

Five Signs You've Outgrown Cassandra

While it's often easy to use Cassandra for your initial NoSQL projects, that can change quickly as your workload and data volumes grow. Service-level agreement (SLA) violations and high ownership costs are common, creating budget overruns, operational instability, and customer service issues. That's when Aerospike can help.

A NoSQL key-value store, Aerospike delivers ultra-fast, predictable performance for read/write workloads at scale, often at substantial savings in total cost of ownership (TCO). Aerospike's speed, high availability, operational efficiency, and self-management features have prompted many Cassandra users to switch to Aerospike. One firm replaced its 450-node Cassandra cluster with a 60-node Aerospike cluster and expects to save \$6.9 million in TCO over three years.

What are 5 signs that you may have outgrown Cassandra?


Sign 1	You're worried about about "server sprawl" and high TCO
Sign 2	Peak loads are disrupting your SLAs.
Sign 3	You're tired of living with cascading failures.
Sign 4	Your operations team keeps growing (and so do your costs).
Sign 5	You're struggling to acquire and retain Cassandra expertise.

Sign 1: You're worried about "server sprawl" and high TCO

It's no secret: Cassandra scales well horizontally but doesn't use resources very efficiently. That's why many production users need large clusters. Vendors love that: they price by the node. But large clusters mean more components and more hardware failures, not to mention more labor to monitor, tune, and manage the environment.

Aerospike's patented Hybrid Memory Architecture™ (HMA) operates so efficiently that it often reduces TCO to a fraction of Cassandra's. We compared the cost of Aerospike and Cassandra over a 3-year period for a 36TB database with a growth rate of 50% and 33% in years 2 and 3, respectively. Our scenario involved a balanced read/write workload operating on Dell R730dx servers with Micron SSDs, and we considered costs for servers, power, cooling, space, and labor. The result? Aerospike cut costs by 90%, saving almost \$9.27 million over three years.

