DIGGIBYTE

CFO Financial Analytics Modernization

Migrating Synapse workloads to a scalable, secure and governed Databricks platform

CLIENT INFO:

NOKIA

Client:
Nokia

Use Case:
Migrating Synapse workloads to Databricks platform

OUTCOME MAP ALIGNMENT

Azure Synapse → Diggibyte → Databricks

~80% Less development effort and code redundancy

70 GB Structured financial data

SLA-ready CFO reporting pipelines

A reusable, configuration-driven data engineering foundation transformed financial reporting workloads into governed, reliable and scalable Databricks pipelines.

DIGGIBYTE

THE CHALLENGE

Modernizing a complex financial analytics estate

Nokia needed to transition its CFO financial analytics platform from Azure Synapse to Databricks while protecting reporting continuity, improving engineering efficiency, and strengthening enterprise governance.

MIGRATION COMPLEXITY

Database logic, stored procedures, Synapse views and notebooks had to be transitioned efficiently

Financial data arrived from ADLS Gen2 and SQL Server and moved through several processing layers.

Pipelines needed safe reprocessing without duplicate financial records.

BUSINESS REQUIREMENTS

CFO dashboards required reliable, consumption-ready datasets within defined SLA timelines.

Delivery teams needed a consistent way to build and onboard pipelines with less code duplication.

Sensitive financial data required fine-grained access, auditability and controlled releases.

BUSINESS OBJECTIVE

Improve finance analytics efficiency with a scalable and governed platform that accelerates delivery while sustaining trusted CFO reporting.

DIGGIBYTE

THE SOLUTION

A reusable financial data platform built on Databricks

Diggibyte implemented a configuration-driven ETL framework and a governed Medallion architecture to standardize processing from ingestion through CFO consumption

WHAT WE DID

01	**Configuration-driven ETL**	A reusable framework standardized pipeline development and enabled new workloads with minimal code changes
02	**Delta Lake processing**	MERGE-based processing enabled safe, recoverable reprocessing without data duplication
03	**Medallion architecture**	Bronze → Silver → Gold → Datamarts separated raw ingestion, standardization, business aggregates and reporting tables.
04	**Performance & cost**	Serverless compute, autoscaling, job clusters, and Delta optimizations aligned resources with workload patterns
05	**Governance**	Unity Catalog with RBAC and ABAC provided fine-grained access, audit logging, and monitoring.
06	**DevOps automation**	Git, CI/CD and Databricks Asset Bundles standardized deployments and environment management

Data footprint: Aapproximately 60-70 GB of structured financial data across processing layers, including raw data, curated datasets, business aggregates, reporting datamarts and data-quality metadata

DIGGIBYTE

IMPACT & OUTCOME

Measured engineering efficiency and dependable CFO reporting

The completed implementation delivered documented gains in engineering productivity, reliability, governance andreporting performance.

BUILD ON DATABRICKS

~80%
Reduction in development effort and code redundancy

60-70 GB
Processed across financial data layers

Within SLA
End-to-end pipeline execution

QUANTIFIED IMPACT

1. **Nearly 80% less development effort and code redundancy.**
2. **End-to-end pipeline execution achieved within defined CFO reporting SLA timelines.**
3. **60-70 GB of structured financial data is processed across the solution.**

SUSTAINABLE VALUE

1. **Reusable framework accelerates onboarding of new pipelines.**
2. **Delta MERGE enables consistent and recoverable processing**
3. **Centralized governance improves control of sensitive finance data.**

REQUIRED BEFORE EXTERNAL SUBMISSION

Prior written approval from Diggibyte is required before the submitted customer story can be used for any public marketing purposes.

Contact Us: www.diggibyte.com sales@diggibyte.com