

CASE STUDY

How Nomura achieved significant improvements in stability and performance

At a Glance

The Customer

NOMURA

Challenge

Nomura faced significant challenges with their large-scale data processing, handling 1.8 billion rows per day in a highly regulated environment with stringent SLAs. They experienced serious production issues due to the limitations of their existing relational databases and outdated Hadoop infrastructure, leading to frequent outages and difficulties in diagnosing problems. Additionally, their Grafana dashboard was rudimentary, and they lacked the tools to properly test and emulate production workloads, which compounded the challenges of maintaining system stability and performance.

Solution

Nomura's solution involved integrating Dremio to enhance their data infrastructure, improving system performance and stability. They upgraded their Grafana dashboard for better diagnostics and created a hybrid data platform combining on-premises and cloud environments. This integration enabled seamless scaling, adaptability, and a 13.9% performance boost, optimizing their data processing capabilities.

Results

Nomura achieved significant performance and stability improvements by integrating Dremio and MinIO into their data infrastructure, resulting in a 13.9% performance increase. The updated Grafana dashboard with executor-level metrics improved system diagnostics and stability, reducing production issues. Additionally, the hybrid data platform allowed for seamless scaling and adaptability across on-premises and cloud environments, optimizing Nomura's data processing capabilities.

The Business

Nomura is a prominent global financial services firm headquartered in Tokyo, Japan. Specializing in investment banking, securities, wealth and asset management, and retail banking, Nomura serves a diverse client base across the world. The company prides itself on delivering comprehensive financial solutions and strategic insights tailored to meet the evolving needs of its clients. With a strong emphasis on innovation and leveraging its extensive market knowledge and networks, Nomura aims to drive growth and create value while maintaining a commitment to responsible business practices.

Founded on a legacy of financial expertise and a deep-rooted presence in global markets, Nomura is dedicated to fostering economic prosperity and contributing positively to society. Through its operations, Nomura strives to uphold high standards of integrity, transparency, and sustainability, aiming to build enduring relationships with clients and stakeholders alike. As a leader in the financial services industry, Nomura continues to innovate and adapt to the dynamic landscape of global finance, positioning itself as a trusted partner in navigating complex financial challenges and opportunities.

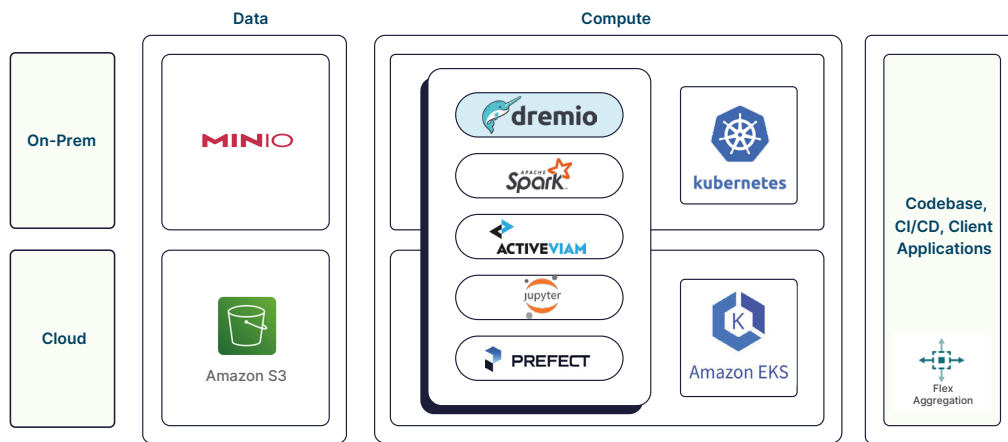
The Challenge

Nomura faced several significant challenges in optimizing its data infrastructure, particularly centered around its use of Hadoop. Initially, Nomura recognized that its relational databases were inadequate to handle the growing data volumes and complex queries required for market and credit risk analysis. This led them to adopt Dremio five years ago, which provided faster query performance and flexibility compared to traditional databases. However, maintaining stability and meeting stringent SLAs became increasingly challenging due to outdated infrastructure and limited insights into system performance and query behavior.

Key challenges included frequent production issues and serious outages, exacerbated by the difficulty in diagnosing problems and the lack of comprehensive monitoring tools. Nomura's Grafana dashboard, initially rudimentary, evolved to include executor-level metrics crucial for monitoring distributed system performance. Additionally, they struggled with the inefficiencies of Hadoop, such as high operational costs and the complexity of managing a monolithic stack, which hindered their ability to scale and innovate effectively.

"The main reason that we went with Dremio was its openness and it is lightning fast. If you build your materialized views, what we call a raw reflection, if you optimize it for querying, it goes like lightning."

**Connor Brennan -
Managing Director -
Risk IT, Nomura**



The Solution

Nomura implemented several solutions to address their data infrastructure challenges. They enhanced their Grafana dashboard to provide executor-level metrics for better system insight. They created a structured metadata layer and developed a query replay system to emulate production workloads in non-production environments, improving regression testing. Nomura also transitioned from Hadoop to MinIO, utilizing Kubernetes for scalability, which allowed them to repurpose existing hardware. Additionally, they adopted Apache Iceberg to simplify their data pipeline and improve data management capabilities.

Dremio played a crucial role in Nomura's solution by serving as their data processing engine. Initially, Nomura adopted Dremio to handle their large datasets and demanding query workloads. Over time, they integrated Dremio more deeply by enhancing their monitoring and diagnostic capabilities with an evolved Grafana dashboard. They also developed a metadata layer to better manage queries and implemented a query replay system for improved testing. Dremio's performance and flexibility were pivotal in Nomura's data infrastructure improvements and transition to new technologies like MinIO and Apache Iceberg.

Results

By implementing their new infrastructure solution, Nomura achieved significant improvements in stability and performance. The integration of Dremio, combined with MinIO over Hadoop, resulted in a 13.9% increase in performance. Enhanced monitoring through the updated Grafana dashboard, which now includes executor-level metrics, allowed for better diagnostics and system stability,

addressing previous issues related to production outages and regulatory compliance.

Additionally, Nomura's solution enhanced efficiency and scalability. The integration of Dremio's query replay system improved query management and regression testing, reducing production issues. The hybrid data platform with Kubernetes and MinIO facilitated seamless scaling and adaptability for both on-premises and cloud environments, ensuring a robust and flexible infrastructure to meet future demands. By leveraging Dremio's capabilities, Nomura was able to optimize their data processing, leading to a more stable and high-performing environment.

Conclusion

Nomura faced several challenges in managing its large and complex data infrastructure. Their traditional relational databases struggled to handle the increasing data load, resulting in production issues and outages. They lacked adequate tools for system diagnostics, making it difficult to identify and resolve issues

To address these challenges, Nomura embarked on a transformation journey focused on enhancing system stability and scalability. This included upgrading their infrastructure with newer hardware and exploring alternatives like MinIO for better storage efficiency and cost-effectiveness. They also pioneered the use of Iceberg to improve data management and achieve versioning capabilities critical for their analytics. Overall, Nomura's approach aimed to modernize their data platform, ensuring it could support their intensive workload and regulatory requirements while paving the way for future hybrid cloud integrations.

ABOUT DREMIO

[Dremio](#) is the unified lakehouse platform for self-service analytics and AI, serving hundreds of global enterprises, including Maersk, Amazon, Regeneron, NetApp, and S&P Global. Customers rely on Dremio for cloud, hybrid, and on-prem lakehouses to power their data mesh, data warehouse migration, data virtualization, and unified data access use cases. Based on open source technologies, including 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.