

Aerospike vs. ScyllaDB - A Brief Comparison

Aerospike's NoSQL database platform has earned an exemplary reputation for delivering exceptional -- and predictable -- performance at scale, superior uptime, and strong data consistency at a much lower total cost of ownership (TCO) than other alternatives. We benchmarked our platform against ScyllaDB, an Apache Cassandra replacement said to deliver "ultra-low latency and extreme throughput." Our tests revealed that Aerospike performed 7x better at 15% of the cost.

Benchmark Summary

We tested Aerospike 4.7.0.5 Community Edition and ScyllaDB Open Source version 3.1.2 using the Yahoo Cloud Serving Benchmark (YCSB) with a 50/50 mix of reads and writes on 4 TB of unique data consisting of 2 billion objects sized at 2K each. With a replication factor of 2, the total user data was 8 TB. Our environment included three Dell® R730xd rack-mounted servers, each with dual sockets and a total of 256 GB of DRAM. Each server utilized four 3.2 TB Micron 9300 Max NVMe SSDs for data and ran Centos Linux version 7.6. We tuned both DBMSs appropriately.

Table 1 summarizes the results when both systems were under maximum load. Aerospike demonstrated extraordinary speed at scale, providing 7.17x better throughput and 99th percentile read/write latencies that were 3.2x lower than ScyllaDB.

	Read Throughput (Transactions per second)	Update Throughput (transactions per second)	99th Percentile Read Latency (milliseconds)	99th Percentile Update Latency (milliseconds)
Aerospike	753,000	753,000	6.5	4.5
Scylla	104,950	104,950	21.0	14.3
Ratio	7.17	7.17	3.2	3.2

Table 1: Summary of throughput and latencies at maximum load

We also compared throughput and latencies when each system was operating a more typical production load of 70% of maximum. Again, as Table 2 shows, Aerospike delivered more than 7.2x the throughput with 2.9x lower data access latencies at the 99th percentile than ScyllaDB.

	Read Throughput (Transactions per second)	Update Throughput (transactions per second)	99th Percentile Read Latency (milliseconds)	99th Percentile Update Latency (milliseconds)
Aerospike	527,100	527,100	1.2	1.5
Scylla	73,465	73,465	3.5	4.6
Ratio	7.2	7.2	2.9x	2.9x

Table 2: Summary of throughput and latencies at 70% maximum load

But raw performance is only one consideration. TCO is another. We calculated the TCO of running Aerospike and ScyllaDB for 3 years with a database of 14 billion keys (28 TB of unique user data) and 20% annual data growth rate. We considered equipment, utilities, space, and labor costs, assuming an annual administrator's salary of \$85,706.

The result? As shown in Table 3, Aerospike was 7x cheaper than ScyllaDB. ScyllaDB's total cost for 3 years was just over \$4.6 million. Aerospike's cost was less than \$650,000 -- a savings of more than 85%.

	ScyllaDB			Total	Aerospike			Total
	Year 1	Year 2	Year 3		Year 1	Year 2	Year 3	
Luster Nodes	142	171	205		20	24	29	
Single server cost over 3 years (\$)	6,496				6,496			
Server Cost (\$)	922,432	1,110,816	1,331,680	3,364,928	129,920	155,904	188,384	474,208
Power Cost (\$)	99,258	119,536	143,304	362,098	13,980	16,777	20,272	51,029
Cooling Cost (\$)	137,867	166,022	199,034	502,923	19,418	23,301	28,156	70,875
Space Cost (\$)	2,461	2,461	2,461	7,383	461	461	461	1,383
Support Cost (\$)	101,418	122,131	146,424	369,973	14,284	17,141	20,712	52,137
Total Costs (\$)	1,263,436	1,520,966	1,822,903	4,607,305	178,063	213,584	257,985	649,632

Table 3: 3-year TCO comparison for ScyllaDB and Aerospike in US Dollars.

Let's turn now to the detailed results of the YCSB benchmark runs.

Load Results (Insert Rate)

To populate our databases, we used a single instance of YCSB with 250 threads. As Table 4 shows, Aerospike loaded data 7.6x quicker than ScyllaDB. We took this measurement during YCSB's load phase, which is not part of the workload phase and simulates a 100% write workload. Unlike Aerospike, ScyllaDB periodically performs compactions. When such operations are taken into account, Aerospike performed 7.9x better.