Cost comparison: Cassandra vs. Aerospike


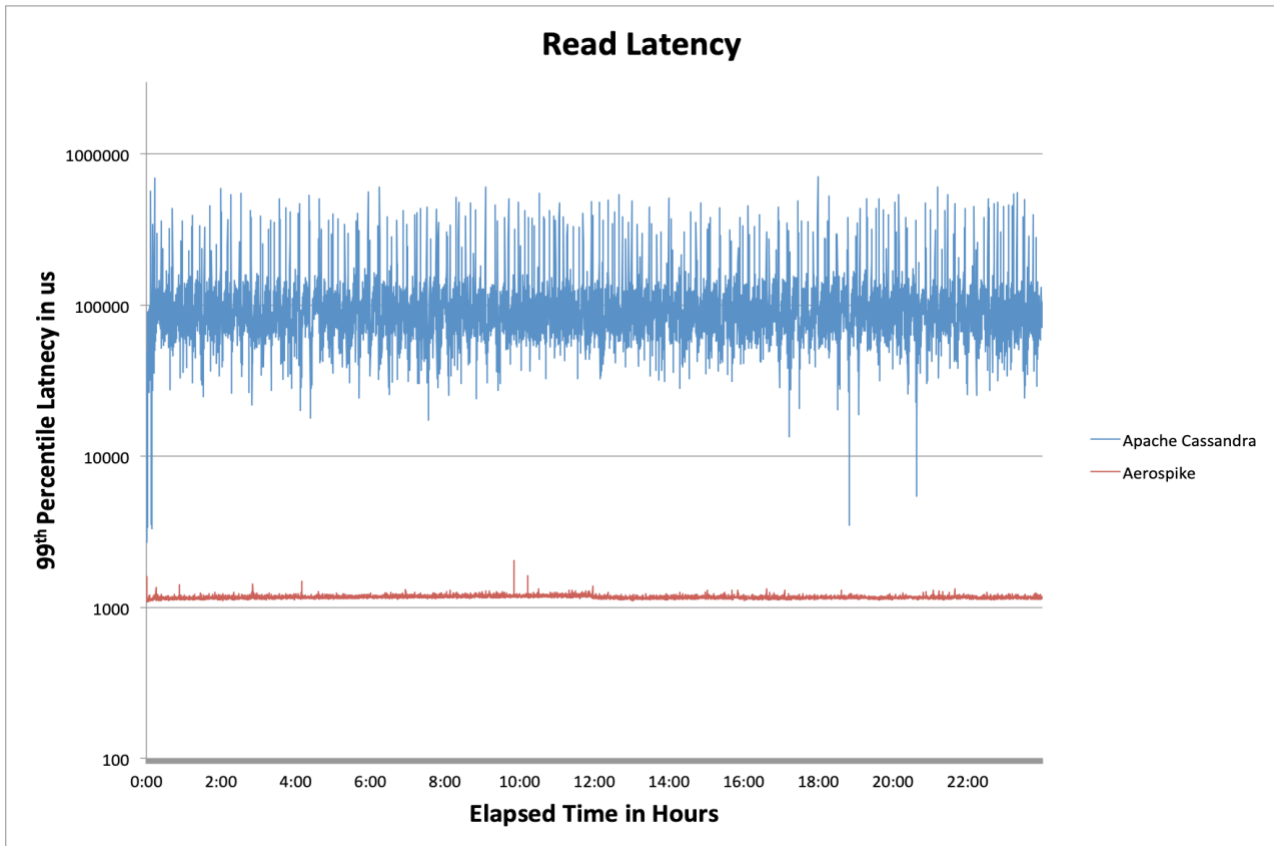


In Year 1, database has 20 billion unique objects; with a replication factor of 2, total data size = 36TB. In Year 2, data grows 50%. In Year 3, data grows an additional 33%. Total Aerospike savings: \$9.27 million.

Sign 2: Peak loads are disrupting your SLAs

Achieving strong, consistent performance is often elusive when running mixed workloads on Cassandra. Written in Java with a log-structured merge tree storage design, Cassandra requires users to tune Java virtual machines, manage garbage collectors, optimize compactions, and carefully construct data models to maximize performance. Mistakes can cause system instability and unpredictable latencies.

Furthermore, any Cassandra node can accept any read or write request even if it doesn't manage the data involved, so extra network traffic is unavoidable. This contributes to latency variances. Indeed, our benchmarks revealed huge differences in read/write latencies for Cassandra while Aerospike's latencies remained more predictable. Aerospike's response times were much faster, too.



Sign 3: You're tired of living with cascading failures

All too often, it's tough to keep Cassandra running smoothly. Cascading performance problems, node failures, and even data loss are issues Cassandra users may face due to:

- Memory pressure caused by hinted handoff during failover.
- Compactions creating caching, I/O, and/or CPU pressure.
- Memory use causing frequent garbage collection, and thus, CPU pressure and high tail latency.
- A large number of tables causing memory pressure.
- Poorly sized partitions causing pressures for the Java heap and garbage collection.
- That's why so many blogs, presentations, articles, and videos discuss such topics. If you're tired of trying to prevent one failure after another -- or fixing one failure only to encounter another -- why not evaluate your options? Aerospike's design avoids these problems, as many former Cassandra users have already discovered.

Sign 4: Your operations team keeps growing (and so do your costs)

As your workload and data volumes grow, expect your Cassandra cluster to expand rapidly. Managing very large clusters isn't easy or cheap. Consider what's required to provision such an environment. As one architect observed, ". . . (W)ith Cassandra there's a lot more configuration and tuning out of the box. Aerospike? Pretty much change a few things and you're good to go."¹

Diagnosing and fixing problems on very large clusters is labor-intensive, too. As one CTO noted, "Before Aerospike, we were spending more and more of our time on the care and feeding of Cassandra and less and less time on the building of new product offerings. With Aerospike, . . . we're just focused on adding new functionality to our platform"²

Sign 5: You're struggling to acquire and retain Cassandra expertise

Technical resources are hard to acquire and retain, and hiring staff with deep Cassandra skills is no exception. Even if you find the right people, expect to pay a premium for them -- 15%-20% more than those skilled in other DBMSs, according to recent surveys.

Chances are you'll need your staff to tune different versions of Cassandra for your workloads. In addition, if you use a Cassandra distributor, you'll need to decide if want your staff to wait until a patch is vendored in, move to a newer Apache release (and abandon your support subscription), or fix the code base you have. Each approach brings challenges.

Maybe you don't have the time or budget to build in-house expertise or hire independent consultants. Aerospike's small server footprint, combined with its self-managing and self-healing clusters, means that your staffing needs are likely to be much lower and easier to fulfill.

¹ <https://youtu.be/ljWnG8dWTms?t=203>

² <https://youtu.be/tGPqvI3BwzY?t=126>

Get ready for the future with Aerospike

If you're struggling to achieve what you want with Cassandra -- or if you've experienced any of the 5 signs just discussed -- explore what Aerospike can do for you. Email info@aerospike.com or download our detailed [white paper](#) on "Five Signs You Have Outgrown Cassandra."

About Aerospike

Aerospike enterprises overcome seemingly impossible data bottlenecks to compete and win with a fraction of the infrastructure complexity and cost of legacy NoSQL databases. Aerospike's patented Hybrid Memory Architecture™ delivers an unbreakable competitive advantage by unlocking the full potential of modern hardware, delivering previously unimaginable value from vast amounts of data at the edge, to the core and in the cloud. Aerospike empowers customers to instantly fight fraud; dramatically increase shopping cart size; deploy global digital payment networks; and deliver instant, one-to-one personalization for millions of customers. Aerospike customers include Airtel, Baidu, Banca d'Italia, Nielsen, PayPal, Snap, Verizon Media and Wayfair. The company is headquartered in Mountain View, Calif., with additional locations in London; Bengaluru, India; and Or Yehuda, Israel.