SCYLLA

SCYLLADB WHITE PAPER

Cutting Database Costs: Lessons from Comcast, Rakuten, Expedia & iFood



Dealing with infrastructure costs typically isn't high on an R&D team's priority list. But these aren't typical times, and lowering costs is unfortunately yet another burden being placed on already overloaded teams.

For those responsible for data-intensive applications, reducing database costs can be a low-hanging fruit for substantial cost reduction. If you're managing terabytes to petabytes of data with millions of read/write operations per second, the total cost of operating a highly-available database and keeping it humming along can be formidable, whether you're working with open source on-premises, fully-managed database-as-a-service or anything in between. Too many teams have been sinking too much into their databases. But, looking on the bright side, this means there's a lot to be gained by rethinking your database strategy.

For some inspiration, here's a look at how several teams significantly reduced database costs while actually *improving* database performance.

60% Cost Savings by Replacing 962 Nodes + 60 Cache Servers

Comcast is a global media and technology company with three primary businesses: Comcast Cable, one of the United States' largest video, high-speed internet and phone providers to residential customers; NBCUniversal and Sky, its British telecommunications arm.

"We reduced our P99, P999, and P9999 latencies by 95% — resulting in a snappier interface while reducing CapEx and OpEx."

— Phil Zimich, senior director of engineering at Comcast

CHALLENGE

Comcast's Xfinity service serves 15 million households with more than 2 billion API calls (reads/writes) and more than 200 million new objects per day. Over the course of seven years, the project expanded from supporting 30,000 to over 31 million devices.

It first began with Oracle, then later moved to Apache Cassandra (via DataStax). When Cassandra's long tail latencies proved unacceptable at the company's rapidly-increasing scale, the team began exploring new options. In addition to lowering latency, they also wanted to reduce complexity. To mask Cassandra's latency issues from users, they placed 60 cache servers in front of their database. Keeping this cache layer consistent with the database was causing major admin headaches while also increasing costs.

SOLUTION

Moving to ScyllaDB enabled Comcast to completely eliminate the external caching layer, providing a simple framework in which the data service connected directly to the data store. The result was reduced complexity and higher performance, with a much simpler deployment model.

RESULTS

Since ScyllaDB is architected to take full advantage of modern infrastructure, allowing it to scale up as much as scale out, **Comcast was able to replace 962 Cassandra nodes with just 78 nodes of ScyllaDB**.

They improved overall availability and performance while completely eliminating the 60 cache servers. The result: a 10x latency improvement with the ability to handle over twice the requests at a fraction of the cost. This translates to 60% savings over Cassandra operating costs, saving **$2.5 million annually** in infrastructure costs and staff overhead.

	Cassandra Node Count	ScyllaDB Node Count
Recording Ring	432	18
Reminder Ring	96	18
Recording Secondary Ring	70	18
History Ring	96	6
Instruction and Lookup Ring	268	18
Total	**962**	**78**

More From Comcast →

2.5x Lower Infrastructure Costs From a 75% Node Reduction

Rakuten allows its 1.5 billion members to earn cash back for shopping at over 3,500 stores. Stores pay Rakuten a commission for sending members their way, and Rakuten shares that commission with its members.

"Cassandra was definitely horizontally scalable, but it was coming at a stiff cost. About two years ago, we started internally realizing that Cassandra was not the answer for our next stage of growth."

— Hitesh Shah, engineering manager at Rakuten

CHALLENGE

Rakuten Catalog Platform provides machine learning (ML)-enriched catalog data to improve search, recommendations and other functions to deliver a superior user experience to both members and business partners. Their data processing engine normalizes, validates, transforms and stores product data for its global operations.

While the business was expecting this platform to support extreme growth with exceptional end-user experiences, the team was battling Apache Cassandra's instability, inconsistent performance at scale and maintenance overhead. It faced JVM issues, long garbage collection (GC) pauses, and timeouts – plus they learned the hard way that a single slow node can bring down the entire cluster.

