

Using Cassandra and Instacluster to deliver always-on, low-latency, analytics-driven reporting, and automation

## Case Study



ADSTAGE

## Overview

AdStage provides an all-in-one online advertising platform to manage Google, Bing, Facebook, LinkedIn, and Twitter ad campaigns together. AdStage is a technology company building advanced tools to help businesses and agencies of all sizes advertise more effectively. The AdStage Platform brings together usability, smart design, and intelligent technology to help advertisers make sense of their ad data and optimize their campaigns.

By taking a cross-network approach and integrating the major search and social advertising networks, AdStage enables customers to spend more time on their businesses rather than reporting and analyzing all of their spend and performance on individual delivery networks.

## Highlights

- Deployment: 3 clusters totaling 80+ nodes
- Sector: AdTech
- Technology: Apache Cassandra® on AWS

“*Instacluster has been instrumental in helping us manage, scale, and optimize our Cassandra deployment, allowing us to focus on delivering a great product to our customers.*”

**Jason Wu,**  
CTO & Co-Founder, AdStage

## Challenge

AdStage analyzes a vast amount of data from five major advertising networks, and its automation tools are constantly monitoring performance trends and optimizing campaigns on behalf of advertisers.

For AdStage, having downtime is not an option, and to add to this constraint any significant latency means that the platform cannot deliver the ideal solution to customers, effectively meaning that the product is down.

AdStage also needed a solution that allowed the platform to scale rapidly. Integrating a new customer to the platform can mean literally adding terabytes of data within a short time span. The company is currently experiencing massive data growth due to rapid growth in customers and the resulting added data needs.

## Solution

Ultimately Apache Cassandra® delivered the operational speed that AdStage needed to ensure that the platform had low latency and the required throughput. Cassandra's high availability architecture ensured that they didn't have any single point of failure in the environment. In addition, Cassandra's ability to simply scale by adding more nodes meant that they had covered all of their requirements.

Having a managed service looking after the application's Cassandra clusters allowed the AdStage engineering team to concentrate on their customers and to get the most value out of customer data sets.

## Migrating to Open Source Apache Cassandra

During the early stages of AdStage development the company participated in the DataStax Startup Program, and their large production cluster was deployed on DataStax Enterprise. While AdStage met the conditions of being a startup there was no licensing fees associated with their production deployment.

As the company went through a considerable growth phase AdStage was nearing the point where the company would no longer qualify for the DataStax Startup Program, and as such were facing the prospect of having to pay considerable licensing fees as the production deployment was now over 80 nodes.

The proposed increase in operational costs for AdStage was prohibitive and unsustainable. A decision was made to transition to the open source version of Apache Cassandra®.

The Instaclustr technical operations team worked with the AdStage engineering team to transition every node in both production and staging environments to the open source version of Apache Cassandra, and with the Instaclustr managed solution this was all achieved with zero downtime.

AdStage have been able to continually rely on the expertise and continued support of the Instaclustr Tech Ops team when dealing with large data ingestion and the application's scaling requirements.

## ■ ■ ■ About ■ ■ ■ Instaclustr

Instaclustr delivers reliability at scale through our integrated data platform of open source technologies such as [Apache Cassandra®](#), [Apache Kafka®](#), [Apache Spark™](#), [Elasticsearch](#), and [Redis](#).

Our expertise stems from delivering more than 70 million node hours under management, allowing us to run the world's most powerful data technologies effortlessly.

We provide a range of managed, consulting, and support services to help our customers develop and deploy solutions around open source technologies. Our integrated data platform, built on open source technologies, powers mission critical, highly available applications for our customers and help them achieve scalability, reliability, and performance for their applications.

Apache Cassandra®, Apache Spark™, Apache Kafka®, Apache Lucene Core®, Apache Zeppelin™ are trademarks of the Apache Software Foundation in the United States and/or other countries. Elasticsearch and Kibana are trademarks for Elasticsearch BV, registered in U.S. and in other countries.