

From Cloud-Modernized Data to Proactive Enterprise Intelligence

Building a governed, AI-ready Azure Databricks data intelligence foundation across multiple business divisions.



About the Client

Our client is a large global enterprise operating multiple business divisions with an established Microsoft Azure data and analytics environment.

Over several years, the organization had progressively modernized data from ERP, CRM, Finance, Commercial, Operations and other enterprise systems. Individual divisions had introduced cloud data platforms, analytical data marts, Databricks workloads and Power BI solutions to address their respective business requirements.

The enterprise was therefore not starting from a legacy-only environment. It had already made substantial progress in cloud modernization.

The challenge was that modernization had occurred division by division and use case by use case. As the estate expanded, data was increasingly centralized in Azure, but business definitions, quality controls,

analytical logic and governance had not evolved as one cohesive enterprise foundation.

This became especially important as the organization began expanding self-service analytics, Chat with Data and AI-driven analytical experiences.

PalTech was engaged to help evolve the environment from a collection of modern cloud capabilities into a governed enterprise data intelligence foundation ready for BI, AI and proactive analytics.

Business Challenges

Data Was in the Cloud, but Business Meaning Remained Fragmented

Critical logic was distributed across data pipelines, SQL, notebooks, analytical marts and Power BI semantic models.

Definitions for metrics such as revenue, margin, customer performance and operational efficiency could vary between business domains, making it difficult to establish which interpretation should be

used for an enterprise-level question.

Cross-Division Analysis Required Manual Reconciliation

Finance, Commercial, CRM and operational teams had strong visibility within their individual domains, but questions crossing those boundaries still required substantial analyst intervention.

A question such as “Why did margin deteriorate this month?” could require multiple datasets and teams before price, volume, product mix, customer or operating-cost drivers could be understood.

Governance Had Not Scaled Consistently

Ownership, lineage, certification, quality monitoring and access controls had evolved differently across workloads. Data teams could not always quickly determine which data product was authoritative or which downstream analytical assets would be affected by an upstream issue.

AI Exposed the Semantic and Trust Gap

Early conversational analytics demonstrated the value of natural-language access to enterprise information, but also exposed the limitations of the underlying architecture.

AI could retrieve technically correct data while applying the wrong KPI definition, hierarchy, relationship or business context.

Trusted enterprise AI therefore required more than connecting an LLM to data. It required governed data, trusted semantics and contextual understanding.

Analytics Was Still Reactive

Modern dashboards improved visibility, but users still had to discover issues themselves, drill through multiple views and manually develop an explanation.

The next stage was to make the platform increasingly capable of surfacing what

deserves attention rather than waiting for someone to find it.

PalTech's Solution Approach

PalTech approached the engagement as an enterprise data intelligence transformation, not simply another migration.

The solution standardized the underlying engineering and governance foundation while allowing individual business divisions to retain the domain-specific models and definitions required by their businesses.

1. Governed Azure Databricks Data Foundation

Azure Databricks and Delta Lake formed the core processing and data intelligence foundation.

Reusable ingestion and engineering patterns supported enterprise data from ERP, CRM, Finance, Commercial, Operations, APIs, files and external sources, using batch, incremental, CDC and streaming patterns where appropriate.

Data was progressively refined through:

Bronze → Silver → Business-Aligned Gold
Data Products

Existing workloads could coexist during modernization, avoiding the need for a disruptive big-bang migration.

The first domains established reusable engineering patterns that subsequent divisions could adopt instead of rebuilding ingestion, validation, security and deployment capabilities.

2. Data Quality, Governance & Continuous Trust

Unity Catalog provided the governance foundation for data and AI assets, supported by PalTech's data-quality and trust patterns.

The solution embedded:

- Ownership and discovery
- End-to-end lineage
- Policy-based access
- Data-product certification
- Quality and reconciliation controls
- Freshness and completeness monitoring
- Auditability and operational visibility

Quality was treated as a continuous platform capability rather than an end-of-pipeline validation exercise.

3. Semantic & Intelligence Layer

PalTech established a business-aligned semantic layer over trusted enterprise data products.


The layer defined and governed:

- Enterprise KPIs
- Business dimensions
- Domain-specific measures
- Business terminology
- Relationships and hierarchies
- Synonyms and context
- Authoritative analytical sources

This provided a consistent analytical contract for both BI and AI without forcing every business division into one physical enterprise model.

Power BI consumed governed analytical datasets through the Databricks SQL serving layer, while Databricks conversational and AI experiences operated against the same trusted data and business context.

Enterprise Information Flow



Information-Flow Principle

Source data primarily flows into the governed platform.

