerospike's self-managing features and exploitation of modern hardware technologies often enable it to deliver comparable or better performance on clusters with fewer nodes, thus improving price performance. Aerospike exploits SSDs as raw block devices to achieve maximal resource efficiency.

3. You need greater flexibility

As customer demands and business conditions change over time, having a flexible data management platform becomes critical. Data modeling and storage/memory management options are two factors worth exploring, as both influence the types of applications that can be supported with cost efficiency and high performance.

Aerospike and DynamoDB both support key-value and document data models. Aerospike also offers high-performance support for graph and vector data, whereas AWS recommends separate products like [Neptune](#) to handle graph data and [OpenSearch Service](#) for vector storage and search. Aerospike also provides more flexible and robust operations on complex data types, such as Maps and Lists, allowing users to apply filtering, sorting, and range queries directly on these structures.

While both Aerospike and DynamoDB are designed for distributed computing environments, Aerospike offers firms more options to tailor memory and storage configurations for their application and budget needs. With Aerospike, firms can choose from hybrid memory (SSDs and DRAM), all DRAM, or all SSD configurations. Aerospike optimizes the use of SSDs as primary data storage. It treats them as raw devices, employing a proprietary log-structured file system rather than relying on file system, block, and page cache layers for I/O. This provides distinct performance and reliability advantages.

DynamoDB distributes and stores data using a [B-tree structure](#), while SSDs are used to persist these structures. It does not permit users to tailor memory and storage usage in the same manner as Aerospike.

DynamoDB and Aerospike also differ in their abilities to operate across data centers. In short, Aerospike data replication and multi-site support offer users more options while maintaining operational simplicity. Aerospike supports both synchronous and asynchronous data replication across data centers, availability zones, and even cloud providers, so firms can customize deployments to meet their resilience and availability needs. Both synchronous and asynchronous data replication are supported via [multi-site clustering](#) and [Cross Datacenter Replication](#), respectively. Async configurations offer selective replication services (i.e., data filtering) for data sovereignty compliance and limiting data transfers to minimize network, bandwidth, and storage costs. Aerospike's synchronous replication automatically enforces strong data consistency.

In contrast, DynamoDB users who want to maintain copies of tables (or subsets of tables) [across different DynamoDB database instances](#) can implement their own solutions using other AWS services (e.g., Streams and Lambda). DynamoDB also allows users to create global tables with replicas stored in different regions. Global tables allow read/write access to occur in any region, and [changes are automatically replicated across regions asynchronously, typically within a second](#). Reads spanning multiple regions of global tables operate under eventual consistency. Write conflicts across regions are resolved with a last-writer-wins strategy. [Transactions are not supported across regions in global tables](#). Note: as of Q4'24, global tables will soon have [strong consistency](#) as well.

Firms can scale up, scale out, or shrink Aerospike clusters as desired with no cluster downtime and no need for auxiliary data services. To scale an Aerospike cluster vertically, firms simply configure nodes with upgraded CPUs, storage, etc. Similarly, firms can add or remove nodes when needed for horizontal scaling, and Aerospike will automatically rebalance data partitions across the cluster with minimal data movement and minimal performance impact.

With DynamoDB, horizontal scaling is achieved with a separate AWS [Application Auto Scaling service](#); users set policies on DynamoDB tables and global secondary indexes, and the service automatically adjusts provisioning based on consumption. There is complexity and risk in setting scaling policies. Vertical scaling is not within user control.

4. You want to simplify your data management infrastructure

Simplifying your IT infrastructure frees your staff to focus on other projects, including those designed to drive new business opportunities or cut costs. Aerospike's flexible deployment options enable firms to standardize on its platform for a wide range of applications, reducing the overall complexity of their data management infrastructures.

Many firms initially deploy Aerospike as a high-speed cache to deliver fast access to data managed by other systems that cannot keep pace with business demands. When configured to manage user data and indexes entirely in memory, Aerospike provides a cost-effective means of enabling new applications to access critical data managed by SQL and NoSQL stores, including those not deployed on AWS. Firms can also leverage Aerospike's built-in persistence features to manage real-time data without relying on a back-end system for storage.

Conversely, DynamoDB is widely deployed as a persistent, distributed store on AWS. An additional specialized caching service (DAX) is available to front-end DynamoDB only if needed.

Aerospike's performance-optimized support for key-value, document, graph, and vector data makes it suitable for a wide range of applications. AWS promotes DynamoDB for key-value and document data but recommends different services for graph and vector data.

5. You don't want to be locked into AWS

Although AWS is one of the leading cloud computing providers, many firms simply aren't comfortable relying too heavily on any one vendor for their essential service needs. Simply put, they want the option to move between different cloud providers as well as hybrid configurations. With DynamoDB, that's not an option – it only runs as a fully automated Database-as-a-Service (DBaaS) on AWS. With Aerospike, you have options beyond AWS.

Aerospike can be deployed:

- on-premises
- as a cloud-managed service on AWS or Google Cloud Platform (GCP);
- in hybrid configurations
- with the Aerospike Kubernetes Operator (AKO), which supports Amazon Elastic Kubernetes Service (EKS), Microsoft Azure Kubernetes Service (AKS), Google Kubernetes Service (GKS), and local deployments

Aerospike's flexible platform options give firms greater flexibility to tailor deployments for different applications and to implement fail-over strategies spanning multiple cloud platforms or cloud and on-premises configurations.

Get ready for the future

A real-time, multi-model data management system, Aerospike delivers ultra-fast, predictable performance for read/write workloads at gigabyte to petabyte scale with low TCO. Firms in financial services, AdTech, telecommunications, e-commerce, gaming, and other industries use Aerospike for systems of engagement and systems of record, simplifying their infrastructures and often reducing costs considerably.

Compared to DynamoDB, Aerospike supports a wider range of deployment platforms, configurations, and data models and more flexible data replication. Aerospike's pricing structure, cost-efficient use of hardware resources, and internal architecture enable firms to enjoy exceptional price/performance at scale by delivering predictable, ultra-low data access latencies on comparatively small server footprints. It can serve as a cache for various SQL and NoSQL systems to speed access to critical data and as a real-time operational data store for key-value, document, vector, and graph data to support many business applications.

If you've experienced any of the five signs just discussed, why not explore what Aerospike can do for you? [Contact Aerospike](#) to estimate TCO savings for your workload, or [try Aerospike](#) and see how you can benefit, too.



Aerospike is the real-time database for mission-critical use cases and workloads, including machine learning, generative, and agentic AI. Aerospike powers millions of transactions per second with millisecond latency, at a fraction of the cost of other databases.

Global leaders, including Adobe, Airtel, Barclays, Criteo, DBS Bank, Experian, Grab, HDFC Bank, PayPal, Sony Interactive Entertainment, The Trade Desk and Wayfair, rely on Aerospike for customer 360, fraud detection, real-time bidding, profile stores, recommendation engines, and other use cases.

Try Aerospike for free: aerospike.com/try-now

©2024 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 | aerospike.com