

Migration of Analytics Workload from Snowflake to Big Query for SaaS-based applicant tracking software and recruiting platform

5 Billion

Events per week

70+ TB

Data in Snowflake

<10 secs

Looker Dashboards

CLIENT

The client is a SaaS-based applicant tracking software and recruiting platform that helps thousands of companies source, hire, and onboard top talent

AT-A-GLANCE

- 4 Billion to 5 Billion events per week
- 73TB+ data in the Snowflake storage
- Improved data freshness

PROJECT CONTEXT

The client's analytics system is based on Snowflake and Google Looker for insights. The pipeline for this system uses Kafka, MongoDB, Postgres, Vertica and FiveTran. This system handles more than 4 Billion to 5 Billion events per week with 70 TB+ data in the Snowflake storage. This system is currently facing query performance issues, ETL latencies, and inefficiencies.

The client would like to consolidate and simplify the data pipeline where possible and improve slow customer-facing queries to improve cost, CSAT, and customer retention.

PROJECT OBJECTIVES

- To demonstrate BigQuery capabilities with low query latencies.
- Simplified and more reliable pipeline and ETL with monitoring and improved data freshness as compared to the current 4hr to 6 hr lag in ETL processing.
- Integrate the existing Looker Email dashboard to BQ to demonstrate better performance.

Technology Stack



SOLUTION DELIVERY

- SquareShift delivered a lower TCO compared to the existing pipeline in Snowflake
- Looker Dashboards developed came up in <10 seconds lag
- Simplified and more reliable pipeline and ETL with monitoring
- Improved data freshness
- Simplified Architecture
- High-Quality Data/Reliability of data