



# 80% Faster to GenAI with DataOS

## A Fortune 500 Case Study



# Executive Summary

## The Challenge

A Fortune 500 global manufacturer struggled to unlock the value of operational and product performance data across more than 150 million devices worldwide. Despite substantial efforts and years of initiatives, the company faced lengthy time-to-insight cycles that prevented them from achieving critical outcomes like predictive maintenance, performance improvements and quality optimization. They were also limited in their ability to capitalize on new revenue opportunities from their vast device data.

## The Solution

By implementing DataOS, accelerated their device intelligence and GenAI capabilities by 80%, delivering key predictive maintenance and quality optimization initiatives several quarters ahead of schedule, and establishing a unified platform for data-driven operational excellence and new monetization opportunities.

## The Results

- 80% reduction in time-to-market for device intelligence and GenAI experimentation
- Multiple quarter acceleration of strategic manufacturing initiatives
- Demonstrable gains in engineering velocity and operational efficiency
- Global-scale implementation across geographies and multiple product ecosystems
- Unlocked new data monetization opportunities beyond traditional operations

***"With our unified data platform, developers can shift from 'just producing tables' to building reliable, discoverable data products with clear ownership and governance—all while supporting both real-time and batch processing needs."***

— Sr. Director of Data & Analytics

# Manufacturing Data: Collected, but Not Connected

Manufacturing generates unprecedented volumes of operational data—reports show that a typical factory creates one terabyte daily, with 90% going unused. DataRPM research shows 85% of industrial companies let data from trillions of data points go completely unused.

## The Manufacturing Intelligence Challenge

This disconnect stems from three fundamental challenges that prevent data from becoming actionable intelligence:

- **Data Volume and Velocity:** Production lines generate millions of sensor readings per hour across temperature, pressure, vibration, and quality metrics. Smart sensors and connected equipment create exponential data growth while machine learning models demand continuous feeds to maintain accuracy.
- **Semantic Ambiguity:** Manufacturing plants use inconsistent naming conventions and data formats across different systems, equipment, and locations. A "temperature sensor" might be labeled as "temp\_01," "T-sensor-A," or "thermal\_reading" in different systems, making it difficult to automatically correlate and analyze related data points.
- **Temporal Misalignment:** Manufacturing data streams operate on different time scales and frequencies - quality sensors might report every second, maintenance logs weekly, and production schedules daily.
- **Integration Complexity:** Manufacturing environments contain heterogeneous data formats spanning structured databases, time-series sensor data, video feeds, maintenance logs, and operator reports. Integration becomes exponentially complex across decades-old legacy systems, modern cloud.

(1) <https://www.industryweek.com/technology-and-iiot/article/21249974/drowning-in-raw-data-lean-principles-can-help>

- **Organizational Barriers:** Data silos prevent cross-functional insights between production, quality, maintenance, and supply chain teams. Skills gaps in data engineering within traditional manufacturing workforces, combined with disconnected operational and information technology systems, limit progress. Budget constraints create a cycle where manufacturers need ROI proof before platform investment but can't demonstrate ROI without the platforms.

While manufacturers excel at generating operational data, most struggle to transform this information into actionable intelligence that drives measurable outcomes like reduced downtime, improved quality, or optimized inventory levels.

# Company Profile: A Global Technology Leader

## Company Overview

This Fortune 500 technology manufacturer operates globally with large-scale consumer and commercial product lines. The company maintains substantial data engineering and analytics teams to support its multinational operations and market-leading position in hardware device manufacturing.

- **Scale:** Multinational operations with global market presence
- **Product Portfolio:** Large-scale consumer and commercial product lines
- **Data Organization:** Substantial data engineering and analytics teams

The company faced competitive market pressures and recognized that unlocking product intelligence was critical for:

- Predictive maintenance and quality optimization
- Operational efficiency and cost management
- Enhanced customer experience and satisfaction
- New data-driven service capabilities

