

The Real-Time Big Data Database

Scylla goes beyond the usual tradeoffs to deliver scale-up performance of 1,000,000 IOPS per node, and scale-out capacity to hundreds of nodes with 99% latency under 1 millisecond.



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

Scylla powers real-time big data for companies across industries, including IBM, AppNexus, Investing.com, mParticle, and many others.



THE BENEFITS OF SCYLLA



10x Higher Throughput. Written in C++ to maximize hardware utilization and achieve up to 1,000,000 read/write operations per node.



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



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.



Compaction, Streaming, and Repair Solved. Scylla's auto-tuning capabilities include an array of dynamic scheduling algorithms that minimize database operation latency jitter and reduce compaction streaming and repair time.



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



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.



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.



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.

“ 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*

TAKING BIG DATA DEPLOYMENTS FROM GOOD TO GREAT



Cassandra



ScyllaDB

High Availability	<p>✓</p> <p>Multi-region and fault tolerant</p>	<p>✓</p> <p>Multi-region and fault tolerant</p>
Scale Out	<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>
Scale Up	<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>
Predictable Low Latency	<p>✗</p> <p>Limited per-node performance. Cannot fully exploit the disk, memory and CPUs</p>	<p>✓</p> <p>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>Auto-tuning, isolation and prioritized workloads minimize operational overhead</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 DOES SCYLLA STACK UP? YAHOO! CLOUD SERVING BENCHMARK

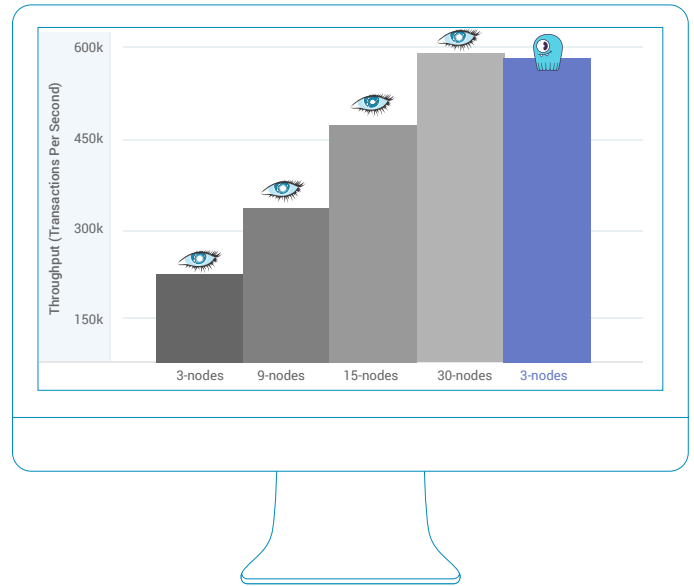
- **Scylla** achieves the same throughput with a 3-node cluster as a **30-node Cassandra cluster**
- Scylla delivers a **1:10 ratio reduction** in the total cost of ownership
- Scylla **improved latency by more than 400% (4x)** in the 99th percentile (read workload), versus a 30-node Cassandra cluster, with no impact on throughput

Read the complete benchmark at scylladb.com/product/benchmarks

GET STARTED TODAY

Spin up a live Scylla cluster and experience extreme performance first-hand at scylladb.com/test-drive.

Scylla is open source. Download it [at scylladb.com](http://scylladb.com) or contact us at: sales@scylladb.com



“ 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
2445 Faber Place, Suite 200
Palo Alto, CA 94303 U.S.A.
Email: info@scylladb.com

Israel Headquarters
11 Galgalei Haplada
Herzeliya, Israel

