

Five Reasons DynamoDB Users Should Consider ScyllaDB

When is DynamoDB or ScyllaDB the right choice for an AWS User?

ScyllaDB provides a close-to-the-metal architecture for linear scalability on cloud infrastructure, achieving millions of operations per second (OPS) with single-digit millisecond P99 latencies. Its fully managed version, [ScyllaDB X Cloud](#), offers extreme elasticity, allowing auto-scaling from thousands to 2 million OPS in minutes. In contrast, DynamoDB features a serverless architecture with costly add-ons for low latency and global replication (not covered by reserve provisioning), leading to higher expenses and administrative complexity.

The ScyllaDB Price Performance Drivers

The difference in architecture and approach to commercialization between ScyllaDB and DynamoDB should compel DynamoDB users with one or more of the following deployment characteristics to consider moving to ScyllaDB:

- Running workloads with combined read and write throughput of over 50 K OPS, particularly with write-heavy workloads
- Running with very tight latency SLAs, particularly with single-digit millisecond P99 latencies
- Running across multiple AWS regions, particularly with multiple regions on multiple continents
- A requirement or need for an option to run DynamoDB apps on other public clouds or on-premises.

Any combination of these characteristics will generally lead to ScyllaDB delivering the same or better throughput and latency at 50% of the cost of DynamoDB. Of course, the specific workload, operational scenarios, and other factors must be evaluated case-by-case to nail down the exact cost comparison figures. Still, differences in fundamental features, functionality, and options to reserve provisioned capacity upfront between ScyllaDB and DynamoDB underpin this cost discrepancy range and act as price-performance drivers, as described in the following table.

Price-Performance Driver	ScyllaDB	DynamoDB
1. Capacity pricing model (ScyllaDB delivers optimized and transparent use of cloud infrastructure, linear scalability, flexibility for modest data growth, and built-in burst capacity)	<ul style="list-style-type: none"> • Server-based – user can optimize underlying cloud resources • Provisioned(w/ Flex Credits) or OnDemand • Performance capacity planning • Cost estimated only by total vCPUs, NVMe storage, and data egress charges (reduced through network compression) 	<ul style="list-style-type: none"> • Serverless – underlying resource use and efficiency hidden from the user • Provisioned or OnDemand (3.48X markup) • Read and Write Capacity Units (RCU/WCU) planning • Cost estimated by total RCUs + WCUs (5X RCU), multiplied by WCU/RCU size in kb / 1kb payload
2. Caching model (ScyllaDB delivers ultra-low single-digit to sub-msec P99 latencies without an external cache – even for throughputs > 1M OPS)	<ul style="list-style-type: none"> • Integrated in-memory cache • No separate API or developer/DBA intervention or external cache is required • Replaces external caches in most use cases 	<ul style="list-style-type: none"> • External cache (separate charge for either AWS DAX or AWS ElastiCache (Redis)) • External cache recommended to reach P99 latencies below 25 msec (multi-region) • Note: Reserved provisioning is unavailable for DAX, making this a higher/unpredictable cost
3. Global replication model (Global replication includes synchronization of global tables without additional charges or external components)	<ul style="list-style-type: none"> • No separate charges for global tables, each individually configurable for global or local quorum • The higher the replication factor, the higher the storage and data egress fees 	<ul style="list-style-type: none"> • Added charge per Global table per region per WCU • Note: Reserved provisioning is unavailable for Global Tables, , making this a higher/unpredictable cost
4. Multi-table support (Significant savings by bundling small workloads/tables)	<ul style="list-style-type: none"> • Multiple tables can be supported • Workload prioritization by table within the provisioned capacity 	<ul style="list-style-type: none"> • Workload is always defined as a single table • Charge per workload using RCUs, WCUs, Global Tables, and DAX
5. Deployment model (Provides cloud flexibility and breaks cloud vendor lock-in)	<ul style="list-style-type: none"> • Runs on AWS, GCP, Azure, OCI, On-prem. • DynamoDB-API and CQL compatible 	<ul style="list-style-type: none"> • Runs only on AWS • Uses only the DynamoDB API

ScyllaDB Price-Performance Advantages Corroboration

Third-Party Benchmarks

ScyllaDB sponsored a benchmark by BenchANT, specialists in NoSQL database testing. They evaluated three distributions—Uniform, Zipfian, and Hotspot—across various throughput scenarios with both ScyllaDB Cloud and DynamoDB on AWS. The detailed report is available [here](#). Figure 2 shows the best results for DynamoDB compared to ScyllaDB for provisioned workloads with a uniform distribution, though real-world scenarios are typically better modeled as Zipfian or Hotspot.

