



Event Driven Data Orchestration

SUMMARY

Our client, a Fortune 500 financial services company, faced challenges in orchestrating batch ETL/data curation pipelines across teams producing and consuming data. The existing approach, using CRON jobs and custom-built processes, was inefficient, leading to delays, increased complexity, and lack of centralized monitoring. A more efficient and standardized orchestration solution was needed to reduce latency for data availability, streamline feedback loops, improve security, and simplify the setup process for teams. Better visibility into performance was also a priority.

SOLUTION

To address these challenges, a pub/sub (publish/subscribe) solution was implemented, enabling consumer groups to trigger downstream processes automatically once data became available. The solution was designed to integrate AWS EventBridge for pub/sub functionality, allowing events to be registered, sent, and subscribed to using Snowflake stored procedures. AWS services like EventBridge, API Gateway, and Lambda were utilized to facilitate communication between Snowflake and downstream AWS processes, centralizing the orchestration of data pipelines. This solution was successfully deployed and data producer and consumer groups have been using the solution in production.



RESULTS

95%

Reduction in orchestration
setup time

11 hours/day

Time saved daily for
cross-team pipeline