

HEXAWARE

Real-Time Data Synchronization on AWS for Financial Capital Markets

Faster insights through modern architecture

Case study



Client

| Client

Challenge

Solution

Benefits

Summary

Delivering Trusted Data Across Global Operations

A leading global financial services enterprise operating in the capital markets sector manages vast volumes of transactional and reference data across multiple business platforms. The enterprise supports institutional clients who depend on timely, accurate information to make critical investment and operational decisions.



Challenge

- Client
- Challenge**
- Solution
- Benefits
- Summary

Keeping Pace with Evolving Capital Market Data Demands

In capital markets, decisions are only as good as the data available at the moment they are made. The customer's existing data synchronization processes struggled to keep pace with the speed and scale of the business.

Large-scale data transfers between the analytics data warehouse and operational systems relied heavily on bulk movement processes that introduced significant latency. As a result, business users and applications often operated on stale information, limiting their ability to respond quickly to market conditions.

The challenge became more complex as data structures evolved over time. New columns, changing data types, and expanding datasets required frequent intervention, increasing maintenance effort and slowing the delivery of new business requirements.

The enterprise also faced growing operational overhead from managing and monitoring multiple data pipelines. Maintaining reliability while supporting increasing data volumes placed additional pressure on technology teams.

To support future growth, the customer needed a solution that could provide near real-time data synchronization, scale seamlessly with demand, adapt to schema changes, and reduce the burden of pipeline management.



Solution

- Client
- Challenge
- | **Solution**
- Benefits
- Summary

A Modern AWS Foundation for Real-Time Data Synchronization

To support its data-driven operating model, a modernization initiative aimed at improving data movement between its analytics and operational platforms was planned. The objective was to create a scalable, resilient framework capable of delivering timely and accurate data while reducing operational overhead and improving flexibility for growth.

Hexaware partnered with the client to design and implement a modern synchronization framework on AWS that aligned with these business objectives. The solution used a serverless, event-driven architecture to automate data movement, eliminate infrastructure management burdens, and enable seamless scalability as data volumes increased.

Real-Time Data Through an Event-Driven Architecture

The new framework was built to efficiently and reliably synchronize data between analytical and operational environments. Using a serverless, event-driven architecture built on AWS, we established a highly automated data movement layer capable of supporting both large-scale migrations and ongoing incremental updates.

Workflows managed execution states, automated retries, and improved resilience, while event-driven processing enabled dynamic execution without manual intervention.

Further, data extraction and transformation activities used automated pipelines, and a scalable intermediary layer supported high-volume data transfers. A metadata-driven configuration model simplified the onboarding of new datasets, reduced administrative effort, and enabled more flexible pipeline management across the data ecosystem.

Large-Scale and Continuous Data Synchronization

To accommodate different business requirements, the framework incorporated two synchronization models.

Full Load Framework

For initial migrations, major refreshes, and schema evolution scenarios, the platform supported bulk data movement and selective table synchronization. Business teams could initiate full-load processes manually or through APIs whenever significant data updates were required.

Incremental Load Framework

To ensure data remained continuously aligned across environments, incremental synchronization processes were executed every 15 minutes. Using watermark-based change tracking, the framework identified and processed only modified records, reducing latency and improving efficiency.

Solution

- Client
- Challenge
- Solution**
- Benefits
- Summary

Key Capabilities of the AWS Solution

This architecture ensures high throughput, low latency, and flexibility in handling diverse data scenarios.

The architecture intelligently selected the optimal transfer method based on data volume. Large datasets exceeding 1.5 million records were routed through Amazon S3 to maximize throughput, while smaller datasets leveraged Foreign Data Wrapper (FDW) technology to accelerate updates and minimize processing overhead.

Additional capabilities included:

- Configuration-driven execution through DynamoDB and AWS Systems Manager Parameter Store
- REST-based APIs for workflow initiation and management
- Automated schema and data validation
- Secure credential management through AWS Secrets Manager
- Automated merge operations within PostgreSQL
- Integrated monitoring, alerting, and exception handling

The result was a highly scalable synchronization platform that enabled the enterprise to efficiently move millions of records, improve data availability across business systems, and create a reliable foundation for real-time analytics and operational decision-making.

The solution was selected due to:

- Serverless-first design: Eliminates infrastructure management
- Dynamic routing (S3 vs FDW): Optimizes performance based on data volume
- Decoupled orchestration: Improves maintainability and scalability
- Config-driven framework: Reduces deployment effort for new dataset
- High fault tolerance: Step Functions retries and monitoring

Compared to traditional ETL tools, this approach is optimized for performance and cost.