# The Data Challenge: Long Timelines, Limited Results

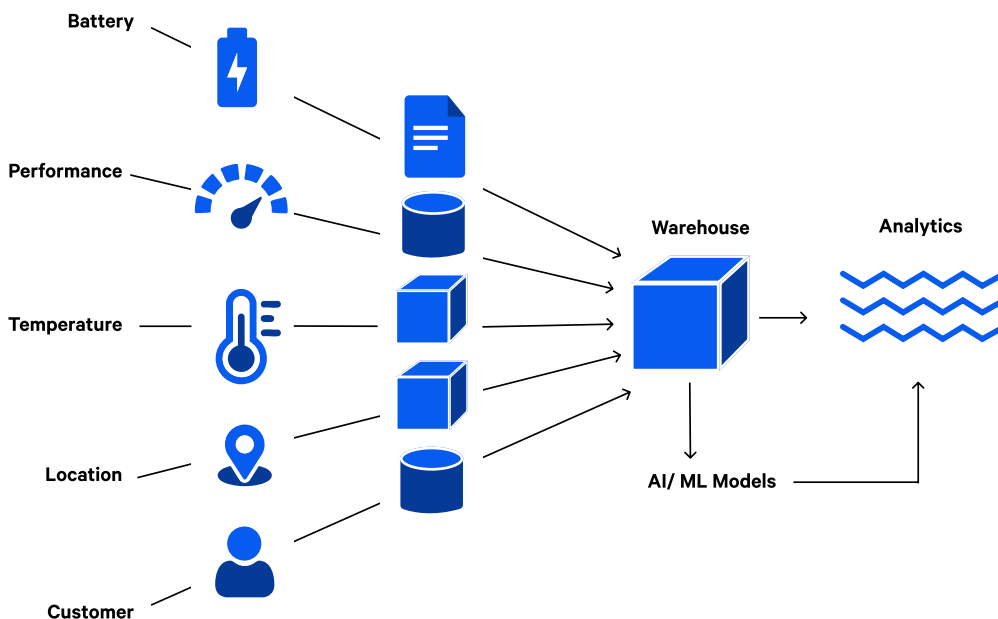
## Previous Attempts and Limitations

Despite recognizing the untapped potential in device telemetry data, previous attempts with various technology partners consistently fell short of expectations, creating a cycle of skepticism that made stakeholders question whether a comprehensive solution could be implemented in less than 24 months. The challenges were both technical and organizational, rooted in the fundamental complexity of unifying vast streams of operational data into actionable business intelligence.

## Initiative History

- Multiple years of data platform initiatives with various technology partners
- Multiple attempts to create unified device intelligence capabilities
- 2-3+ year timelines for each major data project
- Increasing skepticism from leadership about achievable outcomes

## Integration Complexity of Device Telemetry Data



## Underlying Challenges

The company's data challenges stemmed from both technical limitations and organizational barriers that had accumulated over years of incremental solutions. While the organization has substantial data engineering capabilities and had invested significantly in various platforms, the fundamental issues that prevented successful device intelligence could be put into two categories that reinforced each other in a cycle of complexity and delay.

### Technical Challenges

- **Data Silos:** Device telemetry scattered across disconnected systems
- **Integration Complexity:** 9-month+ cycles to connect new data sources
- **Tool Sprawl:** Multiple point solutions creating maintenance overhead
- **Scalability Limitations:** Inability to handle real-time processing at device scale

### Organizational Challenges

- **Skill Set Gaps:** Traditional data engineering approaches inadequate for current requirements
- **Governance Deficits:** Lack of data lineage and quality controls
- **Collaboration Barriers:** Disconnected workflows between data teams and business units

## The Cost of Delayed Innovation

Companies that don't effectively tap into their data for innovation end up playing catch-up in costly ways. They will be fixing things more often after they break instead of seeing problems coming. Warranty claims become a resource drain because of inefficient processing and unexpected costs. Customers get frustrated because issues aren't caught and resolved before they become real problems. And while competitors move faster with data-driven insights, these companies find themselves stuck in slower innovation cycles, always a step behind in the market.

