

Customer data provider gains 5X throughput over MongoDB on half the hardware

CASE STUDY

Introduction

A leading customer data provider that collects and aggregates consumer information globally has chosen Aerospike over MongoDB to rewrite their homegrown JSON data store while moving to a unified data platform. The firm builds and manages unified data layers that help brands optimize customer experience to win new customers, and optimize marketing spend.

The firm found that current document stores have fallen short of providing the type of scale and low latency access their applications and customers need. After evaluating MongoDB - or creating yet another custom datastore - the firm chose Aerospike as their real-time document store. It cited Aerospike's high throughput and hardware efficiency as deciding factors.

According to one of their technical leads, "Aerospike performed at 5x the throughput of Mongo on half the hardware we spec'd against all the hardware we spec'd for Mongo."

The firm's core data store forms the spine that supports all of its ad-tech use cases across new and existing applications for its holding company. This datastore serves data to various applications for the company's Demand Side Platforms (DSPs) and Data Management Platforms (DMPs).

An essential requirement for their solution is a real-time user profile store. User profiles are stored as JSON document data to provide for a modern user profile's structured and semi-structured components.

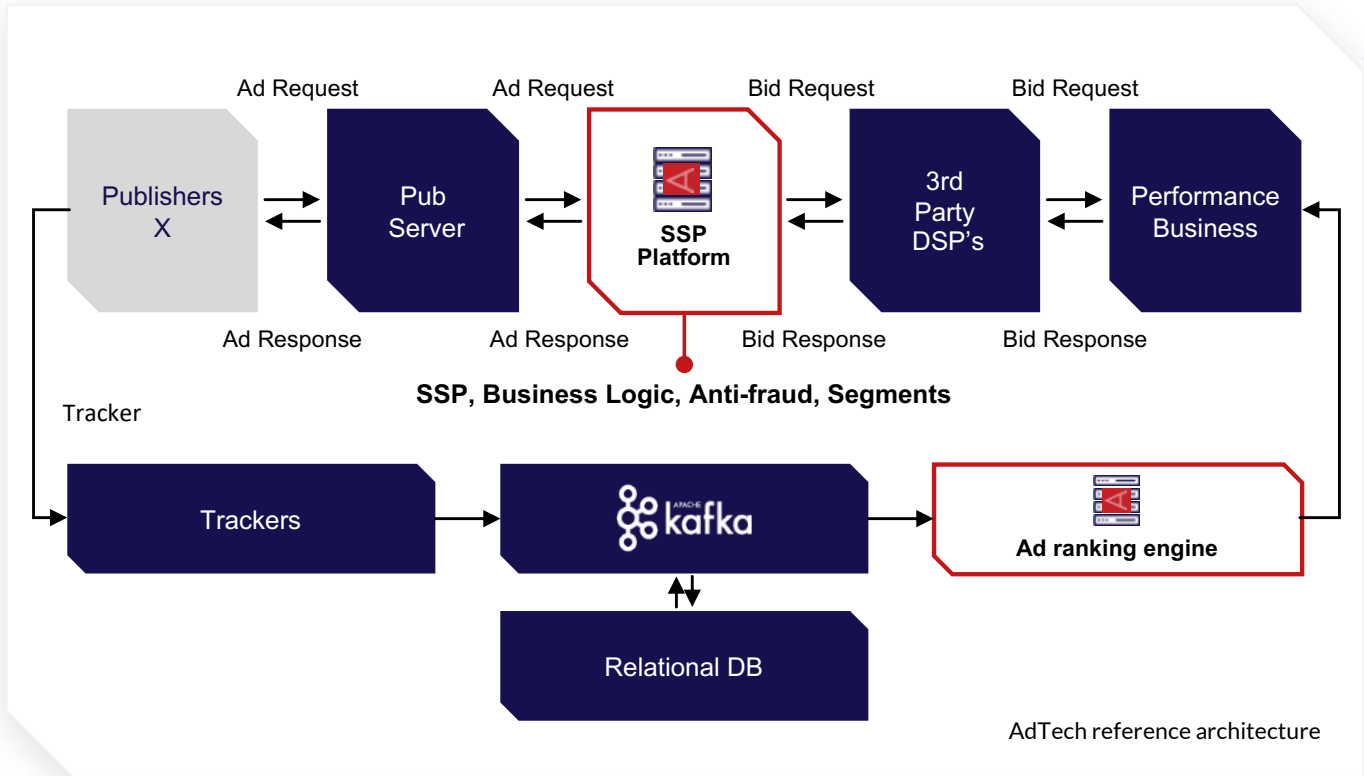
The Challenge

The company was exploring the creation of a new data product and needed to move from a batch-processing model to a real-time implementation. They wanted to rewrite their homegrown JSON datastore while moving to a unified data platform that would serve their needs for years to come.

Right out of the gate, they knew their MongoDB solution would be unable to keep up with the speed, complexity, and scale of this 9TB application.

Highlights

- Marketing services company replacing a homegrown JSON document store
- Needed to move from a batch-processing model to real-time
- Evaluated MongoDB - chose Aerospike for superior throughput and TCO
- Achieved 5X throughput over MongoDB on half the amount of hardware
- Now processing 2.5 billion records per hour, with an SLA of 2 milliseconds
- Loading time of 1/2 a day on Aerospike vs. 2 weeks with MongoDB



Results

With 2.5 billion records per hour, 1000 objects per record, and an SLA of 2 milliseconds, a real-time data platform was imperative. Once Aerospike came into the picture, they could accomplish this with half the hardware it would have taken with MongoDB. According to the technical lead, "We had data loaded in half a day and were able to perform advanced capabilities like XDR (Cross-Datacenter Replication) as a result." This contrasts with data loading times of 2 weeks with MongoDB.

Aerospike has become the de facto platform for the company as they build out its real-time data infrastructure. The key for them is real-time data access.

The company also uses Aerospike Connect for Spark for data migration and enrichment purposes. With the Spark connector, they can load data from their homegrown system into Aerospike in just hours.

Working with Aerospike offers future benefits for the company – they are looking at how they do curation and data enrichment. With Aerospike, they could bring other assets into that workflow.

< AEROSPIKE

Aerospike unleashes the power of real-time data to meet the demands of The Right Now Economy. Global innovators and builders choose the Aerospike real-time, multi-model, NoSQL operational database and data platform for its predictable sub-millisecond performance at unlimited scale with dramatically reduced infrastructure costs.

©2023 Aerospike, Inc. All rights reserved. Aerospike and the Aerospike logo are trademarks or registered trademarks of Aerospike. All other names and trademarks are for identification purposes and are the property of their respective owners.