BI, Chat with Data and agents interact bidirectionally at the query and analysis layer: users or agents submit questions, the governed platform evaluates data and semantics, and contextualized results are returned.

Human feedback or approvals can feed analytical workflows, but the architecture does not assume uncontrolled autonomous write-back into ERP, CRM or other systems of record.

Moving Beyond Chat with Data

Conversational analytics allows a user to ask:

“Why did margin decline?”

Proactive intelligence changes the interaction.

The platform can identify that margin has moved materially, analyze the contributing dimensions and surface the finding before an analyst manually discovers it.

Finance & Commercial

Detect material margin or revenue movement and analyze contributing business unit, product, geography and customer dimensions.

Commercial Intelligence

Surface unexpected changes in order volume, conversion, customer or product performance with relevant analytical context.

Data Operations

Identify stale or incomplete business-critical data products and provide data teams with context around likely dependencies and downstream analytical impact.

The objective is not autonomous business decision-making.

It is to shorten the time between an important change occurring and the right person understanding why it deserves attention.

Accelerators and Reusable PalTech IP

Agentic Data Quality & Trust

Designed to improve data trust continuously as the enterprise platform expands.

Capabilities

- Automated data profiling
- Metadata enrichment
- AI-assisted DQ-rule recommendations
- Human validation of critical controls
- Continuous monitoring and anomaly detection
- Contextual explanation of trust issues

Value

Reduces manual validation effort while helping data owners identify quality issues earlier and understand their business impact.

AI-Accelerated Migration & Modernization

Uses AI-assisted engineering and gated validation to accelerate modernization of existing data and analytical assets.

Capabilities

- Legacy-asset discovery
- Dependency and transformation analysis
- Source-to-target mapping
- AI-assisted code and report conversion
- Automated target validation
- Human approval at critical checkpoints

PalTech's MCP-driven migration approach provides representative proof of this engineering model: automation accelerates execution while validation gates keep people in control.

Lakehouse Optimization & FinOps

Extends modernization into continuous platform optimization.

Capabilities

- Cost and consumption monitoring
- Cost-anomaly identification
- Cluster and workload rightsizing
- Query and pipeline optimization
- Storage-growth analysis
- Budget-risk visibility

Value

Creates stronger accountability for platform performance, consumption and operating cost across business domains.

Business Value Delivered

[~40–50%] Faster Domain Onboarding

Reusable ingestion, transformation, governance and deployment patterns reduced the effort required to introduce new data domains and analytical workloads.

[~50–60%] Lower Manual Data-Quality & Validation Effort

Automated profiling, quality rules, monitoring and reconciliation reduced repetitive data-validation activities and

enabled teams to focus on higher-value exceptions.

[~25–35%] Lower Lakehouse Operating Cost

Improved workload visibility, rightsizing, query and pipeline optimization, and FinOps controls helped reduce unnecessary platform consumption.

Hours or Days → Minutes for Selected Analytical Questions

Governed conversational analytics and AI-assisted investigation reduced the time required to answer routine cross-dimensional business questions.

Higher Trust Across BI and AI

Certified data products, governed business semantics, lineage and embedded quality controls provided a consistent foundation for Power BI, conversational analytics and emerging AI applications.

From Reactive Reporting to Proactive Intelligence

The organization moved toward an analytical operating model where selected business and data signals could be identified, investigated and surfaced without relying entirely on users to continuously monitor dashboards.

Business Value Delivered

The enterprise had already modernized much of its data landscape on Microsoft Azure.

PalTech addressed the harder next-stage challenge: making those investments work together as a trusted foundation for enterprise intelligence.

By combining Azure Databricks, governed data products, embedded data quality and governance, enterprise semantics, Power BI, conversational analytics and

proactive intelligence, the organization evolved from:

Cloud-Modernized Data to Governed, AI-Ready Enterprise Intelligence. The resulting foundation supports continued expansion of trusted analytics, enterprise AI and proactive decision support across business divisions. Explore our Data & Analytics capabilities.

Technology Stack

Layer	Technologies / Capabilities
Cloud	Microsoft Azure
Data Platform	Azure Databricks, Delta Lake, ADLS Gen2
Data Engineering	Lakeflow Connect, Lakeflow Spark Declarative Pipelines, Lakeflow Jobs
Governance	Unity Catalog, Lineage, Security, Certification
Semantics	Unity Catalog Semantics, Metric Views, Governed Business Definitions
Analytics	Databricks SQL, Power BI
AI & Intelligence	Genie / Chat with Data, ML & AI, Analytical Agents
Optimization	Databricks System Tables, FinOps & Workload Optimization
Interoperability	APIs and MCP-enabled integration where appropriate

Contact Us



www.pal.tech



info@pal.tech