HEXAWARE

AWS Cloud-Native Pricing Platform Modernization for an Insurer

Rebuilding the Data Foundation

Case study

Client

A Global Insurer

A leading specialty insurer and reinsurer within the Lloyd's of London ecosystem, headquartered in London with operations across Europe, North America, and Asia, underwriting a diverse risk portfolio. The company drives operations with cloud and advanced analytics for pricing precision, strong compliance, and operational efficiency.

Challenge

Complex Pricing Data Slowed Risk Assessment and Underwriting

The client's legacy pricing platform stored actuarial and risk data in deeply nested JSON structures within on-premises systems, making it difficult to access, transform, and analyze at scale.

Data preparation for pricing analytics, regulatory reporting, and model comparisons relied heavily on manual processes, leading to long turnaround times and operational inefficiencies.

At the same time, the lack of automated, event-driven pipelines, version control, and reliable audit trails limited their ability to recalibrate pricing models quickly or respond to changing market conditions.

As pricing models and data volumes grew, scalability, governance, and compliance requirements further increased complexity—driving the need for a cloud-based solution that could simplify data, improve accessibility, and enable efficient, analytics-ready processing.

Solution

Building a Scalable, Event-Driven Pricing Data Platform

Hexaware led a full-scale modernization of the client's pricing ecosystem by transitioning from legacy, on-prem systems to a cloud-native platform on Amazon Web Services. The focus was to simplify complex pricing data, improve accessibility, and create a scalable foundation for analytics and decision-making.

AWS Services Supporting the Platform

Historical pricing data was migrated from SQL Server into Amazon S3, while metadata-driven transformation pipelines were built using AWS Glue to standardize data into analytics-ready formats. Serverless querying was enabled through Amazon Athena, and governed access was established using AWS Lake Formation—ensuring secure, compliant, and easy data consumption across teams.

Building Event-Driven, Serverless Architecture

To eliminate manual dependencies and improve responsiveness, Hexaware implemented a scalable, event-driven, serverless data ingestion and processing pipeline based on modern data lake principles.

Seamless Data Ingestion across Sources

Both historical and live pricing data were ingested into centralized S3 landing zones—supporting batch inputs from legacy systems and real-time feeds from rating platforms. This ensured continuity between production pricing and analytics.

Automated, Event-Driven Orchestration

Using Amazon EventBridge, workflows were triggered automatically based on events rather than schedules. AWS Lambda functions handled data movement and validation, while AWS Step Functions orchestrated end-to-end processing—creating a reliable, self-managed pipeline.

Decoupled and Resilient Processing

To ensure scalability and fault tolerance, messaging via Amazon SQS enabled asynchronous processing. AWS DataSync automated secure data movement between ingestion layers, reducing operational overhead.

Turning the Data into an Analytics-Ready Lakehouse

At the core of the solution was a modern data lake architecture designed for performance, consistency, and scale.

- Data was transformed and flattened into structured formats using AWS Glue
- Stored in Apache Iceberg tables to enable:
  - » Version-controlled datasets
  - » Incremental processing
  - » High-performance querying

This ensured pricing data was always consistent, traceable, and ready for analysis.

Enabling Real-Time Insights for Business Teams

With serverless access via Amazon Athena and dashboards built on Microsoft Power BI, actuarial and underwriting teams gained near real-time visibility into pricing performance, portfolio risk, and model outcomes—driving faster, data-backed decisions.

Solution

Embedding Governance, Security, and Operational Excellence

To meet regulatory and enterprise requirements, Hexaware embedded governance and security across the platform:

- Fine-grained access control using AWS Lake Formation
- Monitoring and logging via Amazon CloudWatch
- End-to-end deployment through Infrastructure-as-Code (Terraform) and CI/CD pipelines

This ensured a secure, compliant, and repeatable environment across all stages of deployment.

The Outcome: A Future-Ready Pricing Data Platform

The result is a highly scalable, automated, and intelligent data platform—capable of handling growing pricing models, enabling real-time analytics, and supporting faster, more accurate underwriting decisions.

