

## Case Study

# Modern Data Platform Powers Scalable Insights for Industrial Contractor



## Customer Overview

### Customer

The client is a full-service industrial contractor providing a wide range of construction, maintenance, and turnaround services to the chemical, petrochemical, power, automotive manufacturing, and refining industries across the United States.

### Profile

The industrial contractor serves various industries, from chemical and power to automotive and refinery, with construction, turnaround, and maintenance services.

### Industry

Engineering, Procurement and Construction (EPC)

### Services

- Cloud Data Modernization
- Microsoft Fabric Implementation
- Power BI Dashboard Development
- Performance Tuning and Optimization

## Business Need

The industrial construction firm was grappling with several limitations in its legacy, on-premises data architecture. The existing data mart was costly to maintain, required constant hands-on infrastructure management, and could not scale to meet the demands of a rapidly growing and geographically dispersed business.

With construction and maintenance projects running simultaneously across multiple regions, data volume, variety, and velocity continued to increase. The legacy environment, heavily reliant on manual data handling and ad-hoc reporting, struggled to provide timely insights and hindered executives' and project managers' ability to make data-driven decisions.

The company faced critical challenges such as:

- High infrastructure and maintenance costs for aging on-prem systems
- Fragmented data across project management, ERP, HR, and field tracking systems
- Limited ability to scale reporting across new projects or regions
- Dependence on IT teams for report creation, delaying insight delivery
- Lack of real-time visibility into cost overruns, labour allocation, and project performance

The contractor sought to modernize its data ecosystem to stay competitive and ensure operational excellence by migrating to a unified, cloud-native platform. The new solution needed to:

- **Scale dynamically** to support increasing data volumes from field operations, finance, safety, and HR systems
- **Reduce operational burden** by eliminating hardware maintenance and manual report consolidation
- **Leverage Microsoft Fabric's capabilities** for end-to-end data engineering, warehousing, and advanced analytics
- **Integrate seamlessly** with existing systems like Viewpoint (ERP), Procore, Microsoft SQL Server, and timekeeping platforms
- **Improve cost efficiency** through a pay-as-you-go model while reducing reliance on IT for daily reporting needs
- **Enable self-service analytics** through governed, secure Power BI dashboards for various roles across the organization

By adopting a modern Microsoft Fabric-based architecture, the industrial contractor aimed to create a scalable and agile decision-making foundation, empowering corporate and field teams with timely, accurate, and actionable data.

As part of a broader digital transformation journey, the company was also transitioning from Microsoft Dynamics AX to Dynamics 365 Finance and Operations. In parallel, they engaged Synoptek to modernize their analytics platform using Microsoft Fabric, ensuring that historical data, current workloads, and future needs were addressed in a unified architecture.

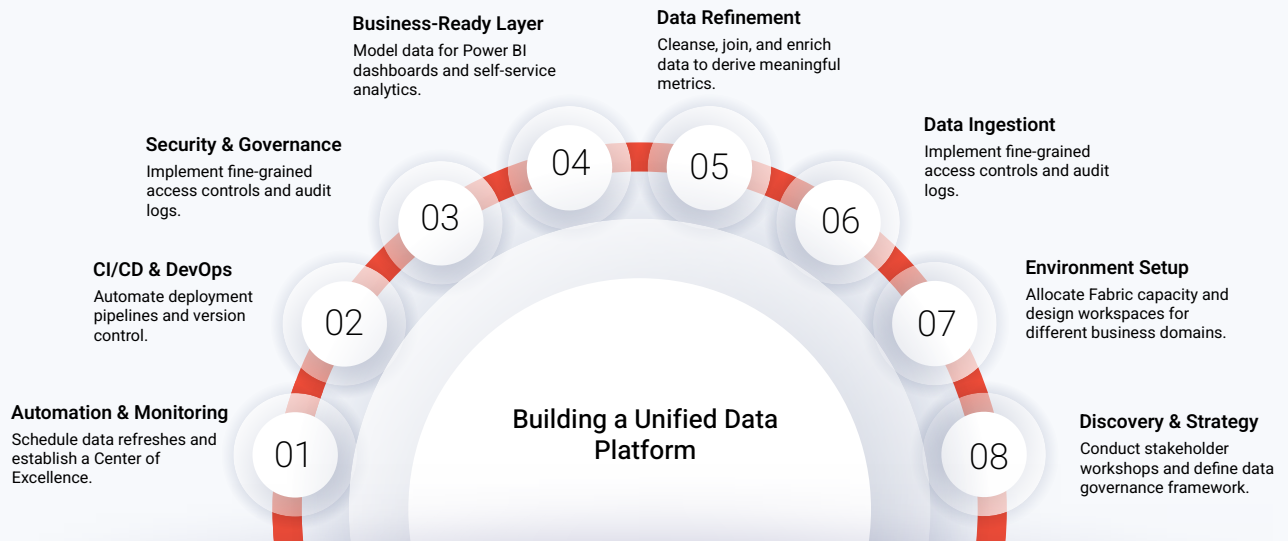
## Solution and Approach

Synoptek partnered with the client to architect and implement a modern data platform using Microsoft Fabric. The initiative was designed to run concurrently with the AX to Dynamics 365 upgrade, ensuring smooth historical data migration and eliminating legacy technical debt.