SOLUTION

Rakuten replaced 24 nodes of Cassandra with 6 nodes of ScyllaDB. ScyllaDB now lies at the heart of its core technology stack, which also involves Spark, Redis and Kafka. Once data undergoes ML-enrichment, it is stored in ScyllaDB and sent out to partners and internal customers. ScyllaDB processes more than 250 million items daily, with 60-90K read queries per second (QPS) and 18k-30k write QPS (78K-120K total).

One ScyllaDB-specific capability that increases Rakuten's database cost savings is Incremental Compaction Strategy (ICS). ICS allows greater disk utilization than standard Cassandra compaction strategies, so the same amount of total data requires less free storage space. With traditional compaction strategies, users need to set aside half of their total storage for compaction. With ICS, Rakuten can use 85% or more of its total storage for data, enabling far better hardware utilization.

RESULTS

Rakuten can now publish items up to five times faster, enabling faster turnaround for catalog changes. This is especially critical for peak shopping periods like Black Friday. Predictably low latencies allows the company to commit to impressive internal and external SLAs. **Moreover, the company has lowered infrastructure costs by 2.5 times following the 4x node reduction.**

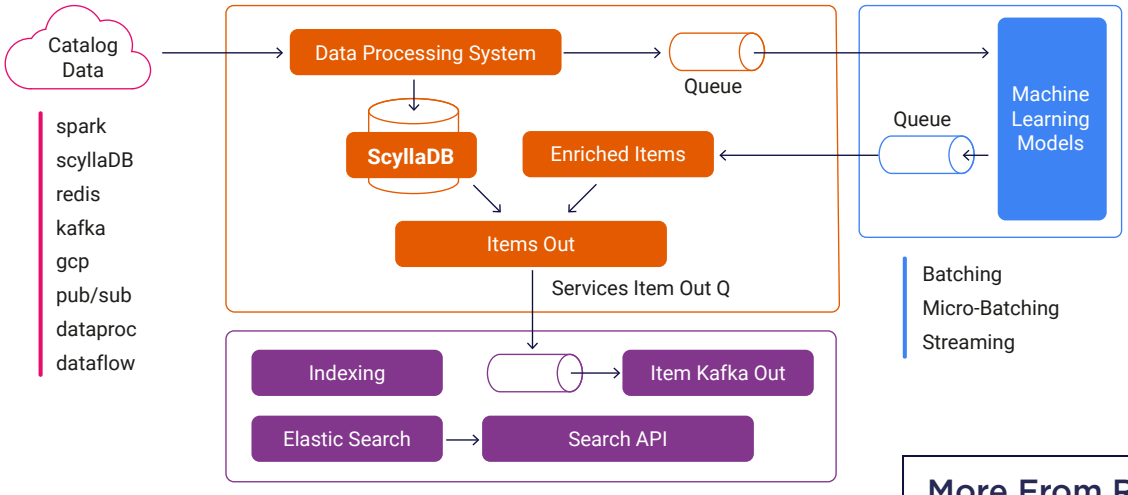

More From Rakuten →

35% Cost Savings by Replacing Redis + Cassandra

Expedia is one of the world's leading full-service online travel brands helping travelers easily plan and book their whole trip with a wide selection of vacation packages, flights, hotels, vacation rentals, rental cars, cruises, activities, attractions and services.

"We no longer have to worry about 'stop-the-world' garbage collection pauses. Also, we are able to store more data per node and achieve more throughput per node, thereby saving significant dollars for the company."

— Singaram Ragunathan, cloud data architect at Expedia Group

CHALLENGE

One of Expedia's core applications provides information about geographical entities and the relationships between them. It aggregates data from multiple systems, like hotel location info, third-party data, etc. This rich geography dataset enables different types of data searches using a simple REST API with the goal of single-digit millisecond P99 read response time.

