DIGGIBYTE

Enterprise Synapse to Databricks Migration

A governed Lakehouse foundation for Pharma analytics

CLIENT INFO:

HEMAS

Client:
Hemas

Use Case:
A governed Lakehouse foundation for Pharma analytics

OUTCOME MAP ALIGNMENT

Azure Synapse → Diggibyte → Databricks

~100
Synapse Pipelines

~311
Synapse Notebooks

~304
Medallion Tables

Customer-validation draft: The supplied source is a Statement of Work. Program scale is documented; realized business outcomes must be confirmed after go live before submission as a success story.

DIGGIBYTE

THE CHALLENGE

From fragmented Synapse workloads to a governed analytics foundation

Hemas required a structured path to modernize a sizeable Pharma data estate while preserving historical data, daily processing, downstream reporting, and business continuity.

BEFORE DATABRICKS

~100
Ingestion and transformation pipelines required assessment and prioritization.

~311
Notebooks needed migration or selective refactoring for Databricks jobs.

~304
Bronze, Silver and Gold tables required schema conversion, validation and reconciliation

Historical data and subsequent incremental processing
had to be supported without disrupting dailybatch operations

BUSINESS LIMITATIONS

Platform performance, scalability, governance, and cost efficiency needed improvement.

Access control, lineage, metadata, and audit capabilities needed centralized governance.

Reporting users needed trusted consumption layers plus modern AI/BI experiences.

Migration scope and sizing needed validation through discovery and data profiling

MODERNIZATION PROGRAM

A three-month modernization program designed to move daily-batch Pharma workloads from Azure Synapse to a unified, governed Databricks Lakehouse

DIGGIBYTE

THE SOLUTION

An end-to-end Synapse migration approach on Databricks

Diggibyte designed a phased delivery model covering discovery, workload migration, governance, validation, consumption enablement, and production stabilization

WHAT WE DID

01	**Discover & prioritize**	Inventory pipelines and sources, validate access, profile data and agree migration success criteria.
02	**Migrate & refactor**	Move Synapse ingestion and transformation workloads; refactor notebooks and jobs where required.
03	**Load & reconcile**	Execute full historical loads, enable incremental processing, migrate DDLs and validate counts and schemas.
04	**Govern & secure**	Implement Unity Catalog, role-based permissions, metadata management, lineage and auditing.
05	**Enable consumption**	Connect reporting tools and prepare 2 AI/BI Genie Spaces and 8 AI/BI Dashboards for the Pharma unit.
06	**Cut over & stabilize**	Support production deployment, go-live validation, sign-off and hypercare.

HOW WE DID

The Databricks capabilities: Databricks Lakehouse, Delta Lake, Unity Catalog, Databricks Workflows, AI/BI Dashboards, AI/BI Genie and Delta Sharing proof of concept

DIGGIBYTE

IMPACT & OUTCOME

Documented program scale and intended value

The SOW documents migration scope and expected platform benefits. It does not yet contain post-production measurements; therefore, no realized savings, acceleration, or SLA improvement is claimed here.

BUILD ON DATABRICKS

3 months	Daily batch	Full + incremental
Planned program	Primary cadence	Data loading

INTENDED TECHNICAL VALUE

1. Unified Bronze / Silver / Gold architecture.
2. Centralized access, lineage, and audit controls.
3. Optimized pipelines designed to meet agreed SLAs.
4. Historical continuity plus incremental processing

INTENDED BUSINESS VALUE

1. Trusted data for reporting and advanced analytics
2. Improved scalability and operational manageability
3. Governed self-service through AI/BI and Genie.
4. A reusable foundation for future Pharma use cases

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