

CASE STUDY

Tanobel Saves 3,000+ Hours Monthly with Dremio's Intelligent Lakehouse Platform

At a Glance

The Customer



Challenge

Tanobel's data landscape was marked by performance bottlenecks, data silos, and manual reporting workflows.

Solution

Tanobel built an architecture that combined Dremio Reflections for accelerating query performance with Apache Iceberg tables for modern data lakehouse functionality.

Results

- **Reporting speed:** Report processing times reduced from 15 minutes to under 1 second
- **Time savings:** 3,000+ employee hours saved monthly on report processing
- **Development efficiency:** Report creation time reduced from 2 weeks to 1 days
- **User adoption:** Supporting 100+ daily users accessing 500+ reports

The Customer:

Tanobel is a publicly traded beverage company in Indonesia, known primarily for its bottled drinking water products. With plans to expand into a broader range of beverages and food products, Tanobel's operations span manufacturing, supply chain, and sales across multiple sites. Managing increasing volumes of data and complex reporting needs, Tanobel recognized the need for a more scalable and efficient data platform to support decision-making and operational agility.

The Challenge:

Before adopting Dremio, Tanobel's data landscape was marked by performance bottlenecks, data silos, and manual reporting workflows. Reports took up to 15 minutes to process, and in some cases, data preparation spanned two weeks, delaying access to critical insights.

Additionally, real-time ERP workloads were frequently disrupted by reporting queries, which led to delays in data input and business operations. "We had to segregate the data just to avoid locking up the ERP," explained Tanaka Murinata, CIO at Tanobel. "It wasn't scalable, and it impacted productivity."

Data inconsistency also plagued the organization. Without a unified system, different teams often relied on outdated or mismatched reports, leading to conflicting interpretations and internal disputes. "It was difficult to explain why numbers differed—different cut-off times led to different results," Tanaka said.

The Solution:

To address these challenges, Tanobel implemented Dremio's Intelligent Lakehouse Platform. The initiative began with Dremio Community Edition (CE) on a single server, eventually scaling to Enterprise Edition with a full production-grade infrastructure built internally by Tanobel's IT team.

Tanobel built an architecture that combined Dremio Reflections for accelerating query performance with Apache Iceberg tables for modern data lakehouse functionality. Data from ERP systems is ingested on a scheduled basis—ranging from every 10 minutes to once daily—depending on business needs.

Tanobel runs Dremio on an infrastructure specifically optimized for its requirements, employing the Red Hat OpenShift Container Platform. For data storage, they utilize Openshift Data Foundation (ODF) configured with SSDs for block and object storage, ensuring a high-performance base for Dremio. User reporting is handled via the Tanobel Intelligence Platform, a web-enabled system offering straightforward report access and scheduled delivery through email and messaging.

By integrating Dremio into their data ecosystem, Tanobel significantly enhanced their analytics capabilities across the organization. The platform enabled near real-time access to data from multiple databases, allowing users to perform complex queries and analysis without being constrained by the limitations of traditional systems. With Dremio's intuitive SQL interface and support for JDBC, Tanobel empowered departments with self-service reporting and analytics, reducing reliance on IT and accelerating decision-making.

Additionally, the move to Dremio Enterprise Edition introduced centralized data governance, giving the IT team the ability to define role-based access controls and ensure that sensitive information was only accessible to the appropriate users. This not only improved security but also helped maintain consistency and compliance across the organization. Dremio's semantic layer and support for

reusable datasets further enabled scalable data sharing, allowing curated data products to be built once and reused across teams—streamlining collaboration and driving analytical efficiency.

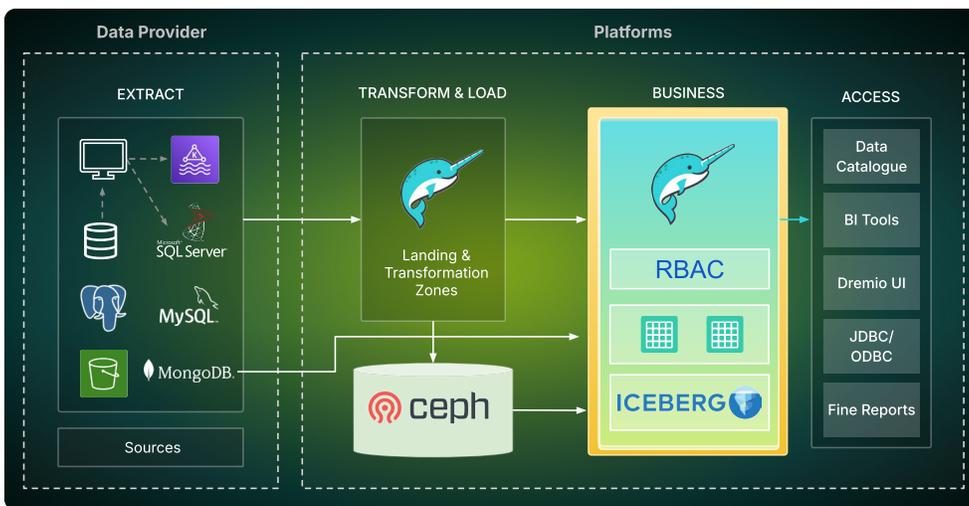
Results:

Tanobel's transformation has been remarkable. Query times for reports that once took up to 15 minutes now complete in under 2 seconds, while data preparation time has dropped from two weeks to 1–2 days. In some cases, minor report adjustments now take only hours.

Through Dremio, Tanobel not only standardized data access and removed the burden of manual data wrangling but also significantly boosted data utilization. Business users are now empowered with easily accessible reports, enabling them to monitor data continuously throughout the day instead of just occasionally. This shift, combined with the over 3,000 monthly man-hours saved from eliminating manual tasks, allows the organization to focus more effectively on strategic, data-driven initiatives.

The platform supports 500+ reports daily accessed by over 100 active users, and has enabled new reporting use cases across supply chain, manufacturing, and finance. Departments like supply chain can now react to

Tanobel Solution: Dremio On-Premises Lakehouse



- Data Volume, Variety, and Velocity**
~1TB of parquet data as 100 Iceberg tables
- Query Execution Time**
 - Query times for reports that once took up to 15 minutes now complete in under 2 seconds,
- Scalability**
 - ~Horizontally scalable
 - ~Consistent performance at scale

Over 100 active users

1TB of parquet data as 100 Iceberg tables

2000 jobs per day, 4-6 GB per query

sales spikes in near real time, adjusting production and inventory faster than ever. During holidays, for example, sales for specific products can suddenly increase—but with instant visibility, teams can quickly adapt.

The adoption of Dremio has also driven major cultural shifts. “Now we think in terms of reusable datasets instead of isolated tables,” said Tanaka. “We can combine and share across departments without bottlenecks.”

Governance has improved significantly. With Dremio, Tanobel can define user roles and control access at a granular level—critical in maintaining internal standards and preparing for Indonesia’s data protection regulations (GDPL).

Conclusion:

Tanobel is actively implementing AI and ML initiatives, building pilot projects using LLMs and vector databases to interpret summarized data automatically. Dremio plays a key role in this strategy by providing fast access to curated datasets that can feed downstream AI models. As the company grows, so will its data needs. Dremio provides a foundation not just for current performance, but also for Tanobel’s vision of scalable, accessible, and intelligent data infrastructure.

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.

Dremio and the Narwhal logo are registered trademarks or trademarks of Dremio, Inc. in the United States and other countries. Other brand names mentioned herein are for identification purposes only and may be trademarks of their respective holder(s). © 2025 Dremio, Inc. All rights reserved.