

# EMPOWERING REAL-TIME ADVERTISING WITH PERFORMANCE AND STABILITY AT SCALE

# 150B

Daily transactions streamed

#### Industry

– Digital advertising

#### Location

– Headquarters: Chicago, Illinois, USA

“It worked better than I would’ve ever imagined. I can sleep at night. We have to be compliant with all financial regulations. We knew governance would be easier and better on the Cloudera platform. Cloudera has a strong focus on key parts instead of trying to cover the entirety of the ecosystem. They’re focusing on the things that we really use and need.”

Patrick Jaromin, Director, Software Engineering, Ad Tech, Conversant

Conversant is the leader in personalized digital marketing, performing more than a trillion real-time decisions daily about what content to put in front of 160 million people, across all major ad exchanges.

## Challenge

Conversant sees tens of billions of transactions every day, and must respond to them in milliseconds to remain competitive. With 30 percent organic annual growth, Conversant needed an infrastructure that could keep pace with the speed and scale of its business.

“We see the entire Internet because we’re interfaced with all major exchanges,” said Patrick Jaromin, director of software engineering, ad tech at Conversant. “We needed a platform that could handle the volumes we’re seeing on a regular basis. Things like relational databases weren’t working. I had what I call a ‘lost year’ when I couldn’t accomplish goals because we spent so much time trying to work with a platform that wasn’t performing. It wasn’t reliable.”

## Solution

Conversant partnered with Cloudera to create an environment to support better understanding of customers and products through real-time processing, analytics and machine learning.

One Cloudera cluster facilitates analytics and machine learning across the entire enterprise. In addition, Conversant maintains 100-plus node clusters that function as operational databases at each of its data centers for stream processing and storing information for 150 billion daily transactional events. These clusters maintain real-time online profiles and deliver query response in milliseconds.

**Solution highlights**

- Modern Data Platform: Cloudera Enterprise
- Workloads: Operational Database, Data Science, Data Engineering
- Key Components: Apache Spark™, Apache Kafka®, Apache HBase®
- Database: Greenplum

**Applications supported**

- Real-time communication personalization and targeted outreach
- Data sources
- Log files from 150 billion transaction events per day

**Impact**

- New features and product innovations drive smarter decisions

**Big data scale**

- 500-node analytics cluster

## Implementation

After struggling with the performance of legacy relational databases for its operational workloads, Conversant invested in a persistent big-data platform, but it was unstable and suffered frequent outages—impacting Conversant's core business.

In 24 hours, Conversant migrated 60 billion consumer profiles into a Cloudera cluster.

"It worked better than I would've ever imagined," said Jaromin. "I can sleep at night. We don't have those same performance issues. We're under a good deal of scrutiny in our business. We have to be compliant with all financial regulations. We knew governance would be easier and better on the Cloudera platform. Cloudera has a strong focus on key parts instead of trying to cover the entirety of the ecosystem. They're focusing on the things that we really use and need."

## Results

Without dedicating time and resources to keep the data management platform stable and performing, Conversant's team can invest its energy into product innovations, introducing new features and making smarter, faster decisions.

"The faster we retrieve the data, the faster we process, the more processing we can do," said Jaromin. "Because we're not suffering from frequent outages of the platform, I can focus on driving our platform forward and that's exactly what we're doing."