CASE STUDY

World Bank Treasury Transforms Capital Markets Trading with AI-Powered Lakehouse

At a Glance

The Customer

WORLD BANK GROUP

Challenge

The World Bank Group is a massive multilateral development institution dedicated to reducing poverty and advancing prosperity across 189 countries, managing an annual portfolio of over $100 billion. Its mission, providing critical financial services and policy advice, relies on securely operationalizing data for evidence-based decision-making.

Solution

The World Bank Treasury modernized its financial data infrastructure with Dremio, implementing the “Finance One Lake” strategy to unify data from 80+ systems using a zero-ETL, schema-on-read architecture. This platform provides a consistent semantic layer for analytics, powers an AI-driven trade automation system, and ensures robust governance through open lakehouse principles.

Results

- 95%+ accuracy in AI-driven trade data extraction from term sheets
- 6-8 hours reduced to 15 minutes per trade processing time
- 80+ data products onboarded across Treasury and Finance operations
- 70% of finance data consolidated into unified lakehouse platform
- Single source of truth established across 189 member countries’ financial operations

The Customer

The World Bank Group is one of the world’s largest multilateral development institutions, working across 189 countries to reduce poverty and advance shared prosperity. Its Treasury division manages portfolio investment, trading operations, and financial services that support approximately 80 developing countries. With a global portfolio exceeding $100 billion, the Treasury depends on timely, reliable, and secure data to guide strategic decision-making, monitor global markets, and deliver financial services with consistency and integrity. As data volumes and complexity continued to grow, the organization recognized the need for a unified, scalable platform capable of powering modern analytics and emerging AI initiatives.

The Challenge

For years, the Treasury operated within a fragmented data landscape spanning more than 80 treasury and finance systems. This environment routinely produced conflicting financial metrics, undermining confidence in reports delivered to senior leadership and exposing longstanding data quality challenges. Without a unified architecture to reconcile disparate sources, teams struggled to maintain a single source of truth for mission-critical insights.

Manual workflows further limited operational efficiency. Traders spent six to eight hours reviewing term sheets, extracting key fields, and entering them into trading platforms—an error-prone process that slowed execution and diverted time away from high-value market analysis. Traditional warehouse-centric approaches added additional burden through redundant data copies, ongoing pipeline maintenance, and delayed access to new attributes or data sources.

Dependence on IT for reporting and data modeling also restricted agility. Analysts frequently waited for technical teams to prepare datasets or build new views, creating bottlenecks that hindered timely decision-making. As global financial markets accelerated, the lack of consistent, accessible, and governed data became a significant obstacle to the Treasury’s ability to operate efficiently and respond quickly to changing conditions.

The Solution

The World Bank Treasury modernized its financial data infrastructure with Dremio, launching the “Finance One Lake” strategy. This initiative unified data from over 80 Treasury and Finance systems, replacing years of fragmentation. Deployed on Azure Kubernetes Service and connected to Azure Data Lake Storage, Dremio established a zero-ETL, schema-on-read

> "By implementing Dremio as our lakehouse foundation, we've created a 'finance village with one lake' – a unified platform where all our data streams converge. The combination of zero-ETL architecture and AI-powered automation has transformed how our Treasury operates, enabling traders to focus on opportunity analysis rather than manual data entry. This represents a fundamental shift in how we deliver impactful financial services with integrity and excellence."
>
> — Swamy Kiran, Senior IT Officer, Data & Information Management, World Bank

architecture, eliminating data copying and enabling direct querying at the source.

A key component is a centralized semantic layer in Dremio, which ensures all analytics tools, including Power BI and Tableau, operate from consistent business logic, offering true plug-and-play flexibility. This unified layer also supports SHASTRA, an AI-powered trade automation system that uses Azure OpenAI for term-sheet data extraction and verified "golden copy" generation.

Governance and security are integrated using Azure Active Directory and Dremio's identity-based controls, ensuring users across various workflows only access authorized data while enabling broad self-service. By standardizing on Parquet and open lakehouse principles, the Treasury created an open, scalable foundation ready to evolve with technologies like Apache Iceberg and Polaris, avoiding vendor lock-in and ensuring long-term flexibility for expanding analytics and AI initiatives.

Results

The Dremio-powered lakehouse initiative has delivered significant and measurable impact across the World Bank Treasury's operations. Trade processing, which once required six to eight hours of manual effort, now takes roughly fifteen minutes end-to-end, including human review. The AI extraction engine consistently achieves 95% accuracy, allowing traders to quickly validate only the remaining edge cases through a streamlined interface.

These efficiency gains have fundamentally reshaped trader productivity. Instead of spending the majority of their day manually parsing term sheets and entering data, traders can now focus on market analysis, pricing strategy, and opportunity identification—activities especially critical during periods of market volatility. This shift has strengthened the Treasury's ability to respond quickly to changing conditions while improving the overall quality of decision-making.

The implementation of a single semantic layer within Dremio has also resolved long-standing data consistency issues. Teams that once produced conflicting financial metrics now rely on a shared source of truth, ensuring that leadership receives accurate, aligned, and timely reports. The platform has scaled rapidly, onboarding more than 80 data products and consolidating approximately 70% of the Treasury's finance data.

Business agility has improved markedly as well. Analysts can now bring their own data, blend it with trusted corporate datasets, and perform advanced analysis without waiting for IT to build models or pipelines. These self-service capabilities operate within a robust security and governance framework, ensuring teams move quickly without compromising compliance.

Finally, the unified lakehouse foundation has positioned the Treasury for accelerated AI adoption. With clean, well-governed data now readily accessible, the organization is preparing new AI-driven capabilities, including natural language querying available in upcoming Dremio releases. Future enhancements will further strengthen the Treasury's role as an innovator in global capital markets operations.

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.