



WHITE PAPER

Total Cost Comparison: Redis Enterprise vs AWS DynamoDB

CONTENTS

Executive Summary	2
Introduction to Redis Enterprise:	2
Scenarios:	2
Scenario 1: 200 GB dataset size, 5k ops/sec throughput	3
Scenario 2: 1000GB dataset size, 30k ops/sec throughput	4
Scenario 3: 3000 GB dataset size, 100k ops/sec throughput	6
Should Cost Be Your Only Consideration?	8
A Note on Scenarios:	9
Customer Example: Whitepages	9
Conclusion	9

Executive Summary

Redis Enterprise is available in every cloud, inside VPCs, and as downloadable software — as close to your application as you need it. By reviewing a number of typical cloud workload scenarios, this whitepaper shows how Redis Enterprise delivers significant cost savings over AWS DynamoDB, particularly when using Redis on Flash technology. Cost calculations leverage source material from the [Redis Enterprise calculator](#) and the [AWS pricing calculator](#).

Introduction to Redis Enterprise:

Redis Enterprise is an enterprise-grade database offering from Redis Labs. It is available as downloadable software or database-as-a-service in major public clouds, as well as inside the VPC environments of major clouds. This whitepaper primarily focuses on cost comparisons for our database-as-a-service offering inside your VPCs, using Redis on Flash databases.

Redis is fundamentally richer in developer features than DynamoDB — with a range of data structures (Hash, Sets, Sorted Sets, Lists, Hyperloglog, Bitfields, Bitmaps, Streams) in addition to key-value. The database also offers modules such as RediSearch, ReJSON, Redis Graph, Redis-ML, ReBloom filters and others, extending Redis to data processing scenarios that would be unviable with DynamoDB.

These **built-in modules** extend Redis to almost every data processing scenario under the sun, at blazing fast performance. For example, built-in search and secondary indexing with RediSearch easily **outperforms standard search engines by a factor of 5X**.

Redis Enterprise's distinguishing features include linear scalability with **exceptional performance**. Its very high throughput (ops/sec) at extremely low latencies with minimal resources make it a cost-effective choice for a range of scenarios. Additional cost savings can be achieved using a combination of RAM and Flash SSDs to store, process and serve your dataset with Redis on Flash. Redis on Flash uses intelligent tiering, with keys and hot values stored in RAM, while cold values get automatically pushed to Flash— without requiring any code changes to your application.

Additionally, Redis Enterprise distributes datasets through an **active-active** mechanism that relies on the cutting-edge **CRDT** (conflict-free replicated data types) approach. This not only delivers local latencies for read/write applications, regardless of their number of geo-replicated regions, but also provides seamless, deterministic conflict resolution for each data type. The mechanism does not rely on quorum, ensuring uninterrupted availability, even when a majority of geo-replicated regions are not available.

Redis Enterprise's enhanced data **persistence** mechanisms overcome the potential disk access bottlenecks associated with multiple Redis instances trying to access a persistent storage device, or during AOF rewrite and snapshot operations. **High availability** is assured through in-memory diskless replication across zones/racks, as well as between clusters located in different regions or clouds, or between on-prem and cloud deployments. Automatic failover is triggered instantaneously and recovery takes place within single digit seconds, **without data loss**.

Scenarios:

Since the data needs of different workloads vary considerably, we used sample workload types to evaluate cost differences. Many successful applications in the entertainment, e-commerce, advertising, travel, social and transportation categories fall into the following scenarios, and their requirements are pretty typical.

We evaluated three different scenarios from our customers, each with a 1:1 mix of reads and writes, and an object size of 1KB:

- Scenario 1: 200 GB dataset size, requiring 5k ops/sec throughput
- Scenario 2: 1000 GB dataset, 30k ops/sec
- Scenario 3: 3000 GB dataset, 100k ops/sec

In all cases, we used replication and Multi-AZ in our calculations, since serious workloads and applications need these availability capabilities.

Scenario 1: 200 GB dataset size, 5k ops/sec throughput

A typical starting scenario is a dataset 100 GB to 200 GB in size, with a 5k ops/second throughput requirement and latency preferably under 1ms. One important thing to keep in mind is that with Redis Enterprise, even if most of the dataset is stored on Flash, you typically see <1ms latency. However, DynamoDB latency is usually in the **10-100 ms range**.

Because DynamoDB indexes every write by default, for our calculations 5k ops/sec with 50% reads and 50% writes translated to 2.5k reads/sec and 5k writes/sec.