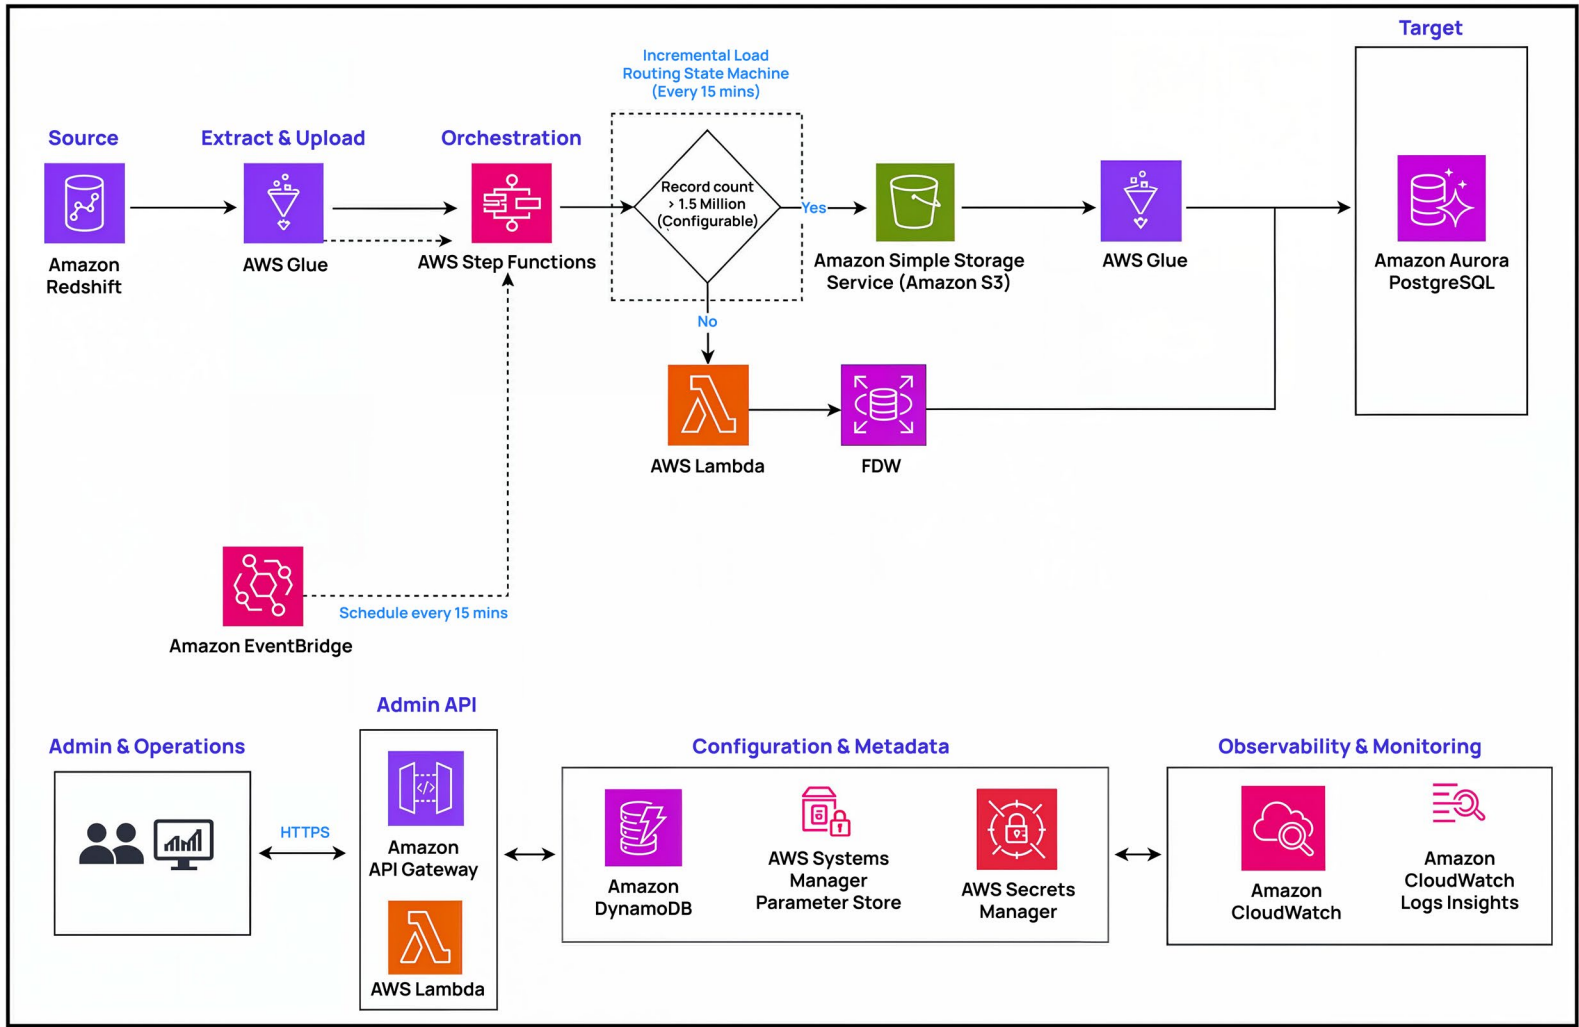


Solution

- Client
- Challenge
- Solution**
- Benefits
- Summary

Solution Architecture

Hexware’s AWS architecture drives scalable and efficient data movement with AWS’ newest capabilities:



- AWS Step Functions: Visual workflows, retries, and state management
- AWS Lambda: Event-driven triggers for incremental loads
- AWS Glue: Scalable ETL and Redshift UNLOAD integration
- Amazon S3: High-throughput intermediate storage
- Amazon DynamoDB: Metadata-driven pipeline configuration
- Amazon Redshift: Source analytics data warehouse
- Amazon Aurora PostgreSQL: High-performance target database
- AWS Secrets Manager & Systems Manager Parameter Store: Secure configuration and secret handling

Benefits

- Client
- Challenge
- Solution
- Benefits**
- Summary

Access to Trusted Market Intelligence at Scale

With a modern, automated synchronization framework in place, critical data is now shared across analytical and operational systems. By ensuring timely access to accurate and consistent information, the solution enabled faster responses to market dynamics.

Up to 70% Lower Data Latency for Faster Decisions

By replacing batch-oriented synchronization processes with near real-time updates, the customer reduced data latency by approximately 70%. Business users and downstream applications now have access to significantly fresher information faster.

Data Refresh Every 15 Minutes across Critical Systems

The incremental synchronization framework delivers updates every 15 minutes, ensures operational platforms and user-facing applications remain aligned with the latest available data from the enterprise data warehouse.

50% Reduction in Manual Operational Effort

Automation across orchestration, validation, monitoring, and execution significantly reduced manual intervention requirements. Technology teams now spend less time managing pipelines and more time focusing on strategic initiatives.

Zero Infrastructure Management for Scalability for Millions of Records

The serverless architecture automatically scales to support large data volumes without requiring dedicated infrastructure administration. This lets the enterprise accommodate future growth while maintaining performance and reliability.

Faster Dataset Onboarding through Configuration-Driven Automation

