

The Real-Time Big Data Database

Scylla is an open source drop-in replacement for Apache Cassandra with scale-up performance of 1,000,000 IOPS per node, scale-out to hundreds of nodes and 99% latency of < 1msec.



Samsung Benchmark shows throughput of Scylla vs Cassandra using 2TB of data for different YCSB workloads

Scylla is trusted by internet and traditional enterprises like IBM, AppNexus, Investing.com, mParticle, and many others.



THE BENEFITS OF SCYLLA



10x Higher Throughput. Written in C++ to squeeze every bit of performance from hardware and allow up to 1,000,000 read/write operations per node.



Optimum Total Cost of Ownership. Scylla maximizes resource utilization including CPU, memory, and disk and network interfaces with linear scale-up out of the box.



Low and Consistent Latency. Lockless implementation and an independent memory management stack remove the inefficient reliance on JVM or Linux page cache and deliver consistently low latency.



Easy to Use. Apache Cassandra's wire protocol, a rich polyglot of drivers, and integration with Spark, Presto, and Graph tools mean resource-efficient and performance-effective coding.



Highly Scalable. Auto-sharding, homogeneous servers, and native multi-datacenter implementation allow seamless linear scaling without compromising on application downtime or performance.



Compaction, Streaming, and Repair Solved. Scylla's Workload Conditioning provides an umbrella of dynamic scheduling algorithms to minimize database operation latency jitter and reduce compaction streaming and repair time.



Always-On Availability. Automatic failover and replication across multiple nodes and data centers provide reliable fault tolerance.



Community Backed. Scylla has been open sourced since day one and is backed by a growing community of contributors, and Scylla leverages the big-data ecosystem around Cassandra, Spark, JanusGraph, etc.

“ ScyllaDB’s NoSQL database offers a powerful combination of low latency and high availability, making it an attractive option for customers of our Watson Data Platform offering. ”

Derek Schoettle, General Manager, *IBM Watson Data Platform*

WHAT MAKES A GOOD VS. GREAT DEPLOYMENT



Cassandra



ScyllaDB

	Cassandra	ScyllaDB
High Availability	<p>✓</p> <p>Multi-region and fault tolerant</p>	<p>✓</p> <p>Multi-region and fault tolerant</p>
High Scalability	<p>✓</p> <p>Homogeneous nodes scale to hundreds of nodes per cluster</p>	<p>✓</p> <p>Homogeneous nodes scale to hundreds of nodes per cluster. Additionally, Scylla scales up with the number of cores in our server</p>
Low Latency	<p>✗</p> <p>Unpredictable and unbounded latency, mainly the result of the JVM's GC</p>	<p>✓</p> <p>C++, shared-nothing and lockless novel design allows < 1 millisecond tail latency</p>
High Throughput	<p>✗</p> <p>Limited per-node performance. Cannot fully exploit the disk/memory/cpus</p>	<p>✓</p> <p>Fully asynchronous core engine with shard-per-core architecture allows perfect scale-up</p>
Low Management	<p>✗</p> <p>Intricate tuning, endless compactions and GC storms makes it hard to maintain</p>	<p>✓</p> <p>Out-of-the-box tuning, isolation and SLA between background and foreground jobs free DevOps time during the day and lets them sleep at night</p>

“ When we heard of a Cassandra drop-in-replacement we were skeptics. But very quickly we found it is all true—not only were the latency and GC issues completely gone, better hardware utilization allowed us to shrink the cluster size by half. ”

Gabriel Mizrahi, CTO, *Investing.com*

HOW DO WE COMPARE? YAHOO! CLOUD SERVING BENCHMARK

We compared a 3-node Scylla cluster against 3-, 9-, 15- and 30-node Cassandra clusters for throughput (operations per second), latency, and cost:

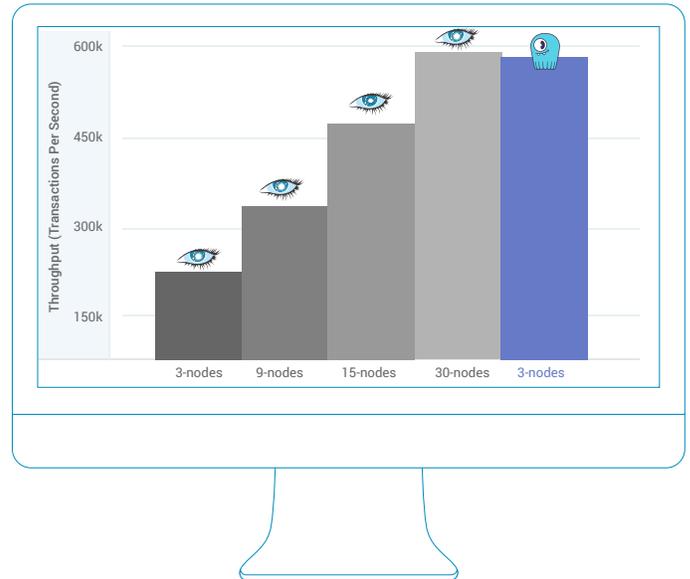
- A **3-node Scylla cluster** delivers the same throughput of as **30-node Cassandra cluster**
- Scylla achieved **1:10 ratio reduction** in the total cost of ownership
- 3-node Scylla cluster **improved latency by more than 400% (4x)**, without adversely affecting throughput, when compared to a 30 node Cassandra cluster in the 99th percentile (read workload)

Read the full benchmark [here](#).

GET STARTED TODAY

Download Open Source ScyllaDB [here](#) or contact us at: sales@scylladb.com

Take Scylla for a 1-hour [Test Drive](#): Quickly spin up a running Scylla cluster to see its performance for yourself.



“ We have observed a huge reduction in read latency by migrating from Cassandra to Scylla which enabled us to easily meet our SLAs. ”

Andrew Katz, CTO & Founder, mParticle

SCYLLADB.COM

SCYLLA.

United States Headquarters
1900 Embarcadero Rd
Palo Alto, CA 94303 U.S.A.
Phone: +1 (650) 332-9582
Email: info@scylladb.com

Israel Headquarters
11 Galgalei Haplada
Herzeliya, Israel