# The DataOS<sup>®</sup> Approach

## Outcome-driven Product vs. Project Thinking

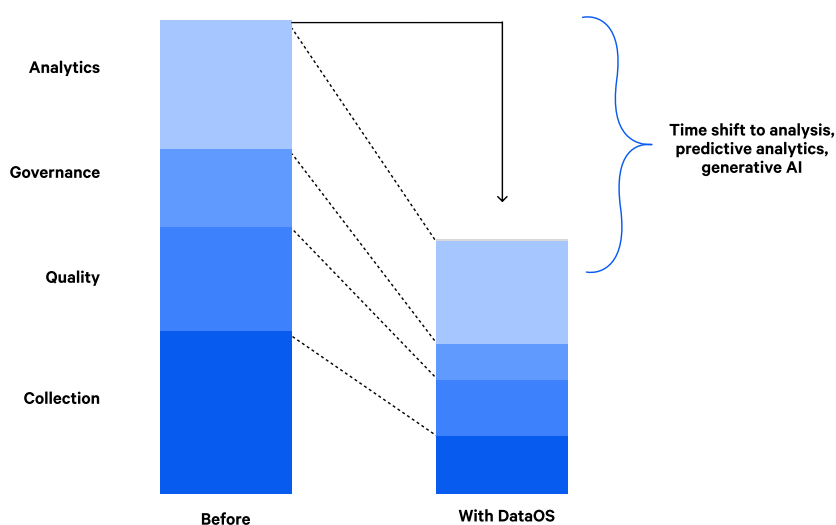
DataOS offered a different approach to the company's data challenges. Rather than requiring infrastructure replacement or lengthy implementation cycles, the platform was designed to integrate with existing systems and demonstrate value through rapid proof-of-value deployments—offering the potential to deliver results in weeks instead of the typical quarters-long timelines.

DataOS represents a fundamental shift from traditional data warehousing to a converged data management platform that delivers:

### Converged, Not Centralized Architecture

- **Maximize, not rip and replace:** Layer into existing infrastructure, at any maturity level, without disrupting current operations.
- **4-6 week production-ready deployments** vs. typical 12-18 month traditional timelines.
- **Governance by design:** Deploy enterprise-wide access controls and usage monitoring with a few lines of code.

## Shifting from Processing to Impact: Before and After DataOS Implementation



## Full Data Product Lifecycle Management

- **Industry-leading approach** to create, build, use and re-use data products based on the latest software engineering principles.
- **Built-in data quality and validation** with automated monitoring and alerting.

## AI-Ready Data Activation

- **90% faster data activation** with production-grade, business and AI-ready data based on data products.
- **Built-in semantic layer** enabling natural language interfaces and rapid AI agent deployment
- **Multiple consumption interfaces** (SQL, GraphQL, REST APIs) for maximum flexibility.

## The DataOS Advantage

### Key Differentiators

- 1. Time Compression:** An average of 80% reduction in development cycles through pre-built accelerators.
- 2. Productivity Enhancement:** Up to 50% improvement in data usage and engineering efficiency.
- 3. Platform Consolidation:** Single solution can replace multiple point tools.
- 4. Developer Experience:** Modern software engineering practices applied to data.

# Implementation Journey: From Skepticism to Success

## Phase 1: Rapid Prototyping (Week 1)

### The Five-Day Breakthrough


- **Friday:** Received nearly 50 sample telemetry data files
- **Weekend:** DataOS team analyzed data structure and requirements
- **Tuesday:** Working demonstration environment ready
- **Following Week:** Comprehensive stakeholder demonstration

### Demonstration Components

1. **Device360 Data Product:** Unified view of device performance metrics
2. **Natural Language Interface:** LLM-powered querying capabilities
3. **Predictive Dashboard:** Real-time risk assessment and alerting

**Stakeholder Response:** The rapid prototype fundamentally shifted the conversation from "Can this work?" to "What can we achieve?" Leadership discussion immediately moved to immediate deployment next steps and future opportunities.