**All prices below are computed for AWS US East (Virginia) and with HA & Multi-AZ*

		On Demand	Reserved
DynamoDB provisioned writes	5,000 writes/sec	\$2,360/mo	\$1,092/mo
DynamoDB provisioned reads	2,500 reads/sec	\$235/mo	\$108/mo
DynamoDB dataset storage	300 GB	\$69/mo	\$69/mo
DynamoDB index storage	30 GB	\$8/mo	\$8/mo
DynamoDB data transfer in	6,570 GB		
DynamoDB data transfer out	6,570 GB	\$591/mo	\$591/mo
DynamoDB Price		\$4.47/hr \$3,263/mo \$39,152/yr	1,868/mo \$22,413/yr

With Redis Enterprise VPC, we chose Flash databases, which are cost-optimal for this scenario. In addition to the cost of Redis Enterprise, we included the minimum cost of the underlying infrastructure (our technology has evolved over the years to maximize performance with the least possible computing resources). Note that Redis on Flash databases still have a RAM component, which ensures <1ms latency. Furthermore, Redis on Flash databases also need additional persistent storage, which is factored into the costs below. Even with all these components, Redis Enterprise VPC is over 11-35 % less expensive than DynamoDB.

Redis Labs Price Estimate

		On Demand Price	Reserved Price
Size	200 GB		
Throughput	5,000 ops/sec		
Total Redis on Flash shards	4 shards		
Redis Labs Price		\$1.23/hr \$900/mo \$10,800/yr	\$810/mo \$9,720/yr

Cloud Resources

		On Demand Price	Reserved Price
i3.xlarge compute instance	2	\$911/mo	\$580/mo
m4 large compute instance	1	\$73/mo	\$42/mo
Storage	2,215 GB	\$222/mo	\$222/mo
Cloud Provider Price		\$1.65/hr \$1,206/mo \$14,466/yr	\$844/mo \$10,129/yr
Total Price		\$2.88/hr \$2,106/mo \$25,266/yr	\$1,654/mo \$19,849/yr

You can see the summary comparison in the following table. We show monthly costs, because workload needs might vary cyclically and your costs can easily adapt to this with both Redis Enterprise and DynamoDB.

Monthly Price Comparison	Redis Enterprise VPC + Cost of AWS instances	AWS DynamoDB	Redis Enterprise VPC + Cost of AWS Instances	AWS DynamoDB	On Demand Savings with Redis Enterprise	Reserved Savings with Redis Enterprise
	<i>On Demand</i>	<i>On Demand</i>	<i>Reserved</i>	<i>Reserved</i>		
Database type	Flash		Flash			
	\$2,106	\$3,263	\$1,654	\$1,868	\$1,157	\$214
					35%	11%

Scenario 2: 1000GB dataset size, 30k ops/sec throughput

Another familiar scenario occurs after your application is moderately successful, and has acquired a large number of users who interact with it often during the day. In these cases, DynamoDB becomes an even more expensive choice.

*All prices below are computed for AWS US East (Virginia) and with HA & Multi-AZ

		On Demand	Reserved
DynamoDB provisioned writes	30,000 writes/sec	\$14,223/mo	\$6,553/mo
DynamoDB provisioned reads	15,000 reads/sec	\$1,421/mo	\$649/mo
DynamoDB dataset storage	1,500 GB	\$369/mo	\$369/mo
DynamoDB index storage	150GB	\$38/mo	\$38/mo
DynamoDB data transfer in	39,420 GB		
DynamoDB data transfer out	39,420 GB	\$3,402/mo	\$3,402/mo
DynamoDB Price		\$26.65/hr \$19,452/mo \$233,425	\$11,010/mo \$132,120/yr

By comparison, Redis Enterprise VPC is up to 65% cheaper with the Flash option, not to mention its much lower latencies:

Redis Labs Price Estimate

		On Demand Price	Reserved Price
Size	1,000 GB		
Throughput	30,000 ops/sec		
Total Redis on Flash shards	20 shards		
Redis Labs Price		\$6.16/hr \$4,500/mo \$54,000/yr	\$4,050/mo \$48,600/yr

Cloud Resources

		On Demand Price	Reserved Price
i3.2xlarge compute instance	4	\$1,822/mo	\$1,161/mo
m4 large compute instance	1	\$73/mo	\$42/mo
Storage	4,415 GB	\$442/mo	\$442/mo
Cloud Provider Price		\$3.20/hr \$2,337/mo \$28,039/yr	\$1,644/mo \$19,733/yr
Total Price		\$9.37/hr \$6,837/mo \$82,039/yr	\$5,694/mo \$68,333/yr

When you compare the two, there's a 48-65% cost savings with Redis Enterprise:

Monthly Price Comparison	Redis Enterprise VPC + Cost of AWS instances	AWS DynamoDB	Redis Enterprise VPC + Cost of AWS instance	AWS DynamoDB	On Demand Savings with Redis Enterprise	Reserved Savings with Redis Enterprise
	<i>On Demand</i>	<i>On Demand</i>	<i>Reserved</i>	<i>Reserved</i>		
Database type	Flash		Flash			
	\$6,837	\$19,452	\$5,694	\$11,010	\$12,615	\$5,316
					65%	48%

Scenario 3: 3000 GB dataset size, 100k ops/sec throughput

Once businesses attract a steady set of users interacting with their applications, they accumulate a ton of data that needs to yield real time insights or trigger actions in the application. Implementing such a scenario in DynamoDB is very expensive, latencies are well over 10-100ms, and only the simplest data processing scenarios can be implemented.