The team was using a multilayered approach with Redis as a first cache layer and Cassandra as a second persistent data store layer, but they grew increasingly frustrated with Cassandra's technical challenges. Managing garbage collection and making sure it was appropriately tuned for the workload at hand required significant time, effort and expertise. Also, burst traffic and workload peaks cut into P99 response time, requiring buffer nodes to handle peak capacity, which drove up infrastructure costs.

SOLUTION

The team migrated from Cassandra to ScyllaDB without modifying their data model or application drivers. As Singaram Ragunathan, cloud data architect at Expedia Group put it: "From an Apache Cassandra code base, it's frictionless for developers to switch over to ScyllaDB. There weren't any data model changes necessary. And the ScyllaDB driver was compatible, and a swap-in replacement with Cassandra driver dependency. With a few tweaks to our automation framework that provisions an Apache Cassandra cluster, we were able to provision a ScyllaDB cluster."

RESULTS

With Cassandra, P99 read latency was previously spiky, varying from 20 to 80 milliseconds per day. **With ScyllaDB, it's consistently around 5 ms**. ScyllaDB throughput is close to three times Cassandra's. **Moreover, ScyllaDB is providing 35% infrastructure cost savings.**

ScyllaDB Benchmark Results

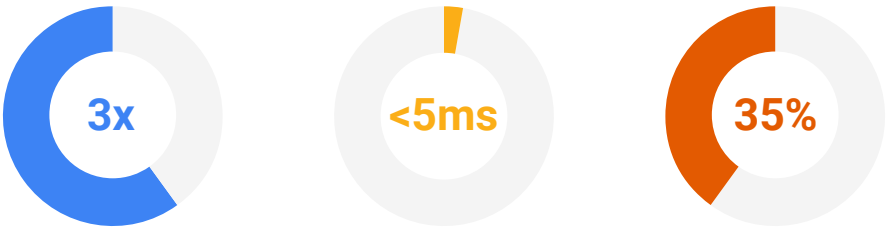

Throughput | P99 Read Performance | Infrastructure Savings

More From Expedia →

Moving Off DynamoDB to Scale with 9X Cost Savings

iFood is the largest food delivery company in Latin America. It began as a Brazilian startup and has since grown into the clear leader, with a market share of 86%. After becoming synonymous for "food delivery" in Brazil, iFood expanded its operations into Mexico.

"One thing that's really relevant here is how fast iFood grew. We went from 1 million orders a month to 20 million a month in less than two years."

— Thales Biancalana, backend developer at iFood

CHALLENGE

The short answer: online ordering at scale, with PostgreSQL as well as DynamoDB.

Each online order represents about five events in its database, producing well over 100 million events on a monthly basis. Those events are sent to restaurants via the iFood platform, which uses Simple Notification Service (SNS) and Simple Queue Service (SQS). Since internet connections are spotty in Brazil, it relies on an HTTP-based polling service that fires off every 30 seconds for each device. Each of those polls invokes a database query.

After the company hit 10 million orders a month and experienced multiple PostgreSQL failures, the team decided to explore other options. They moved to NoSQL and selected DynamoDB for iFood's Connection-Polling service. They quickly discovered that DynamoDB's autoscaling was too slow for their spiky traffic patterns. iFood's bursty intraday traffic naturally spikes around lunch and dinner times. Slow autoscaling meant they could not meet those daily bursts of demand unless they left a high minimum throughput, which was expensive) or managed scaling themselves, which is work that they were trying to avoid by paying for a fully-managed service.

SOLUTION

iFood transitioned its Connection-Polling service to ScyllaDB Cloud. The team was able to keep the same data model they built when migrating from PostgreSQL to DynamoDB.

RESULTS

iFood's ScyllaDB deployment easily met its throughput requirements and enabled the company to reach its mid-term goal of scaling to support 500,000 connected merchants with one device each. Moreover, moving to ScyllaDB reduced the database cost of the Connection-Polling service from $54,000 to $6,000 — a 9x savings.