### The Device360 Data Product powered by DataOS: Unifying Trusted Data



## Phase 2: Proof of Value (Weeks 2-5)

### Environment Setup

- Full DataOS deployment within company's cloud environment
- Integration with existing data sources and infrastructure
- Collaboration with internal team

### Key Achievements

- **4-week completion** of originally 6-week timeline
- **Successful integration** with production telemetry streams
- **Stakeholder alignment** across technical and business teams

### Technical Breakthroughs

- **GraphQL Integration:** Eliminated need for separate API gateway
- **Real-time Processing:** Demonstrated capability for live device monitoring
- **Governance Integration:** Automated data lineage and quality controls

## Phase 3: Strategic Alignment (Weeks 6-12)

**Executive Presentation Results:** A 30-minute presentation to senior leadership resulted in:

- Immediate enterprise-wide adoption decision
- Platform expansion vision across device categories
- OEM opportunity exploration for customer offerings
- Design partnership agreement for future capabilities

### Quantified Benefits

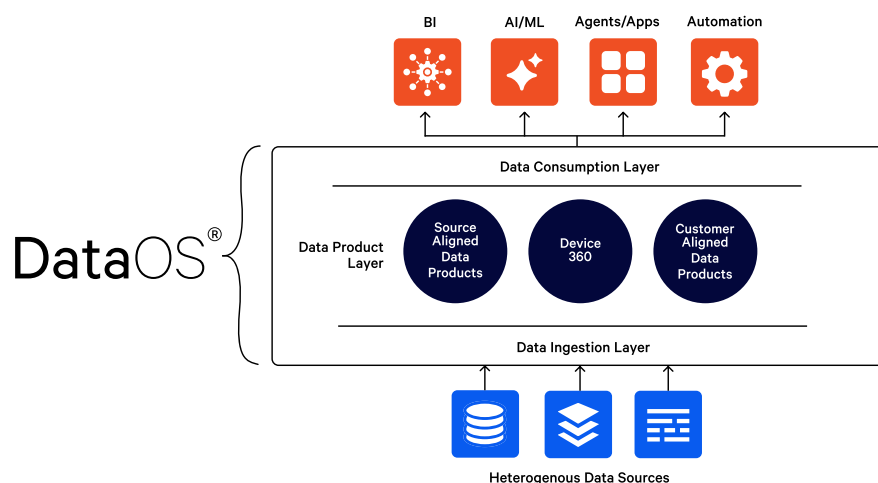
- **Time Savings:** 3-4 quarter acceleration of strategic initiatives
- **Cost Reduction:** Consolidated tool licensing and maintenance
- **Revenue Enablement:** New data-driven service offerings
- **Risk Mitigation:** Improved device reliability and customer satisfaction

# Technical Architecture and Capabilities

DataOS delivered on its promise through three core architectural innovations that fundamentally changed how the company approached data.

## Unified Data Platform Architecture

Rather than forcing the company to rip and replace existing systems, DataOS created a unified layer that connected everything.



### Data Ingestion Layer

- Real-time telemetry stream processing
- Batch historical data integration
- Multi-format data parsing and validation
- Automatic schema evolution and management

### Consumption Layer

- GraphQL APIs for flexible querying
- REST endpoints for application integration
- SQL interface for analytical workloads
- Real-time streaming for monitoring applications

### Data Product Layer

- YAML-defined data contracts and interfaces
- Version-controlled data product definitions
- Built-in data quality and monitoring
- Automated testing and validation pipelines

## Key Technical Innovations

DataOS delivers enterprise-grade data infrastructure through a combination of modern software engineering practices and proven data technologies. By treating data as code and integrating semantic layers directly into the architecture, DataOS creates a system that bridges the gap between technical implementation and business value, while maintaining scale and reliability.

**Data as Code Approach:** All data products defined in version-controlled YAML configurations, enabling:

- Reproducible deployments across environments
- Collaborative development using standard software practices
- Automated testing and quality assurance
- Rollback capabilities for production safety

**Semantic Layer Integration:** Built-in semantic modeling provides:

- Business-friendly data definitions accessible to non-technical users
- Consistent metrics across all consumption interfaces
- LLM-ready data descriptions for natural language querying
- Automated documentation generation and maintenance

### Consumption Layer

- GraphQL APIs for flexible querying
- REST endpoints for application integration
- SQL interface for analytical workloads
- Real-time streaming for monitoring applications

### Demonstrated Capabilities

- 150M+ devices telemetry processing
- Sub-second latency for real-time queries
- Petabyte-scale data storage and processing

# Business Outcomes and Impact

## Organizational Transformation

The transformation unfolded quickly across several dimensions. Team productivity metrics tell the story: 50% reduction in repetitive data engineering tasks, multiple quarter acceleration of strategic initiatives, and 80% faster new use case development.