The engagement focused on unifying project, workforce, and financial data from legacy systems and Dynamics 365 FandO into a scalable, AI-ready platform, enabling real-time insights, reduced IT dependency, and a foundation for predictive analytics.

## Our Solution

The industrial contractor's transformation journey was planned and executed in four strategic phases, enabling a smooth migration from legacy systems to a scalable, modern data ecosystem powered by Microsoft Fabric and Power BI.



## Phase 1: Discovery and Data Strategy

Phase 1 focused on understanding the client's data landscape and aligning analytics efforts with business needs. This phase laid the foundation by identifying key use cases, data sources, and governance requirements.

- Conducted stakeholder workshops across finance, HR, field operations, and safety teams to understand reporting pain points and analytics needs.
- Identified priority use cases: project cost tracking, time and attendance reports, earned value analytics, labor productivity, and safety compliance.
- Mapped data sources (D365 FandO, D365 CRM, Enterprise Content Management System, TMS, EHS) and created a data ingestion blueprint.
- Defined governance framework including data ownership, retention policies, and row-level security requirements.

## Phase 2: Microsoft Fabric Implementation

The second phase focused on building the core data foundation by leveraging Microsoft Fabric's modern Lakehouse architecture. This phase was critical in transforming a fragmented data landscape into a unified, scalable, high-performance platform supporting real-time analytics and reporting.

### 1. Environment Provisioning and Workspace Setup

- **Microsoft Fabric Capacity Allocation:** Reserved appropriate Fabric capacity (F64 SKU) to handle large volumes of structured and semi-structured data generated from ongoing projects across hundreds of sites.
- **Workspace Design:** Created dedicated workspaces for business domains such as Finance, Operations, HR, and Safety. Each workspace followed best practices for data lineage, separation of duties, and governed access.

- **Lakehouse Creation:** Configured Lakehouse containers within Microsoft OneLake to store raw, curated, and gold-layer datasets aligned to a medallion architecture.

## 2. Data Ingestion and Integration (Bronze Layer)

A key priority was centralizing data from multiple on-prem and cloud-based systems and preparing it for analytics in a governed, performant manner.

### Ingestion from SQL Server and On-premises Systems

- Utilized Microsoft Fabric Data Pipelines to connect with SQL Server databases hosting timekeeping, cost control, procurement, and payroll data.
- Developed parameterized pipeline templates to ingest full loads initially, followed by incremental loads based on timestamp or change tracking.

### Delta Lake Tables and File Formats

- All ingested data landed as Delta tables in OneLake to support ACID transactions, time travel, schema enforcement, and concurrent reads/writes.
- Partitioned data by project ID and fiscal period to improve query performance for large fact tables like labour hours, cost journals, and safety incidents.

## 3. Data Refinement (Silver Layer)

- Used Fabric Pipeline in Fabric to cleanse, join, and enrich data. Examples include:
  - Deriving earned value metrics from actual vs. planned cost/labour
  - Mapping project hierarchy (company > region > project > job)
  - Calculating payroll liability and labour burden
- Applied data validation rules to identify anomalies in timesheets, duplicate job codes, or unapproved cost entries.
- Created gold-layer curated tables used as a source for Power BI models.

## 4. Business-ready Layer (Gold Layer)

The Gold Layer (Presentation Zone) produced final datasets for business consumption:

- Modelled conformed dimensions and fact tables for Power BI dashboards and self-service analytics.
- Applied row-level security and usage metadata to ensure trusted and role-specific access.
- Fed Gold datasets directly into Power BI semantic models published to executive, HR, operations, and finance dashboards.

## 5. Security, Auditing and Governance Setup

- Defined data domains and implemented fine-grained access using Microsoft Entra ID security groups and Power BI workspace roles.
- Applied row-level security (RLS) to restrict visibility based on user roles — executives could view enterprise-wide reports, while project managers saw only their assigned jobs.
- Logged all ingestion and transformation activities for auditability, ensuring transparency and traceability across the pipeline.

## 6. CI/CD and DevOps Enablement

- Used deployment pipelines within Microsoft Fabric for promoting dataflows, semantic models, and reports across Dev -> QA -> Prod environments.
- Integrated with Azure DevOps Git repositories to version-control notebooks, data models, and Fabric artifacts.
- Automated alerting and logging for pipeline failures, row count mismatches, or schema drift issues.