Multiple polling systems running in parallel

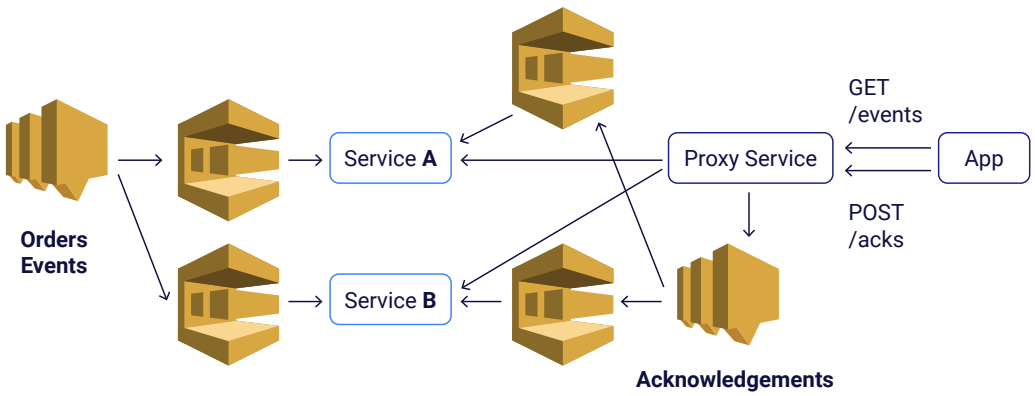

More From iFood →

Wrap Up

These examples are just the start. There are quite a few ways to reduce your database spend:

- Improve your price-performance with more powerful hardware and a database that's built to squeeze every ounce of power out of it, such as ScyllaDB's highly efficient shard-per-core architecture and custom compaction strategy.
- Reduce admin costs by simplifying your infrastructure (eliminating external caches, reducing cluster size or moving to a database that requires less tuning and babysitting).
- Tap technologies like workload prioritization to run OLTP (transaction processing) and OLAP (analytical) workloads on the same cluster without sacrificing latency or throughput.
- Consider Database as a Service options that allow flexibility in your cloud spend rather than lock you into one vendor's ecosystem.
- Move to a DBaaS provider with pricing better aligned with your workload, data set size and budget (compare multiple vendors' prices with this DBaaS pricing calculator).

And if you'd like advice on which, if any, of these options might be a good fit for your team's particular workload, use case and ecosystem, the architects at ScyllaDB would be happy to provide a technical consultation.

ABOUT SCYLLADB

ScyllaDB is the database for data-intensive apps that require high performance and low latency. It enables teams to harness the ever-increasing computing power of modern infrastructures – eliminating barriers to scale as data grows. Unlike any other database, ScyllaDB is built with deep architectural advancements that enable exceptional end-user experiences at radically lower costs. Over 400 game-changing companies like Disney+ Hotstar, Expedia, FireEye, Discord, Crypto.com, Zillow, Starbucks, Comcast, and Samsung use ScyllaDB for their toughest database challenges. ScyllaDB is available as free open source software, a fully-supported enterprise product, and a fully managed service on multiple cloud providers. For more information: ScyllaDB.com

SCYLLADB.COM

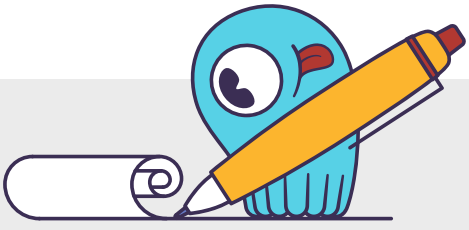

SCYLLA

United States Headquarters
1309 S Mary Ave
Sunnyvale, CA 94087 U.S.A.
Email: info@scylladb.com

Israel Headquarters
11 Galgalei Haplada
Herzelia, Israel

Copyright © 2022 ScyllaDB Inc. All rights reserved. All trademarks or registered trademarks used herein are property of their respective owners.