CASE STUDY

# PL Developments Transforms Data Architecture with Dremio: From Complex ETL to Real-Time Analytics and AI-Ready Infrastructure

## At a Glance

### The Customer

### Challenge

PLD's legacy data warehouse and complex ETL created severe challenges in performance, scalability, and governance. The expensive, proprietary architecture lacked the real-time access and consistency needed for modern agentic AI.

### Solution

PLD implemented Dremio's Agentic Lakehouse with unified data access, enabling secure, governed queries across Postgres and their Iceberg lake environment. They extended this architecture with enterprise-grade catalog controls, AI-ready connectivity, and flexible views that abstract underlying system complexity.

### Results

- 6 seconds complex analytical query time reduced.
- 30% of ERP system resource consumption from legacy ETL processes was eliminated.
- 5,000+ product SKUs unified access point across multiple data sources
- Sub second query performance enabled LLM-powered analytics and automation.

## The Customer

PL Developments (PLD) is a leading manufacturer, packager, and distributor of over-the-counter (OTC) pharmaceutical products and consumer healthcare goods, producing approximately 5,000 items for major national retailers. If consumers purchase generic ibuprofen at Walmart or CVS, it was likely manufactured by PLD. With three decades of experience, PLD has grown to serve almost every retailer in the nation and many globally.

## The Challenge

As a high-volume pharmaceutical manufacturer, PL Developments (PLD) relies on timely, accurate insights into production, component availability, and order fulfillment. However, their legacy Oracle-based environment was slow, brittle, and unable to keep pace. The system had evolved into a closed architecture with over ten overlapping data layers and hundreds of ETL jobs consuming nearly 30% of ERP resources. Simple analytics were complex, requiring heavy transformations and fragile pipelines. This architecture lacked the foundation for real-time analytics or modern AI. Oracle BI and the underlying warehouse lacked the semantic consistency and fine-grained security controls needed to safely broaden data access beyond a small technical team. Attempts to introduce self-service analytics, such as Alteryx, immediately highlighted security gaps and the risk of overexposing sensitive financial data.

Shifting to a modern stack—Kafka CDC, Postgres, and Apache Iceberg—brought short-term performance gains but created new operational burdens like metadata sprawl and failed compactions. Open-source query engines also lacked the required role-based security and governance.

The core issue became clear: PLD lacked a unified, governed access layer. This layer was necessary to provide consistent semantics, enforce granular security, and deliver high-performance analytics across all systems. Without this architectural foundation, PLD could not future-proof its data strategy or confidently expand data access to business users, automated workflows, or emerging AI-driven decision-making.

## The Solution

PL Developments modernized its analytics architecture by adopting Dremio as a unified data access layer, creating a centralized semantic foundation and a single point of entry for all analytics and insights. Dremio delivered the fine-grained security controls that their previous Alteryx environment lacked,

> "Dremio became our Swiss army knife – it solved our immediate lakehouse performance challenges, gave us a federated SQL engine to query across our data warehouse and lakehouse, and provided the single point of access we needed for AI and self-service analytics. It was the key that made everything work."
>
> — Andrew Balaban, Head of IT & IS, PL Developments

while providing a consistent SQL interface across operational systems and the lakehouse, greatly simplifying how teams interact with data.

The organization embraced a "respect the line" architectural approach, keeping fast-changing operational tables in Postgres while maintaining large analytical datasets in their Iceberg lakehouse. Dremio's Intelligent Query Engine seamlessly connected these environments, allowing users to join warehouse and lakehouse data in real time while ensuring each system continued to operate within its optimal performance zone.

With secure, unified access established, PL Developments expanded into AI-driven analytics by integrating tools such as Cursor AI directly with Dremio's MCP. Using Dremio's Open Catalog, the team implemented enterprise-grade governance with schema-level and table-level permissions that now enforce consistent access policies across every AI and BI platform. Dremio's semantic layer further insulated users and AI agents from backend complexity, exposing stable data models through views that remain unaffected by changes to the underlying architecture.

To support large-scale Iceberg operations, the team also adopted Dremio's autonomous table management capabilities. Previously, Iceberg snapshots accumulated rapidly, causing metadata bloat, heavy compaction workloads, and unpredictable query slowdowns. Dremio fully automated compaction, clustering, file-size optimization, garbage collection, and metadata cleanup, ensuring consistently performant tables and eliminating the operational burden of manual maintenance. This foundation now enables PL Developments to run lakehouse analytics and AI workloads reliably and efficiently at scale.

## Results

These advancements created the foundation for a new generation of AI-enabled business processes. PL Developments deployed LLM applications capable of generating customer communications in under a minute by automatically pulling together order histories, cancellation patterns, and production schedules. Quality teams can now query FDA databases for vendor warning letters in real time and immediately place affected products on hold. Legal teams similarly benefit from instant access to ERP data during mergers and acquisitions, enabling faster and more thorough evaluations of potential targets.

The modernization effort also eliminated the operational strain caused by legacy ETL pipelines. Historically, these pipelines consumed nearly 30% of ERP system resources, slowing critical operations and creating unnecessary bottlenecks. By retiring them and shifting to Dremio's virtualized, zero-copy architecture, PL Developments reduced system load, accelerated core processes, and freed capacity across its operational landscape.

Dremio has also become the trusted foundation for all AI access within the organization. Rather than exposing databases directly to LLMs or agentic systems, PL Developments uses Dremio as a secure intermediary that enforces governance, provides consistent semantics, and ensures that AI tools interact only with validated, policy-compliant data. With a unified, secure access layer in place, the company is now positioned to scale its AI initiatives with confidence. Dremio gives PL Developments the stability, governance, and performance needed to pursue modern analytics and AI innovation without compromising control or data integrity.

**ABOUT DREMIO**

Dremio is the intelligent lakehouse platform for the business, serving hundreds of global enterprises, including Maersk, Amazon, Regeneron, NetApp, and S&P Global. Based on open-source technologies like Apache Iceberg and Apache Arrow, Dremio provides an open lakehouse architecture enabling the fastest time to insight and platform flexibility at a fraction of the cost.
Learn more at www.dremio.com.