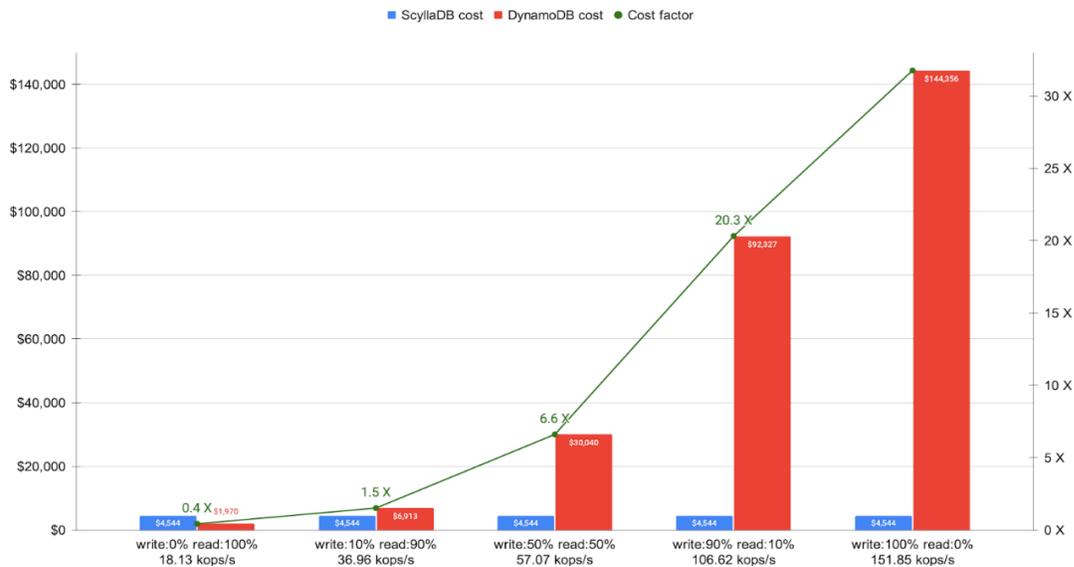


Figure 1 - Price comparison for uniform distribution - Provisioned for ScyllaDB and DynamoDB

At around 50K OPS, ScyllaDB costs half as much as DynamoDB using provisioned list pricing. ScyllaDB's pricing model is based on capacity, so prices remain constant despite increased throughput. Above 50K OPS, the cost of DynamoDB compared to ScyllaDB ranges from 2X to 40X for any range of read/write mixes, dataset size, and replication factor.

Customer Migrations and Improved Outcomes

Many users have shifted from DynamoDB to ScyllaDB in the past two years. Most continue to use ScyllaDB Cloud on AWS, while others run DynamoDB applications on ScyllaDB Cloud for Google or deploy ScyllaDB Enterprise on-premises or in private clouds. The table below highlights common scenarios for former DynamoDB users now using ScyllaDB.

Company (Case study link)	Business	Rationale for moving away from DynamoDB	Reasons for moving to ScyllaDB	Improved Outcomes After Migration
	Brazilian/S. American Food Delivery App	Rapid scaleup over 2 years on DynamoDB was too costly	Far less expensive yet scaled better with lower latency	Reduced monthly subscription cost by 9X, & P99 latency from 80 to 3 msec
	Social Media Digital Native Publisher	Cost of reducing latency through front-end Redis cache, bill shock as their data volume grew 2X per year	Retain DynamoDB API and reduce latency in AWS/GCP multi-cloud environments	R/W to 1 M OPS/100K OPS across AWS and GCP without a front-end cache at under 50% the cost of DynamoDB
	Adtech Services AI-based auctions	Zero downtime migration from DynamoDB to ScyllaDB without changes to the code	Reduced costs by 50% GCP integration without risk	Avoided cost trajectory that would be impossible to absorb as they scale up
	Indian Media Conglomerate - OTT platform	Difficulty running in multiple regions, Bill shock as DynamoDB workloads grew	Predictable monthly costs for global scale capacity to scale	Reduced costs by 5X for global operations with P99 latencies of 4 msec (read) and 0.8 msec (write)

Learn More

Start scaling with the world's most affordable and scalable high-performance NoSQL database. Embark on the path to ScyllaDB expertise for free through [ScyllaDB University](#) and [get started with ScyllaDB Cloud](#), our NoSQL DBaaS on AWS and Google Cloud. Want to better understand your DynamoDB spend? Use our [DynamoDB cost calculator](#) to analyze real-world workloads including, uneven traffic patterns, traffic surges, caching strategies, and Global tables. Interested in a guided, detailed comparison of your DynamoDB spend and ScyllaDB savings? [Contact us](#) for a technical evaluation.