



[DATA ANALYTICS AND DATA MANAGEMENT](#) > [NO SQL](#)

## Personalized Marketing Starts with a Powerful Database: MongoDB

Customer engagement platform Braze turns to MongoDB to optimize personalized marketing.

[Karen D. Schwartz](#) | Sep 13, 2019



You order a pizza, and almost immediately, you get a text message notifying you that your pizza is in the oven. As you're shopping for clothes online, you see a

Braze, a customer engagement platform, aims to help consumer brands humanize and personalize communications with customers through all channels: smartphones, smart watches, TVs, websites and more.

So what does this have to do with storage? Everything. Providing personalized marketing requires the ability to collect, process, segment and store half a trillion pieces of data (and counting). It also means being able to send more than 1 billion messages daily on behalf of its customers. That takes a lot of processing power along with virtually unlimited scalability and lightning-fast performance.

When Braze first launched about eight years ago, CTO Jon Hyman initially thought [PostgreSQL](#) could handle the load but quickly realized that the relational database wouldn't suffice.

“We needed something that could scale very large because we grow in a step-wise manner,” he said. “Every time we start working with a new brand or customer, we essentially acquire their user base and the activity in their mobile app, website and email.”

Performance was also important. “We wanted to make sure our customers’ engagement strategies were in real time, which to us means fast enough for it to feel like you’re interacting or having a conversation. Our requirements are around processing things in milliseconds, not minutes,” he said. “For example, if you make a purchase, you should get the receipt nearly instantaneously.”

Finally, Braze needed a way to segment its customers’ attribute data in any way they required. For example, a music listening app might send Braze information about its customers’ favorite artists, genres or playlists, while a lifestyle or wellness app might send information about customers’ meditation habits. That meant Braze

Fairly quickly, all signs pointed to MongoDB, a [NoSQL](#) database that stores data in documents instead of in a relational format. MongoDB supports a variety of search types, enables indexing and replication, and achieves load balancing by splitting data across multiple [MongoDB](#) instances, allowing it to scale horizontally. The database also can run over multiple MongoDB servers, balancing the load or duplicating data when necessary.

“MongoDB has flexible schema so you can write data like it’s a document instead of having rigid rows and columns like SQL databases,” Hyman said. “That allows us to summarize customer behavior and customer messaging histories very quickly. So for each of our customers, we can store a history of the messages we’ve sent them or their interaction with those messages, or the actions they have taken. And we can look those up very quickly because we have stored the data in a way where it’s grouped by customer.”

[Scalability](#) and performance — two of the most important requirements — also are working well with MongoDB. Today, Braze performs roughly 350,000 MongoDB operations every second and has more than 11 billion user profiles, all stored in MongoDB. Braze processes tens of billions of events every day, and its MongoDB database is growing by half a trillion rows every month.

“We have a tremendous volume of information stored in MongoDB,” Hyman said. “Every time someone visits a website that our software development kit [SDK] is on or downloads a mobile app we’re in, we create a user profile for that consumer to help the brand message them.”

MongoDB also allows Braze to segment data in any number of ways, which, in turn, allows its customers to slice and dice data any way they require to use for personalized marketing. For example, a Braze customer can choose to target users

While MongoDB had the capabilities Braze needed, it wasn't easy to fine-tune the system to reach Braze's ambitious goals. It took a lot of time working with Braze database administrators to get it to this point.

“Our job is to ingest all of our customers’ data and send whatever messages they require as fast as possible by finding the user cohorts they specify,” Hyman explained. To solve that challenge, Braze has thousands of database shards, or different servers, across many database clusters. The company has dozens of terabytes of RAM for its MongoDB databases, allowing it to store most of the data in memory for quick access.

In addition, Braze’s database administrators have improved throughput by deploying the systems in a [multi-tenant, multi-cluster](#) model. Different clusters of different sizes allow the company to isolate the data of each customer based on their size, needs and throughput requirements. That level of isolation also allows Braze to perform statistical analysis of its customers’ campaigns to determine how they can add indices to improve results even more.

While MongoDB works extremely well for Braze, Hyman says it’s not the right choice for all storage needs. It’s about using the right tool for the job, he said. On the other hand, the company retained PostgreSQL for its email analytics, which works best for tabular and columnar data.

## 0 COMMENTS

## RELATED

---

MENU

**ITProToday**<sup>TM</sup>

NEWSLETTER SIGN UP

Q SEARCH

LOG IN

REGISTER

ITProToday

[About](#)

[Advertise](#)

[Contact Us](#)

[Sitemap](#)

[Ad Choices](#)

[CCPA: Do not sell my personal info](#)

[Privacy Policy](#)

[Terms of Service](#)

[Cookie Policy](#)

Follow us:



tech

© 2020 Informa USA, Inc., All rights reserved

[Privacy Policy](#) | [Cookie Policy](#) | [Terms of Use](#)