

Modernizing Data Orchestration and Warehouse for a Technical Support Company

At a glance

This leading fully remote organization mainly offers tech support solutions and security software to the North American market. They partnered with TVS Next to migrate and modernize data infrastructure. They achieved accelerated growth through advanced data analytics at optimal cost.

Services

Intelligence
 Legacy Data Modernization
 ETL Automation
 Self-Serving Reports
 Databricks Datalake
 Cloud Data Warehouse
 Batch Processing
 Real-Time Processing

Industry

Technology

[Start a Conversation](#)



Problem

The company serves millions of customers monthly, and the business units constantly analyze these transactions. They rely heavily on insights from call logs, support pages, and other data sources to make critical business decisions. They relied on siloed data architecture and warehouses to shoulder 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 was growing fast, and so were its needs regarding data storage and compute resources. They 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 organization's leaders. We analyzed the as-is state system performance and understood how their architecture prevented them from meeting their business KPIs.

We crafted a roadmap to modernize the data infrastructure, build a Databricks Lakehouse, migrate their data to the AWS cloud, and quickly did a proof of concept (POC). Satisfied by the agile, efficient, and seamless POC, the company partnered with TVS Next for the entire migration and modernization journey.



Transformation Journey

First, we estimated the effort to translate the SQL scripts and identified SQL differences between the existing and new data warehouse. We then modernized the company's data infrastructure and migrated the data warehouse workloads AWS cloud. When we measured query response time for different pipelines used in the client environment while comparing the time with Redshift performance, we found that Databricks lakehouse significantly reduced the time needed to run jobs.

Since Databricks Job Cluster's pay-per-use pricing model is vastly different from what the organization was used to, we carried out a dedicated cost/performance analysis to estimate and optimize the cost of running their ETL. We reduced the overall cost also through diligent optimization efforts.

Business Outcomes

Instant availability

Of new data for advanced data analytics

Improved capabilities

Ability to run multiple workloads parallelly

10X cost reduction

By leveraging pay-per-use model

Increased productivity

By eliminating downtime

270 Minutes to 80 Minutes

Reduced time to run ETL jobs