



Case Study: Modernizing Business-Critical Data Operations with Azure Databricks

Legacy data environments evolve over years, either by adding new processes, reporting requirements and integrations as the business grows, or a combination of all those things. Over time, what was once a reliable architecture can become increasingly difficult to maintain and evolve.

A long-standing financial services client partnered with DecisivEdge to modernize approximately 50 business-critical SQL Server workloads by migrating them to Azure Databricks. The goal was to simplify operations, improve performance, centralize data processing and establish a scalable foundation for future analytics and AI initiatives - all while preserving the business processes users depended on every single day.

CHALLENGE

The client relied on a complex ecosystem of SQL Server stored procedures, SSIS packages, scheduled jobs, spreadsheets, and supporting infrastructure to power critical business processes across multiple departments across the organization. While these processes did a good job of delivering the information the business depended on, maintaining multiple technologies in different permutations took away valuable time.

The organization set out to modernize approximately 50 SQL Server jobs by migrating them to Azure Databricks. Their objectives extended to the:

- Consolidation of disparate technologies into a centralized platform
- Simplification of ongoing maintenance and operations
- Improvement of processing performance
- Reduction in infrastructure complexity
- Preparation of their data platform for future analytics and AI initiatives
- Preservation of existing business processes (with minimal disruption throughout the migration)

SOLUTION

DecisivEdge partnered closely with the client's data team to modernize their data platform while ensuring business continuity throughout the migration.

Across approximately 50 workloads, the team converted SQL Server stored procedures into Spark SQL and Python notebooks, migrated supporting tables and dependencies, rebuilt ETL workflows, reconfigured file transfers from shared drives to Azure Blob Storage, automated scheduling and

notifications, and validated every migrated process against the legacy environment before production cutover.

The project extended well beyond code conversion. Many workloads included external file ingestion, automated report generation, scheduled email delivery and interconnected job dependencies that required careful sequencing and reconciliation. Each migrated process was reviewed in detail to ensure users received the same trusted outputs they relied on before migration.

The migration also created an opportunity to examine years of legacy code with fresh eyes. By reviewing every process individually, the team identified opportunities to simplify existing logic, improve performance and, in some cases, correct legacy business logic that had accumulated over time.

RESULTS

The project delivered far more than a successful migration. It modernized the client's operational data environment, simplified ongoing administration and established a foundation for future innovation. Key outcomes included:

- **IMPROVED PERFORMANCE**

Many migrated workloads executed significantly faster after moving to Databricks, with some processes completing up to 30 times faster than their SQL Server counterparts. In one example, a process that previously required approximately 90 minutes completed in roughly five minutes, enabling the client's ETL processes to run more efficiently and making business data available sooner for reporting and operational decision-making.

- **SIMPLIFIED AND CENTRALIZED OPERATIONS**

Previously, workloads were distributed across SQL Server, SSIS packages, shared drives, MoveIT, schedulers and multiple supporting applications with numerous dependencies and permissions to manage. Following the migration, notebooks, jobs and data were centralized within Azure Databricks, reducing operational complexity and simplifying ongoing maintenance.

- **ENHANCED MONITORING AND GOVERNANCE**

By leveraging Delta Lake, the client gained built-in version history, logging and monitoring capabilities. Instead of maintaining separate audit, staging and archive tables, many of these capabilities became native features of the platform, improving operational visibility while reducing administrative overhead. Native job monitoring and automated failure notifications also made it easier for support teams to identify and resolve issues quickly.

- **A PLATFORM THAT'S READY FOR WHAT'S NEXT**

Beyond solving immediate operational challenges, the migration established a modern, scalable data platform capable of supporting future analytics, cloud-native development and AI initiatives. Standardizing on Databricks, Spark and Python provides the client with greater flexibility as business requirements continue to evolve.

The success of the project expanded DecisivEdge's role beyond its traditional application work, leading to additional Databricks modernization projects, engagement with new business units and a growing

backlog of future work.

ABOUT DECISIVEDGE

DecisivEdge leads organizations through their transformation journey. We believe that the right mix of people, processes, and technology will drive modernization, business optimization and automation. To learn more, please visit our website at www.decisivedge.com, email us at getresults@decisivedge.com or contact us at 302.299.1570.

* CASE STUDY *