Case Study

Achieving 2ms Query Response Times and Managing 1.3 Billion Transactions via a Governance-Backed Big Data Platform

Client Overview

High-volume financial services and home finance providers handle immense, multi-channel transactional workloads where a lack of consolidated storage frameworks can slow analytical speeds and result in transaction drops. This case study covers the implementation of a governed big data warehouse platform for a leading housing finance NBFC. By combining Databricks Unity Catalog with an 8-node Hadoop cluster, automated multi-tier pipelines, and Hive discovery gateways, the project successfully brought together 6 complex source systems—reducing query latency to 2ms and optimizing processing across a 1.3 billion transaction footprint.

The Challenge?

Decoupled Data Infrastructure

Lack of a centralized data warehouse to combine streaming and batch inputs, preventing a clear single source of truth.

Performance Processing Scale Limits

Inability to run deep, high-volume data analytics on legacy technology platforms without system slowdowns.

Transactional Drop Risks

Processing errors causing payment transaction drops, affecting consumer satisfaction and payment gateway success rates.

Limited Advisory Support

Inability to generate clear business intelligence reporting to support merchants using the payment platform.

Solution Delivered

Unity Catalog Control Plane

Implementation of Databricks Unity Catalog to serve as a secure metastore and user-permission dashboard across distributed workspaces.

Multi-Layered Access Controls

Application of RBAC and ABAC rules to dynamically protect information at the catalog, schema, table, and column tiers based on regional branch or user role.

Robust Cluster Engineering

Deployment of a high-performance Hadoop infrastructure running an 8-node, 2-master cluster setup designed to process 16 TB scale data effortlessly.

Hive Analytical Engine

Integration of Hive query tools to provide a scannable, SQL-like entry point across distributed databases and file systems.

Architecture

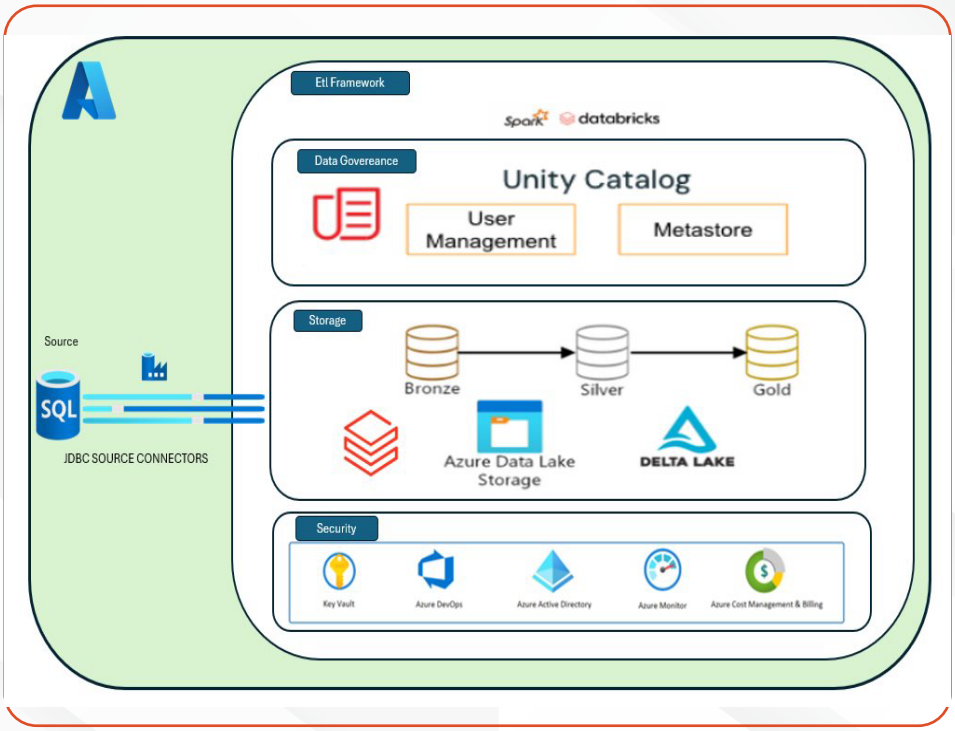

The Impact

Rapid Analytical Execution

Realized guaranteed data accessibility with a **2ms query response time,** removing traditional business intelligence reporting delays.

Massive Volume Management

Successfully processed and secured an enterprise volume of **~1.3 billion financial transaction counts.**

Optimized Data Pipeline Schedules

Cut down integration processing windows, completing full data loads in **7 hours** alongside an automated **~20-minute incremental refresh cycle.**

High Streaming Capacity

Supported continuous streaming ingestion workloads tracking at **~6 m.s** across an active **~16 TB data footprint.**

Improved Payment Metrics

Reduced transaction drops across the core gateway, resulting in higher checkout success rates.

Upgraded Merchant Advisory Power

Provided partners with comprehensive BI advisory capabilities, strengthening merchant platform adoption.

Outcome

Unifying multi-channel financial systems does not require compromising data governance or pipeline processing speeds. By combining fine-grained access controls with an elastic Hadoop cluster setup, this prominent home loans provider successfully eliminated historical operational limits. The solution provides the precise speed, scale, and dynamic security required to lower manual management costs, fulfill compliance requirements, and build a highly competitive financial analytics ecosystem.

Interested in a customized solution designed to meet your unique requirements and achieve comparable results?

At Exponentia.ai, we partner with businesses to address complex data challenges and build trust in enterprise data. Our approach helps organizations streamline data management, improve visibility, and empower teams with reliable insights for better decision-making.

Engage with Us

www.exponentia.ai

engage@exponentia.ai

US UK UAE IND SGP