	Records Inserted	Load Time (minutes)	Compaction Time (minutes)	Total Prep Time (minutes)	Insert Rate (ktps) = Records Inserted / Load Time
Aerospike	2 billion	48	-	48	696K
ScyllaDB	2 billion	368	10	379	91K
Ratio		7.6		7.9	

Table 4: Insert rates for Aerospike and ScyllaDB

Read Results (Throughput and Latency)

Aerospike serviced reads more quickly and with greater predictability than ScyllaDB regardless of whether both systems were running at maximum load or 70% of capacity. Let's examine throughput results first.

Aerospike provided 7.17x better throughput at maximum and 7.2x better at 70% capacity, as shown in Figs. 1 and 2. Aerospike's maximum read throughput was 753,000 OPS with a 99th percentile read latency of 6.5 ms compared to ScyllaDB's 104,950 OPS with a 99th percentile read latency of 21 ms. At 70% load, Aerospike serviced 527,100 reads per second with a 99th percentile 1.2 ms latency compared to ScyllaDB's 73,465 reads with a 99th percentile 3.5 ms latency.




Figure 1: Read throughput for Aerospike and ScyllaDB at maximum load




Figure 2: Read throughput for Aerospike and ScyllaDB at 70% load

In addition to throughput, predictable read latency is critical in operational scenarios. Figs. 3 and 4 explore how the two systems fared when we measured latencies at the 99th percentile for reads. Not only did Aerospike deliver 3.2x lower read latencies at maximum capacity and 2.9x lower read latencies at 70% capacity, it did so with far fewer variances in response times. Simply put, Aerospike's performance was much more predictable -- an important consideration for mission-critical applications.




Figure 3: Read latencies for Aerospike and ScyllaDB operating at maximum load




Figure 4: Read latencies for Aerospike and ScyllaDB at 70% load

Update Results (Throughput and Latency)

Aerospike exhibited significant performance advantages over ScyllaDB for updates as well. Indeed, Aerospike delivered faster throughput with lower (and more predictable) latencies than ScyllaDB when both systems were under maximum load as well as when both systems were operating at 70% capacity.

As Fig. 5 indicates, Aerospike's maximum update throughput was 753,000 OPS with a 4.5 ms 99th percentile latency compared to ScyllaDB's 104,950 OPS with a 99th percentile latency of 14.3 ms. Fig. 6 depicts results when operating at 70% capacity; Aerospike serviced 527,100 writes per second with a 1.5 ms 99th percentile latency compared to ScyllaDB's 73,465 writes with a 4.6 ms 99th percentile latency.




Figure 5: Update throughput for Aerospike and ScyllaDB at maximum load




Figure 6: Update throughput for Aerospike and ScyllaDB at 70% load

In addition to better update throughput, Aerospike delivered lower and more predictable update latencies than ScyllaDB, as shown in Figs. 7 and 8. We measured latencies at the 99th percentile for updates and found that Aerospike delivered 3.2x lower update latency at maximum capacity and 2.9x lower update latency at 70% capacity. Furthermore, in both situations, Aerospike's latencies were far more uniformly low than ScyllaDB's.




Figure 7: Update latencies for Aerospike and ScyllaDB operating at maximum load




Figure 8: Update latencies for Aerospike and ScyllaDB operating at 70% load

What Does This Mean for You?

Our comparative benchmark with ScyllaDB shows Aerospike's clear performance advantages: 7.17x greater throughput, 3.2x lower read/write latencies, and 7.9x greater insert throughput when running at maximum load. Aerospike is predictably performant and more easily scalable than ScyllaDB at 7x lower TCO. Aerospike requires fewer clusters, and these clusters can process more transactions at lower latencies. That reduces the complexity and costs of supporting your applications.

	Aerospike Advantage over ScyllaDB
Read/write throughput	7.17x
Read/write latency	3.2x
Insert throughput	7.9x
Total Cost of Ownership (TCO)	7x

Table 5: Aerospike advantages over ScyllaDB at maximum load.

About Aerospike

Aerospike is the global leader in next generation, real-time NoSQL data solutions at scale. Its patented Hybrid Memory Architecture™ unlocks the full potential of modern hardware to deliver incredible value from vast amounts of data, powering critical applications for Airtel, Baidu, Banca d'Italia, Nielsen, PayPal, Snap, Verizon Media, Wayfair, and many other firms.