**All prices below are computed for AWS US East (Virginia) and with HA & Multi-AZ*

		On Demand	Reserved
DynamoDB provisioned writes	100,000 writes/sec	\$47,438/mo	\$21,844/mo
DynamoDB provisioned reads	50,000 reads/sec	\$4,743/mo	\$2,163/mo
DynamoDB dataset storage	4,500 GB	\$1,119/mo	\$1,119/mo
DynamoDB index storage	450 GB	\$113/mo	\$113/mo
DynamoDB data transfer in	131,400 GB		
DynamoDB data transfer out	131,400 GB	\$10,017/mo	\$10,017/mo
DynamoDB Price		\$86.89/hr \$63,429/mo \$761,146/yr	\$35,255/mo \$423,058/yr

Consider instead Redis Enterprise:

Redis Labs Price Estimate

		On Demand Price	Reserved Price
Size	3,000 GB		
Throughput	100,000 ops/sec		
Total Redis on Flash shards	60 shards		
Redis Labs Price		\$18.49/hr \$13,500/mo \$162,000/yr	\$12,150/mo \$145,800/yr

Cloud Resources

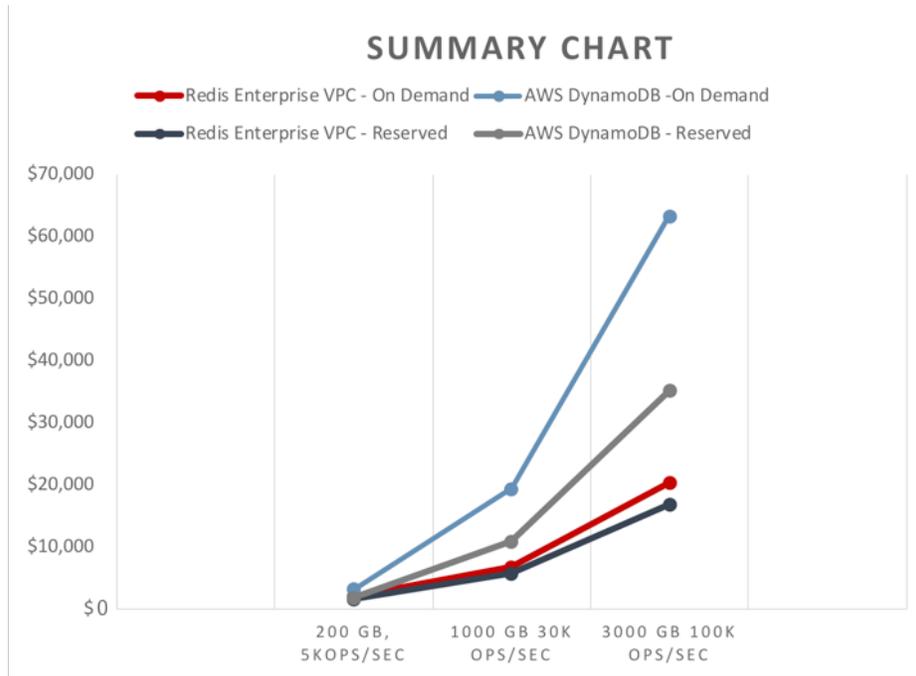
		On Demand Price	Reserved Price
i3.2xlarge compute instance	12	\$5,466/mo	\$3,482/mo
m4 large compute instance	1	\$73/mo	\$42/mo
Storage	13,215 GB	\$1,322/mo	\$1,322/mo
Cloud Provider Price		\$9.40/hr \$6,861/mo \$82,329/yr	\$4,846/mo \$58,149/yr
Total Price		\$27.89/hr \$20,361/mo \$244,329/yr	\$16,996/mo \$203,949/yr

Here again, the total savings from Redis Enterprise are in the 52- 68% range, and customers gain the advantage of sub-milli-second latencies and the ability to use Redis' versatile data structures and modules for a range of processing scenarios.

