



Case Study

Modernizing Data Governance: Hive Metastore to Unity Catalog Migration for a Leading NBFC

Client Overview

A leading Indian Non-Banking Financial Company (NBFC) serving millions of customers. Under RBI regulation, the client required a modernized governance layer to handle sensitive financial data across multiple business units without compromising on security or performance.

The Challenge?

Data Fragmentation



Siloed workspaces (AMC, IHT, MRHFL) prevented seamless collaboration and created data duplication.

Infrastructure Risks



Manual SAS token management and uncontrolled cluster configurations led to security vulnerabilities and subnet exhaustion.

Operational Bottlenecks



Manual library installations and a lack of centralized audit trails hindered IT efficiency.

Data Loss Vulnerability



HMS lacked the robust recovery and time-travel features required for mission-critical financial data.

Solution Delivered

Centralized Metastore



Replaced fragmented HMS instances with a single Unity Catalog Metastore for standardized governance.



UCX Automation

Leveraged the UCX toolkit to migrate 3,800+ tables, 140+ views, and 68+ scheduled jobs with minimal downtime.



Storage & Compute Governance

Registered External Locations using Azure Managed Identities and enforced cluster policies to control costs and machine configurations.



Enhanced Auditability

Leveraged UC system tables for query history, lineage tracking, and monitoring.

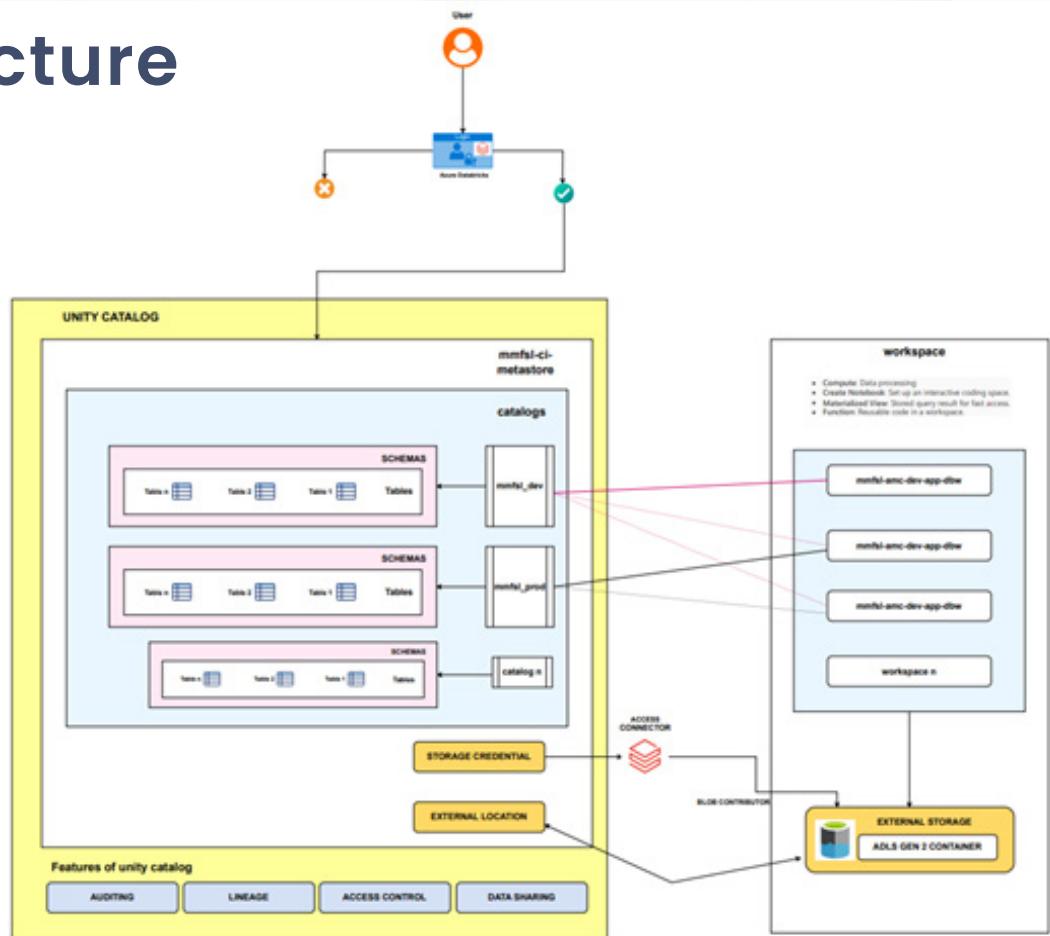
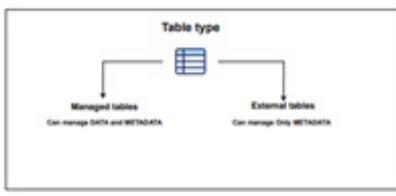


Operational Simplification

Introduced Databricks Volumes for secure library management and shared file configurations.

The Architecture

Admins		
Role	Permissions	Permissions
Account Admin	Manage Account & All Workspaces	Can create a Workspace Assign roles to workspace Configure global security
Workspace Admin	Manage Unity Catalog	Create & manage catalogs, schemas, and tables Set data access permissions Manage data governance & lineage Configure & review alerts
Workspace Admin	One specific Workspace	Administrators from workspace Assign workspace roles Manage settings Monitor usage



The Impact

Projected 25-30% Reduction

in table storage sizes through optimized formats.

Rapid Deployment

Migration of complex environments (IHT_Prod) completed in under 20 days.

Strict Compliance

Delivered an RBI-compliant environment with fine-grained RBAC and end-to-end data lineage.

Risk Mitigation

Introduced advanced safeguards, including table recovery features and secure, admin-controlled library installations.

Outcome

By migrating to Unity Catalog, the institution successfully moved from fragmented, high-risk data silos to a unified, secure architecture. The transformation didn't just move data; it established an RBI-compliant foundation that supports secure AI experimentation and allows the organization to scale analytics with absolute confidence in their data lineage.



Want to get the same benefits for your business?

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

