HEXAWARE

Financial Services

Global Investment Management Firm

A New Unified Analytics Experience: Migrating an Investment Data Lake from AWS to Fabric

Business Problem

- Enterprise Strategy: US Based client wants to migrate Data Lake to MS Fabric as per Enterprise Strategy.
- Existing AWS-based metadata-driven ingestion estate (Python, Glue, Step Functions, Lambda, crawlers) created operational complexity and multi-service overhead.
- Need to preserve proven Python business logic, minimize rewrites to reduce migration risk, and maintain Power BI reporting continuity with minimal semantic changes.

Solution

- Rehosted metadata-driven Python ingestion framework into Microsoft Fabric Python notebooks to maximize code reuse.
- Replaced AWS Glue/Step Functions/Lambda with Fabric notebooks, pipelines, and schedule-based triggers (DAGs).
- Implemented Medallion OneLake (Bronze / Silver / Gold) and centralized configuration (11 tables + mapping spreadsheet) in the OneLake catalog.
- Governed paths from curated Gold datasets into the Lakehouse and onward into Fabric SQL for robust reporting and SQL-driven analytics.
- Repointed and validated three Power BI reports to the new Fabric data layers.
- Implemented monitoring, alerting, and governance for operational resilience and auditability.

25%
Faster time-to-market through high code reuse and automated pipelines

Up to 35%
Reduced cloud run and operational costs by consolidating services

Improved
Developer productivity with Fabric-native notebooks and SQL endpoints

HEXAWARE

Utilities

Global Energy Trading and Renewables Company

Centralizing Near Real-Time Data Access with Microsoft Fabric

Business Problem

- Existing data platform lacked medallion architecture causing unreliable, fragmented, and non-real-time data processing.
- Data latency, quality issues, and absence of standardized data layers hampered decision-making and reporting efficiency.
- Need for a centralized, scalable, and governed data platform supporting near real-time analytics for Commercial Line of Business.

Solution

- Implemented Microsoft Fabric-based medallion architecture (Bronze / Silver / Gold layers) organizing ~300,000 records into a canonical gold layer dataset.
- Incorporated Change Data Capture (CDC) and data classification for near real-time updates with 3 to 5-minute latency, enforcing data accuracy through primary keys, deduplication, and automated quality checks.
- Developed end-to-end data pipelines ensuring reliability, fast query performance, and automated deployment with schema evolution and testing.
- Established a phased roadmap with initial implementation focusing on latency SLAs and data accuracy; Microsoft Purview governance planned in later phases.

70%
Reduction in latency

35%
Increase in operational efficiency through automation, minimizing errors and delays

25%
Improvement in data quality