*"DataOS fundamentally changed how we think about data. Instead of building one-off solutions, we're now creating reusable data products that serve multiple business needs simultaneously."*

## Strategic Business Enablement

### The DataOS platform unlocked four critical capabilities

1. Proactive device failure prediction through predictive maintenance
2. Data-driven warranty optimization for improved cost management
3. Real-time device performance monitoring for enhanced customer experience
4. Data-informed product development insights for faster innovation

These weren't just technical achievements—they represented new revenue streams, cost avoidance through reduced warranty claims, and accelerated innovation cycles.

### Operationalizing Intelligence: From Device Data to New Revenue Streams

Comprehensive device intelligence doesn't just optimize current operations—it creates entirely new revenue opportunities. **Real-world device performance data can reveal insights that are impossible to accurately capture in a lab.** For example, instead of expensive controlled battery degradation tests, actual usage patterns can be analyzed across millions of devices to understand exactly how different environmental conditions, usage behaviors, and device configurations impact battery life. This intelligence becomes a strategic asset: informing product development, enabling premium warranty offerings based on actual risk profiles, and creating new services for partners to optimize their own products based on proven performance analytics

The success came down to leadership alignment with clear executive sponsorship, a technical approach that prioritized rapid prototyping over extensive planning and change management that invested in training 40+ engineers on new approaches while fostering a cultural shift from project to product thinking.

## Speed as a Strategic Advantage

The lessons extend beyond manufacturing. Speed has become a strategic advantage—organizations that compress data innovation cycles by 80% gain significant advantages in market responsiveness and operational efficiency. Product thinking scales where project thinking fails, providing economies of scale through shared infrastructure and faster innovation through reusable components. Most importantly, data products have become business assets with clear ownership, user experience focus, and direct connection to business outcomes.

## Map Your Route to Accelerated Outcomes

**For organizations beginning their journey:** start with a 4-6 week, low-risk Proof of Value (PoV) using DataOS and applying a real-world use case focused on business outcomes. For more established teams: we can help you consider how DataOS could consolidate your current platform complexity—evaluate where time compression could accelerate your existing initiatives and identify opportunities to transform current bottlenecks into competitive advantages.

The technology and practices demonstrated in this case study are available today to maximize your existing data stack without disruption. DataOS layers into your current infrastructure to unlock significant additional value from investments you've already made. A DataOS Proof of Value can quickly validate whether 80% time compression is achievable for your specific data challenges and business objectives—delivering working results in weeks, not quarters.

**Ready to explore what DataOS could accomplish with your data?** Contact us to discuss a customized Proof of Value that addresses your organization's unique challenges and demonstrates measurable time compression in your environment.

# About The Modern Data Company

The Modern Data Company is redefining enterprise data management for the AI era, transforming data from a technical challenge into an organizations' most powerful business asset. The company's flagship platform, DataOS, serves as the essential business and AI-ready data layer for any data stack to solve an organization's toughest data challenges. DataOS transforms fragmented, ungoverned data into AI-ready data products in just weeks, delivering trusted, context-aware data that makes AI systems smarter while providing native integration support including built-in APIs and GraphQL for seamless connectivity to existing tools, workflows, and LLMs.

Fortune 1000+ enterprises using DataOS are accelerating their AI adoption by up to 90% while reducing total cost of data ownership by up to 50%, delivering dramatic increases in business agility. The company's rapidly expanding customer base includes global category leaders across multiple industries who trust DataOS to power their AI and business transformation.

## Learn More

- [www.themoderndatacompany.com](http://www.themoderndatacompany.com)
- [info@tmdc.io](mailto:info@tmdc.io)

*This case study is based on a real-world implementation with a Fortune 500 technology manufacturer. Company and individual names have been anonymized to protect confidential business information. Results and timelines are factual and representative of typical DataOS implementations.*

## Document Information:

- Version: 1.0
- Date: May 6, 2026
- Classification: Public Distribution