BENCHMARK

How Did Aerospike Do on Cassandra's Own Benchmark?

Speed at Scale: Better performance at a dramatically lower cost.

According to a recent post on the DataStax website, End Point, an independent database firm, ran YCSB benchmarks to stress-test Apache Cassandra, HBase, MongoDB, and Couchbase on operations typical to real-world applications. For a variety of workloads typical to modern Web, Mobile and IOT applications, Cassandra performed six times faster than HBase, and 195 times faster than MongoDB.

However, this test didn't include Aerospike, the high-performance NoSQL database. Aerospike ran three of the five benchmarks that best represent Speed at Scale workloads on our database. This paper details the results of our tests. In summary, not only did Aerospike dramatically outperform Cassandra on the benchmarks we tested; we also came in at about 1/5th of the cost of Cassandra.

An Overview of the End Point Benchmark

End Point scaled database clusters from 1 to 32 nodes using Amazon's EC2 i2.xlarge instances. The benchmark used the Yahoo! Cloud Serving Benchmark (YCSB) to generate workloads. It ran 5 different workloads for each database tested:

1. Balanced workload: 50/50 read/write.
2. Read-mostly workload: 95/5 read/write.
3. Insert-mostly workload: 10/90 read/insert.
4. Read-modify-write workload: 50/50 read/read-modify-write.
5. Mixed operational and analytical workload.

The workloads used objects with a *fieldcount* of 10 and *fieldlength* of 10 for a total object size of 100 bytes. The benchmark pre-loaded test data uniformly over the database. Each test ran until 9 million operations were performed for each workload, or for an hour if that operation count wasn't reached by then. The stated object count was 500 million objects per node, but the actual object count varied for each test; these numbers were not published by End Point.

An Overview of Aerospike on the End Point Benchmark

Aerospike used a three-node cluster to perform three of the workloads from the End Point test. The workloads performed were:

1. Balanced workload: 50/50 read/write.
2. Read-mostly workload: 95/5 read/write.
3. Read-modify-write workload: 50/50 read/read-modify-write.

AEROSPIKE

The Aerospike benchmark was set up the same way as the End Point benchmark:

- A three-node Aerospike cluster was created using Amazon's EC2 i2.xlarge instances.
- The replication factor was 1 (identical to Cassandra).
- The storage engine was configured to use the 800G SSD.
- The database was uniformly pre-loaded with objects that consisted of a *fieldcount* of 10 and a *fieldlength* of 10.
- The tests ran until 9 million operations were executed for each workload, or for an hour if that operation count wasn't reached by then.
- The database was loaded with 400 million objects per node for each of the tests performed.

Here are results for the three tests:

Balanced Read/Write

Operations per Second		Average Latency (µs)	Nodes	Annual AWS Cost
Aerospike	119,013	2,008	3	$22,416
Cassandra	119,013	22,075	11	$82,195

Read-Mostly

Operations per Second		Average Latency (µs)	Nodes	Annual AWS Cost
Aerospike	138,298	1,769	3	$22,416
Cassandra	138,298	30,918	16	$119,556

Read-Modify-Write

Operations per Second		Average Latency (µs)	Nodes	Annual AWS Cost
Aerospike	83,460	3,840	3	$22,416
Cassandra	83,460	32,697	11.4	$84,959

Using the benchmark results, the costs of running a service on cloud computing servers and bare metal servers can be further explored.

Costs vary depending on the use case. The following example will use the "Read-Mostly" workload as a use case to generate estimated hardware and labor costs for running a cluster. The labor costs are based on an average DevOps engineer's salary of $130,000 USD and on the assumption that a single DevOps engineer can manage 350 virtual servers..

Let's assume the "Read-Mostly" workload will run on Amazon Web Services. The service will run 24/365 and use an i2.xlarge EC2 instance. The i2.xlarge instance costs $0.85 USD per hour. Extending the cost out for a year produces the following numbers:

Cost Comparison: AWS

	Aerospike **(3)**	Cassandra **(16)**
i2.xlarge	$22,416	$119,556
Labor Cost	$1,114	$5,941
Total:	$23,530	$125,497

If the service grows to a Cassandra cluster size of 650 nodes, the Aerospike cluster will grow to 122 nodes. The cloud computing charges will increase substantially as the cluster size grows:

Cost Comparison: AWS

	Aerospike **(122)**	Cassandra **(650)**
i2.xlarge	$911,618	$4,856,982
Labor Cost	$45,313	$241,428
Total:	$956,931	$5,098,410

Bare Metal Comparison

If this use case requires bare metal servers rather than virtual servers, there will be equipment costs, maintenance costs, utility costs, and engineering labor costs. The bare metal server for this example is a Dell PowerEdge R220 rack server. The server-to-engineer ratio for full metal servers is 1 engineer for 75 servers. In our experience, Aerospike makes better use of bare metal than virtual servers, but to be as conservative as possible, we assume parity. Here are the results:

Cost Comparison: Dell PowerEdge

	Aerospike **(3)**	Cassandra **(16)**
R220 Server	$16,209	$86,448
Electricity	$252	$1,344
Labor Costs	$5,200	$27,733
Total:	$21,661	$115,525

As the new service adds additional nodes, labor costs will have a larger impact on the cost of running the bare metal clusters. If the Cassandra cluster grows to 650 servers, the equivalent Aerospike cluster would be 122 servers. The Cassandra cluster would take 8.67 engineers to maintain, while the Aerospike cluster would be maintained with 1.60 engineers. The following costs reflect the growth of the database cluster:

Cost Comparison: Dell PowerEdge

	Aerospike **(122)**	Cassandra **(650)**
R220 Server	$659,166	$3,511,950
Electricity	$10,500	$54,600
Labor Costs	$211,466	$1,126,666
Total:	$881,132	$4,693,216

End Point titled its benchmark paper "Benchmarking Top NoSQL Databases". The results from the benchmark tests presented in this paper clearly demonstrate that for the workloads we tested, Aerospike is the best-performing database at the lowest cost.

Sources

http://www.datastax.com/wp-content/themes/datastax-2014-08/files/NoSQL_Benchmarks_EndPoint.pdf

http://www.indeed.com/salary/q-Dev-Ops-Engineer-I-San-Francisco,-CA.html

http://instagram-engineering.tumblr.com/post/13649370142/what-powers-instagram-hundreds-of-instances

http://perfcap.blogspot.com/2012/03/ops-devops-and-noops-at-netflix.html

http://www.computerworld.com.au/article/352635/there_best_practice_server_system_administrator_ratio_/

https://f5.com/about-us/blog/articles/what-we-can-learn-about-scale-from-virtualization

AEROSPIKE

2525 E. Charleston Road, Suite 201
Mountain View, CA 94043

Tel: 408. 462. AERO (2376)
www.aerospike.com

Aerospike is the high-performance NoSQL database that delivers Speed at Scale. Aerospike is purpose-built for the real-time transactional workloads that support mission-critical applications. These workloads have the mandate to deliver informed and immediate decisions for verticals like Financial Services, AdTech, and eCommerce. The unique combination of speed, scale, and reliability can deliver up to 10x performance or 1/10th the cost compared to most other databases.