Monthly Price Comparison	Redis Enterprise VPC + Cost of AWS instances	AWS DynamoDB	Redis Enterprise VPC + Cost of AWS instance	AWS DynamoDB	On Demand Savings with Redis Enterprise	Reserved Savings with Redis Enterprise
	<i>On Demand</i>	<i>On Demand</i>	<i>Reserved</i>	<i>Reserved</i>		
Database type	FLASH		FLASH			
	\$20,361	\$63,429	\$16,996	\$35,255	\$43,068	\$18,259
					68%	52%

In Summary

The below charts visually summarize the three scenarios and the savings in each:



Data Characteristics	Redis Enterprise VPC - On Demand	AWS DynamoDB - On Demand	Redis Enterprise - VPC - Reserved	AWS DynamoDB - Reserved
200 GB, 5kops/sec	\$2,106	\$3,263	\$1,654	\$1,868
1000 GB, 30k ops/sec	\$6,837	\$19,452	\$5,694	\$11,010
3000 GB, 100k ops/sec	\$20,361	\$63,429	\$16,996	\$35,255
Savings (avg)	66%		49%	

Should Cost Be Your Only Consideration?

The quick answer is no. You must consider other factors when selecting a database, such as growth vectors for your business, portability of your data to other environments, staying ahead of your competition and reducing your dependence on single large vendors. You should think about whether your applications are likely to generate and process increasing quantities of data as they acquire more users. Is the profile of your data requirements likely to change? If so, be sure to model future scenarios, as well.

Among portability requirements, think of scenarios where you might need to change your operating environments. For example, if you decide to move your applications between clouds or even on-premises for greater control, security and compliance. You should be able to move the data your applications are using easily, without translation or heavy lifting. A third factor to consider is your competition. What types of technologies and cost profiles will help you stay ahead? Burgeoning costs from a single large vendor might hamper your ability to move in the agile fashion needed to grow.

In all such cases, you will find that choosing Redis Enterprise protects your business. It allows you to set up hybrid clusters that span clouds and on-premises environments, thus avoiding being tied down to any one provider or location. Replicating to different regions, clouds or on-premises is also relatively simple, since Redis Enterprise reduces your worries about data portability between environments.

The **Active-Active geographic distribution** implemented in Redis Enterprise simplifies the development and deployment of

highly performant, highly resilient applications. Redis Enterprise's usage of CRDT's cutting edge conflict resolution technology ensures you'll stay ahead of the game without losing any performance advantages. While DynamoDB is also available in Active-Active mode, it doesn't use CRDTs, so conflict handling needs to be done at the application layer. In addition, Redis Enterprise modules offer the flexibility of a range of processing scenarios without needing specialty databases for each type of operation. In particular, **RediSearch** offers simultaneous indexing and search for real time scenarios with **unparalleled performance** compared to leading industry search engines.

A Note on Scenarios:

The scenarios above are representative of what we see among our 8500+ enterprise customers. Several of them in industries like marketing, advertising, media, entertainment, travel and transportation have many more intense scenarios. In general, the more intense the scenario, the bigger the savings. Redis Enterprise natively **outperforms other NoSQLs** by several times, and the enhancements made to run it with multiple cores, multi-tenancy and Flash SSDs allow that performance to be achievable cost effectively.

Customer Example: Whitepages

Customer experience is a top priority for Whitepages, and its innovative people search tools are powered by more than five billion (and growing) global identity records. On average, the proprietary Identity Graph™ solution makes 150,000 to 200,000 calls per second to a three-terabyte database, even surpassing this number during peak hours.

As Whitepages expands its identity dataset to achieve global reach, it relies on Redis Enterprise to handle time-sensitive customer transactions and maintain a sub-hundred millisecond latency for its application. As a result, the company provides a consistent end-user experience through its digital identity verification services for businesses and consumers.

Initially, Whitepages housed the IDs to all that data in a key-value store on Amazon ElastiCache.

But as the company expanded into other countries and saw its database balloon rapidly, this solution became unwieldy and resource-intensive. Other NoSQLs were all evaluated as replacements for ElastiCache, but, ultimately, none could handle the full dataset size and still provide the single-digit latency Whitepages was seeking.

"Our Identity Graph product handles large amounts of data, and because of its size, our applications are extremely latency-sensitive," said Heather Wade, VP of engineering with Whitepages. "Redis Enterprise provided the single-digit latency we required and allows us to meet our customer SLAs every time. At Whitepages, customer experience is our priority, and we are able to deliver top-of-the-line services without compromising performance to all our customers."

Conclusion:

AWS offers a range of data services, and DynamoDB is one of many that seems an attractive proposition at first glance because of the integrated experience, but can turn into golden handcuffs for IT professionals. Not only does it cost more, it restricts customer portability and flexibility. IT professionals interested in growing their business as efficiently as possible, while keeping strategic options open, should embrace Redis Enterprise's flexibility to reduce costs and deliver multi-cloud and on-premises portability.



700 E El Camino Real, Suite 250
Mountain View, CA 94040
(415) 930-9666
redislabs.com