



MANAGING AND OPTIMIZING A FLEET OF TRUCKS THROUGH REAL-TIME STREAMING DATA

IMPACT

- Ingesting data from nearly 200,000 trucks, with plans to increase five times to 1 million trucks on the road by 2025
- \$800,000-\$1 million savings per year in Microsoft Azure storage costs due to COD auto scaling
- Potential \$600,000 yearly operating cost savings

Challenge

This North American trucking company underwent a digital transformation as it needed a platform enabling the management and optimization of its fleet of trucks. The ability to perform predictive maintenance with efficient routing and fueling would minimize cost and provide analytics to the companies using the trucks. An automated platform was an extremely critical project to the company as it would also reduce day-to-day management and increase efficiencies.

The company's Telematics platform provides real-time connectivity with hundreds of thousands of trucks. This enables them to provide a number of solutions including remote diagnostics, over the air updates, predictive maintenance, and fuel and safety performance analytics. Fleet managers are provided real-time updates and visibility into whether trucks can finish their routes or create alternative plans if they cannot. The platform's open architecture has enabled several third-party partners and customers to leverage and integrate the system into their own proprietary telematics platforms.

Telematics is a core product offering to the trucking company and essential to its business. If they experience downtime it can result in revenue loss, increased operational costs for customers and increased warranty costs.

Solutions

The Telematics platform was built around Cludera's HDP using Apache Storm and Apache HBase to support more than 200,000 trucks on the roads every day keeping, among other things, logistics flowing during COVID-19. Storm was used for real-time streaming data processing and HBase was the data store of record for the entire platform.

They chose to deploy an east-west architecture with full data replication by migrating to CDP Public Cloud on Microsoft Azure. By bringing in data from different areas and replicating in two different regions the goal was to create two different production and testing zones - acting as a backup for failover.

The trucking company decided to become a beta customer on CDP Operational Database experience (COD) as it promised to provide a fully autonomous HBase deployment with built-in auto scaling capabilities. This would enable the company to improve its agility and reduce operational complexity. The PaaS like nature of COD combined with co-location in the Azure cloud made this the perfect solution for such a high bandwidth, low latency deployment. By moving to COD, the trucking company was able to re-architect the legacy system to separate Storm and HBase, as well as introduce Kafka (and are migrating to Flink in the near future for processing data coming in from the trucks). Traditionally, the company had all services running on the same two clusters which caused a lot of resource contention.

1.5M

Yearly cost savings (\$)

About Cloudera

At Cloudera, we believe that data can make what is impossible today, possible tomorrow.

We empower people to transform complex data into clear and actionable insights.

Cloudera delivers an enterprise data cloud for any data, anywhere, from the Edge to AI.

Powered by the relentless innovation of the open source community, Cloudera advances digital transformation for the world's largest enterprises.

Learn more at cloudera.com

Results

This new architecture enabled them to distribute load across both platform deployments, provide continuous availability in the case of an outage or disaster should an application or Azure region go down, and enable seamless upgrades and maintenance activities to happen without disrupting the services to their customers.

So far, the company has experienced a lot of success with efficiently replicating their current environment. It is expected to lead to the following results:

- Currently ingesting data from nearly 200,000 trucks, with plans to increase five times to 1 million trucks on the road by 2025
- \$800,000-\$1 million savings per year in Microsoft Azure storage costs due to COD auto scaling
- Potential \$600,000 yearly operating cost savings