The metadata-driven framework allows new datasets to be onboarded quickly through configuration changes rather than custom development, accelerating delivery timelines and improving agility.

Improved Reliability and Fault Tolerance Across Data Pipelines

Built-in retries, workflow orchestration, monitoring, and automated error handling provide greater resilience and visibility, reducing the risk of data synchronization failures and business disruption.



Summary

- Client
- Challenge
- Solution
- Benefits

| **Summary**

Resilient and Future-ready Capital Market Operations

Today, the customer operates a modern, serverless data synchronization platform that delivers near real-time updates across critical capital markets applications. Data now fluidly moves between Amazon Redshift and Amazon Aurora PostgreSQL through an intelligent, event-driven architecture designed for scale and resilience.

Business users benefit from more accurate and timely information, enabling faster decisions and improved user experiences across downstream applications. With a scalable foundation capable of processing millions of records, the customer is well positioned to support future growth, adapt to evolving business requirements, and continue delivering reliable, data-driven experiences across the enterprise.



Ready to modernize your data ecosystem?

Discover how [Hexaware](#) and [AWS](#) build scalable, serverless AWS solutions that deliver real-time insights, lower operational costs, and future-ready data platforms. To learn more, reach out to marketing@hexaware.com.

About Hexaware

Hexaware is a global technology and business process services company. Every day, Hexawarians wake up with a singular purpose: to create smiles through great people and technology. With this purpose gaining momentum, we are well on our way to realizing our vision of being the most loved digital transformation partner in the world. We also seek to protect the planet and build a better tomorrow for our customers, employees, partners, investors, and the communities in which we operate.

With offices across the world, we empower enterprises to realize digital transformation at scale and speed by partnering with them to build, transform, run, and optimize their technology and business processes.

Learn more about Hexaware at <https://www.hexaware.com>.

Get in touch