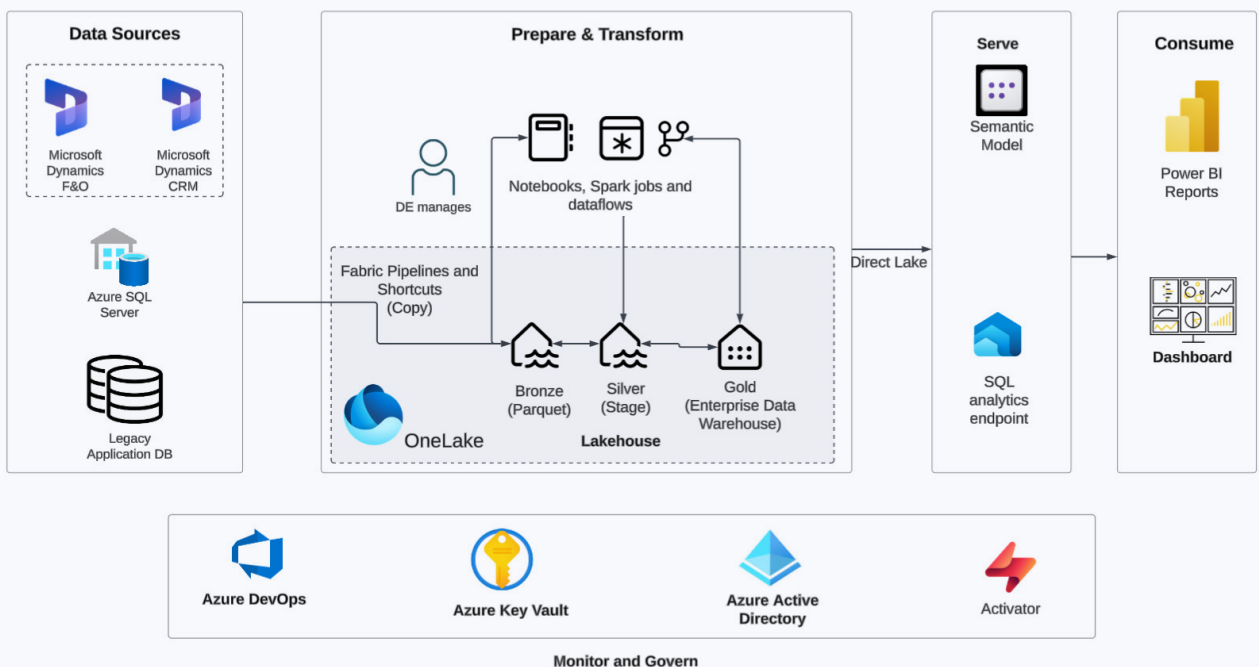
## Phase 3: Semantic Modeling and Power BI Reporting

- Developed reusable semantic models in Power BI for Finance, HR, and Operations domains.
- Created curated Power BI datasets and published them to centralized workspaces for controlled access.
- Built executive dashboards (cost trends, profitability by project, project progress) and operational reports (daily labour logs, safety incidents, earned value analysis).
- Applied role-based access using Power BI Row-Level Security to ensure field managers, regional heads, and finance leaders see only relevant data.

## Phase 4: Automation and Monitoring

- Scheduled incremental data refresh pipelines using Fabric Data Pipelines to ensure near real-time reporting without performance degradation.
- Integrated Active Directory for single sign-on and implemented access monitoring through Power BI admin API.
- Established a Center of Excellence (CoE) to train business users on Power BI self-service capabilities, reducing IT dependency.

## Modernizing Data Platform with MS Fabric – Reference Architecture



## Technology Stack

- Microsoft Fabric (Delta Lakehouse, Data Pipelines, Dataflows Gen2)
- Power BI (Service, Desktop, Semantic Models, RLS)
- Azure Data Factory (legacy ingestion)
- Azure Logic App
- Azure Automation
- Delta Lake, OneLake, Apache Spark Notebooks

## Business Benefits

Synoptek's solution delivered measurable improvements in efficiency, data accessibility, and decision-making capabilities for the contractor. Key business benefits included:

- Enabled 70% faster report generation time, reducing project status cycle from days to minutes
- Built a single source of truth across project cost, labor, and material management
- Empowered business users with self-service BI and real-time project dashboards
- Improved compliance with secure, auditable data access across roles
- Optimized costs by reducing legacy licensing and maintenance expenses
- Delivered a future-ready platform scalable for predictive analytics and AI integration



**40%** performance improvement



**Data consistency and efficiency**



**70%** faster report generation

## About Synoptek

Synoptek delivers accelerated business results through advisory-led, transformative full-life-cycle systems integration and managed services. We partner with organizations worldwide to help them navigate the ever-changing business and technology landscape, build solid foundations for their business, and achieve their business goals.