

Data Infrastructure Migration and Modernization

Driving Better Decision-making for a Global Ecommerce Company



At a glance

This company is one of the world's largest loyalty programs, shopping rewards, and cashback company. TVS Next partnered with them to modernize their data infrastructure, streamline data analytics and accelerate their growth.

Services

Intelligence

- Big Data Legacy Modernization
- Data Lake
- Real Time Streaming
- Data Virtualization
- Data Replication
- Machine Learning
- Recommendation System
- Market Basket Analysis
- Feature Engineering
- Offer Personalization

Industry

Ecommerce

Problem

The company serves millions of customers monthly, and the business units constantly analyze these transactions for shopper behavior, commissions, pricing, and offers. They leverage these insights to provide the best shopping experience and guarantee cart conversions. They relied on a siloed, on-premise, legacy system to do this massive task.

As the data volume continued increasing, multiple business units started running various workloads and ad-hoc analytics on the platform. These jobs began to take a toll on the disks and processors, which reached their maximum threshold. The applications couldn't handle the scale and frequently crashed, leading to severe business disruptions. The constant maintenance and troubleshooting reduced the business units' productivity and eventually affected the company's revenue.

The company needed to build a scalable data infrastructure to process their massive data volume and to ensure that all their business units received the required data insights on time.

Strategy

The TVS Next team conducted **Accelerated Discovery Workshops** with the reward company's executives. We understood the as-is state system performance and how it prevented the company from meeting their business KPIs.

We crafted a roadmap to modernize the data infrastructure and migrate their data warehouse to the AWS cloud. Our team did a POC in 4 weeks where we migrated 15TBs of on-premise data to the cloud while supporting their existing environment. Satisfied by the agile, efficient, and seamless migration, the company partnered with TVS Next for the entire migration and modernization journey.



Transformation Journey

Our team modernized the company's data infrastructure and migrated the data warehouse workloads from the 270 TB on-premises Cloudera Hadoop cluster to the AWS cloud. Using Snowflake, we created data marts on AWS and leveraged AtScale to provide optimized query performance and virtualized data view.

We capitalized on our **Enterprise Data Platform Solution Accelerator** to build a job orchestration platform to ingest, process, and store data in the data warehouse. The data ingestion layer ingests data from various sources like PostgreSQL, MySQL, and Apache Kafka to Snowflake on AWS. The data processing layer reads data from Snowflake and processes and writes it into a final data warehouse layer. We designed the orchestration platform to maximize usability and scalability.

Even a new user to the system can now easily schedule and configure a job. Approximately 600 jobs were migrated and are currently running in the cloud using the job orchestration framework. The new data platform can process near real-time jobs and batch jobs on data, thereby enabling business units to make better decisions using the insights.

Business Outcomes

100% availability

Of data to perform ML and other advanced analytics

30% drop

In load on the on-premises computing cluster

0 business disruptions

Due to auto-scale during peak usage

98% faster

Deliveries and release cycles

68% improvement

In real-time feedback and live monitoring