The New AWS Platform Architecture Diagram

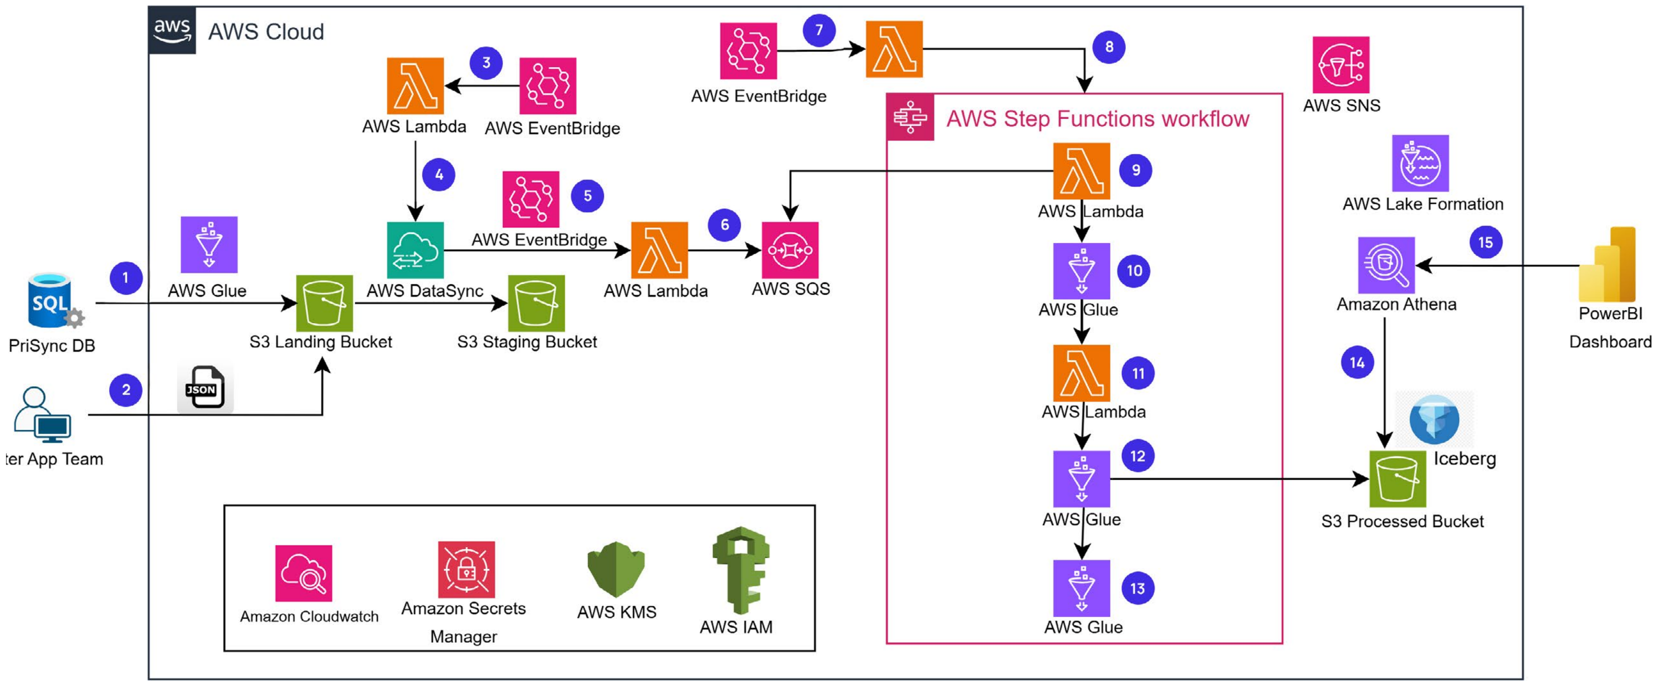

Solution

This architecture is a scalable, event-driven, serverless pipeline on AWS that ingests JSON data, converts it into a canonical Hx format, validates it, and stores it for analytics using modern data lake principles.

1. Data comes from PriSync SQL Server and Hx Live (Rater) into the S3 landing bucket via AWS Glue and direct ingestion.
2. EventBridge and Lambda automate DataSync to move data to the staging bucket. Once completed, events trigger downstream processing through SQS and Step Functions, ensuring decoupled and reliable execution.
3. Glue jobs log metadata in Iceberg tables, then process JSON data by flattening, validating, and applying data quality checks. Any issues trigger SNS notifications. Clean data is stored in Iceberg tables in the S3 processed bucket.
4. Iceberg maintenance optimizes performance. Data is queried via Amazon Athena and visualized in Power BI for near real-time insights.
5. Security, monitoring, and governance are handled through IAM, KMS, Secrets Manager, CloudWatch, and Lake Formation.

How Hexaware's Expertise in Data and Analytics Helps

Hexaware's Data and Analytics expertise brought structure and scalability to a highly complex environment. Years of deeply nested historical JSON were migrated from SQL Server to AWS through careful extraction, schema inference, and reconciliation—ensuring no data loss. To handle evolving pricing models and frequent schema variations, Hexaware designed metadata-driven Glue PySpark jobs that could adapt dynamically, maintaining consistency across the data landscape.

At the same time, Hexaware orchestrated EventBridge, DataSync, Lambda, SQS, and Step Functions into a reliable, event-driven pipeline with built-in failure-recovery. Iceberg tables were implemented with threshold-based batching and tuning to control small files and maintain query efficiency, while Lake Formation enforced fine-grained governance. Terraform and CI/CD pipelines ensured secure, consistent deployments, resulting in a scalable, high-performing data and analytics platform.

Benefits

Measurable Impact Across Insurance Pricing, Compliance, and Growth

The new cloud platform made pricing data easier to access and use. It helped teams work faster, stay compliant, and scale as the business grows.

Up to 45% Faster Data Preparation
Centralized, analytics-ready datasets significantly reduced time spent on data discovery and transformation—accelerating actuarial workflows.

Up to 60% Reduction in Compliance Effort
Automated validation, reconciliation, and audit trails minimized manual intervention—lowering risk and operational overhead.

Up to 40% Improvement in Pricing Accuracy
Standardized transformations and consistent historical data improved model reliability and pricing precision.

Up to 50% Faster Model Deployment
Event-driven pipelines and automation accelerated testing-to-production cycles for new and recalibrated pricing models.

3× Data Scalability Without Rework
The cloud-native architecture seamlessly scaled across growing datasets, pricing models, and business product classes.

Summary

A Real-Time, Scalable Insurance Pricing Engine in Action

Today, the client operates on a modern, cloud-native pricing platform that not only processes data continuously but also keeps pricing decisions aligned with real-time market shifts. Underwriting and actuarial teams can act faster, test models more frequently, and respond to risk with greater confidence—without being slowed down by data bottlenecks.

Over time, this creates a compounding advantage: compliance becomes easier with built-in traceability, scaling new products or models no longer requires rework, and the business can adapt quickly to changing regulations and market conditions. What was once a reactive, effort-heavy process is now a proactive, scalable system that supports sustained growth and sharper competitive positioning.

HEXAWARE

Scale Your Pricing Platform with AWS

Build on AWS with Hexaware for faster insights, scalable pricing models, and solidified compliance. Get cloud-native solutions to help you move from legacy complexity to real-time decisioning.
Learn more: https://hexaware.com/partner/aws/

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