



TubeMogul Delivers Big Data Insights at Enterprise Scale Using Qubole Data Service

Challenges

TubeMogul's software enables brands and advertisers to buy video advertising in real-time across multiple devices and formats using one comprehensive platform. The company's video advertising business employs Big Data analysis from ad serving, targeting, optimization and brand lift measurement. While TubeMogul launched its own fixed cluster while at the same time running on Amazon EMR, this consumed valuable engineering resources.

One of the biggest priorities at TubeMogul with Big Data was to make data easily accessible to non-engineering groups. Big Data was in high demand across the whole organization. Marketing sought insights to educate customers and establish insights around best practices and trends. Marketing research analysts needed to understand trends in viewability, inventory available, price and much more. Business Development needed Big Data to work more effectively with publishers and data partners. The business intelligence team demanded better visibility into business metrics, brand safety and inventory analytics. The machine-learning team sought to make TubeMogul's advertising more relevant by optimizing the process of presenting the right ad to the right person at the right time.

TubeMogul's fixed cluster and Amazon EMR required significant engineering time to accommodate all of its business users. The company needed engineering involvement to setup and manage clusters, to determine which nodes to run, to access keys and to complete account management tasks. Adam Rose, TubeMogul's CTO explains, "We had very large datasets and had to run queries every day to populate smaller normalized tables so that we could analyze data over time, for instance, historical price trends for the past six months. This required writing Python scripts, which can be time-consuming for our business analysts."

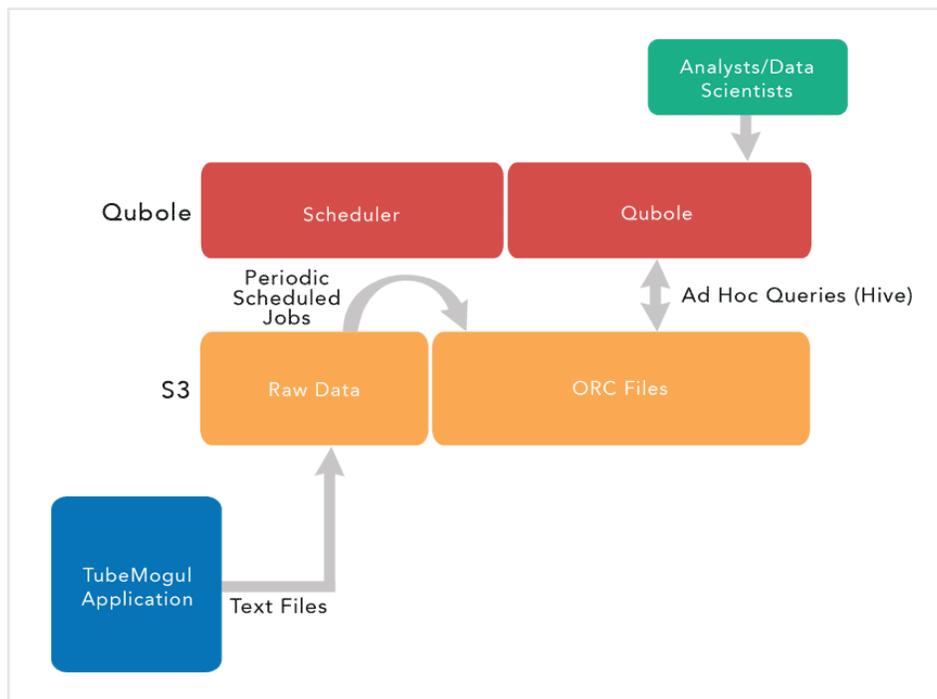
“ Qubole gave us the usability and scalability we need to drive insights from adoption of big data across our entire enterprise.



ADAM ROSE
Chief Technology Officer

Overview

- Video Advertising Platform uses Hadoop for data analysis
- Enterprise rollout to 30 users across marketing, business intelligence, machine learning and engineering teams
- Applications: ad-hoc queries, cluster management and auto-scaling for video, inventory, brand safety, and partner analytics, audience modeling
- Ability to run five times more queries at one third the cost with QDS auto-scaling
- Twice as much data processed in half the time



Why Qubole Data Service?

When the company heard about Qubole Data Service (QDS), they jumped at the opportunity for zero cost of entry, no upfront capital expenses and no long-term commitment. Specifically, QDS' auto-scaling, automated cluster management and the web-based tools that allowed its business users to efficiently schedule queries were particularly appealing.

Results

Through QDS, TubeMogul's data operations easily reached enterprise scale. This remained the case even as the company watched its data volumes swell more than 15 terabytes of log data daily.

QDS offered ease of use. QDS's web-based console simplified the processes of setting up and scaling clusters and managing accounts. Analysts then used SQL and scheduled queries to be executed at pre-defined intervals instead of writing scripts or Python code. "This was the key to our successful enterprise deployment," comments Adam Rose. "Making our business users self-sufficient got us where we wanted to be."

QDS' auto-scaling saves time and money. With this feature, TubeMogul can scale up to meet the demands of queries against large data sets when as many as 30 users are running queries simultaneously. The company no longer has to provision resources to guarantee consistent capacity.

After deploying QDS, TubeMogul enjoyed the following benefits:

- Enterprise rollout to 30 users across marketing, business intelligence, machine learning and engineering teams
- Up to a 33% cost reduction and five-fold increase in query execution with QDS auto-scaling
- Doubled the amount of data processed in half the time

Qubole's high levels of responsiveness to requests for technical support and feature enhancements further assisted TubeMogul's successful transition into a fully enterprise-scale operation.

The Future

Now that TubeMogul has successfully rolled out QDS, the company is focused on deriving even more insights through faster queries. TubeMogul will also use Presto-as-a-Service, a SQL engine that QDS has integrated to perform real-time interactive queries.

About Qubole

Qubole is passionate about making data-driven insights easily accessible to anyone. Qubole customers currently process nearly an exabyte of data every month, making us the leading cloud-agnostic big-data-as-a-service provider. Customers have chosen Qubole because we created the industry's first autonomous data platform. This cloud-based data platform self-manages, self-optimizes and learns to improve automatically and as a result delivers unbeatable agility, flexibility, and TCO. Qubole customers focus on their data, not their data platform.