

Gaining the scalable performance to process up to 15 billion calling and messaging events per day with Couchbase

Industry

- [Telecommunications](#)

Customer application

- Real-time calling and messaging

NoSQL solution

- [Customer 360](#)

Use case

- User profile store

Product

- [Couchbase Server](#)

Cloud Provider

- AWS

Key features

- Multi-Dimensional Scaling
- Cross datacenter replication
- In-memory database



[Viber: Massive Data Lakes on AWS](#)

Rakuten Viber

The Viber app connects more than 1 billion users worldwide through high-quality audio and video calls, messaging, and more. To process up to 15 billion events per day, Viber needs scalable database performance. The company implemented Couchbase in a multi-layered AWS architecture. Couchbase updates user profiles in near-real time, delivering a responsive user experience. By replacing MongoDB™ and Redis with a single Couchbase database, Viber also reduced the number of servers from 300 to 120.

CHALLENGES

- Accommodate rapid, continuous growth, storing petabytes of data in a data lake
- Process 10 to 15 billion events per day – more than 300,000 events per second
- Implement a database that could provide scalable performance – while reducing the datacenter footprint and minimizing complexity

OUTCOMES

- Replaced Redis and MongoDB stack with Couchbase
- Reduced total number of servers from 300+ to ~120
- Delivers a responsive experience with real-time user profile updates
- Increased performance with half the database servers on AWS
- Simplified management with a single Couchbase tier

“With less than half the servers, we’re able to increase performance and gain a much better, scalable architecture.”

— Amir Ish-Shalom

Senior Director of Operations and Chief Architect, Viber

Maintaining a real-time vibe

Viber gains the scalable performance to process up to 15 billion calling and messaging events per day with Couchbase.

Founded in 2010, Viber is committed to providing free, high-quality calling and messaging for anyone, anywhere through its ultra-secure app. By eliminating communication costs, providing a wealth of fun stickers and GIFs, and maintaining a strong focus on privacy and security, the app has become incredibly popular. Today more than 1 billion users across the globe use Viber for audio and video calls, instant voice and video messages, and group chats.

The Challenge: Processing millions of daily events

“We operate at a massive scale,” says Amir Ish-Shalom, Senior Director of Operations and Chief Architect at Viber. “We’re processing 10 to 15 billion events per day, which peaks at over 300,000 events per second. We’re storing many petabytes of data in our data lake.” Viber previously used a combination of MongoDB and Redis NoSQL databases to support its app. But the company’s rapid growth forced system architects to reassess that decision. As the number of calls and messages grew, the company added application servers connected to the database cluster. But MongoDB did not scale well and performance lagged. Meanwhile, Redis could not provide the scalability, reliability, or sharding capabilities required. The Viber team needed a database that could provide scalable performance—without requiring a huge data center footprint.



The Solution: Scaling up with Couchbase and AWS

After evaluating a number of NoSQL databases, the Viber team selected the Couchbase data platform, which could better support the company’s tremendous growth. The team chose Couchbase in part for its easy and transparent sharding and scaling model. Compared with MongoDB, Couchbase could deliver several times the throughput using far fewer nodes.

“With Couchbase, we’re now able to handle hundreds of thousands of ops per second—an order of magnitude more than our previous solution,” says Ish-Shalom.

For Viber, their Couchbase and AWS are better together for them due to the ability to build sophisticated, scalable applications in the cloud, by leveraging the strengths of both solutions. The Couchbase implementation, which

runs on Amazon Elastic Compute Cloud (EC2) within Viber's AWS environment, is divided into multiple clusters, each dedicated to a distinct type of operation. Viber also uses a backup Couchbase cluster, which keeps copies of critical data from the other clusters, giving Viber a failsafe in case of large-scale problems.

These Couchbase clusters are part of an intricate multi-layered architecture that incorporates an Apache Storm event processor plus a wide variety of Amazon services, including Kinesis Data Streaming and Kinesis Data Firehose, Simple Storage Service (S3), Redshift, Athena, and EMR. This integrated architecture handles a variety of real-time, batch, and aggregation processes. For example, the company updates user profiles stored in Couchbase in real time, batches messages into files for storage in an S3 data lake, and aggregates data for business intelligence.

Updating user profiles in real time

The scalable performance offered by Couchbase enables Viber to continuously update user profiles—which might include account information, activity logs, and preferences— even as the number of users and volume of messages grow. “Kinesis takes each event and validates it, making sure the event is properly written,” says Ish-Shalom. “Each event then updates the user profile in the Couchbase NoSQL database in near real time.” With Couchbase, Viber can deliver a responsive user experience to its enormous user base.

Reducing the infrastructure footprint

Making the move from MongoDB and Redis to Couchbase helped significantly consolidate the infrastructure footprint. By consolidating cache and database tiers, Couchbase enabled Viber to reduce the number of servers from 300 to just 120. “With less than half the servers, we're able to increase performance and gain a much better, scalable architecture,” says Ish-Shalom.



At Couchbase, we believe data is at the heart of the enterprise. We empower developers and architects to build, deploy, and run their mission-critical applications. Couchbase delivers a high-performance, flexible and scalable modern database that runs across the data center and any cloud. Many of the world's largest enterprises rely on Couchbase to power the core applications their businesses depend on. For more information, visit www.couchbase.com.