DATA ECONOMY

Modernizing Compensation Data with DevOps & Cloud Operations on AWS

Category

DevOps & CloudOps

Table of Contents

About the Customer

The customer is one of the leading mutual life insurance companies in the United States, serving individuals, families, and businesses for more than 160 years. The company offers a broad portfolio of life insurance, disability income insurance, dental and vision benefits, annuities, and investment solutions while maintaining a strong commitment to financial stability, customer experience, and operational excellence.

As part of its digital transformation strategy, The customer's benefits business initiated the development of a centralized Group Compensation Data Hub (GCDH) to modernize its compensation processing ecosystem. The objective was to consolidate data from multiple legacy systems into a scalable cloud-native platform capable of supporting standardized compensation processing, enterprise reporting, operational governance, and future business growth.

The modernization program required a highly available, secure, and scalable AWS architecture capable of processing hundreds of compensation data elements, supporting metadata-driven data pipelines, enabling automated deployments, and providing centralized operational visibility across the platform. The solution established a unified foundation for both modern software delivery practices and resilient cloud operations.

Customer Challenge

The customer's existing compensation distribution landscape relied on multiple legacy applications and data sources that independently generated compensation information for brokers and internal sales organizations. The fragmented architecture increased operational complexity, introduced technical debt, and made it difficult to maintain consistent compensation processing, reporting, and governance.

The organization required a centralized platform capable of consolidating more than 300 compensation-related data elements from multiple enterprise systems into a single source of truth. The platform needed to support inbound data ingestion, standardized transformation pipelines, integration with Varicent ICM for broker compensation calculations, and generation of standardized outbound datasets for downstream business applications.

In addition to modernizing data processing, The customer required automated deployment workflows, version-controlled development, workload scheduling, centralized monitoring, operational governance, auditability, and resilient cloud infrastructure capable of supporting future compensation programs while reducing operational overhead and improving platform reliability.

Partner Solution

DataEconomy designed and implemented the Group Compensation Data Hub (GCDH) on AWS to centralize compensation data processing using a scalable, metadata-driven architecture. The solution combines Amazon S3, Databricks, automated orchestration, CI/CD practices, and cloud operations best practices to provide a unified platform for compensation data management.

The architecture implements a Medallion data model consisting of Bronze, Silver, and Gold layers on Amazon S3 to standardize data ingestion, validation, transformation, reconciliation, and publishing. Metadata-driven processing frameworks and reusable utility components enable consistent execution across inbound, outbound, and downstream data pipelines while supporting evolving business requirements.

Databricks workspaces deployed within a secured Amazon VPC process compensation data using standardized transformation frameworks, validation logic, reconciliation processes, audit logging, and Delta Lake capabilities. Integration with Varicent ICM enables automated compensation calculations, while standardized outputs are delivered to downstream reporting platforms and business applications.

From a DevOps perspective, the platform incorporates automated software delivery through GitHub-based source control, Jenkins CI/CD pipelines, metadata-driven pipeline development, automated testing and validation, and Control-M orchestration for scheduled workload execution. Version-controlled deployments and standardized development practices improve software delivery consistency while simplifying ongoing platform enhancements.

From a Cloud Operations perspective, the solution emphasizes operational excellence through centralized monitoring with Amazon CloudWatch, audit logging using AWS CloudTrail, resilient Multi-AZ deployment, scalable Amazon S3 storage, metadata-driven operational governance, workflow monitoring, exception management, configuration management, and operational KPI visibility. These capabilities enable secure, reliable, and scalable day-2 cloud operations across the compensation platform.

AWS Services Used

AWS DevOps Competency Services:

- Amazon S3
- AWS Lambda
- Amazon VPC
- Amazon EC2
- Amazon EBS
- Amazon CloudWatch

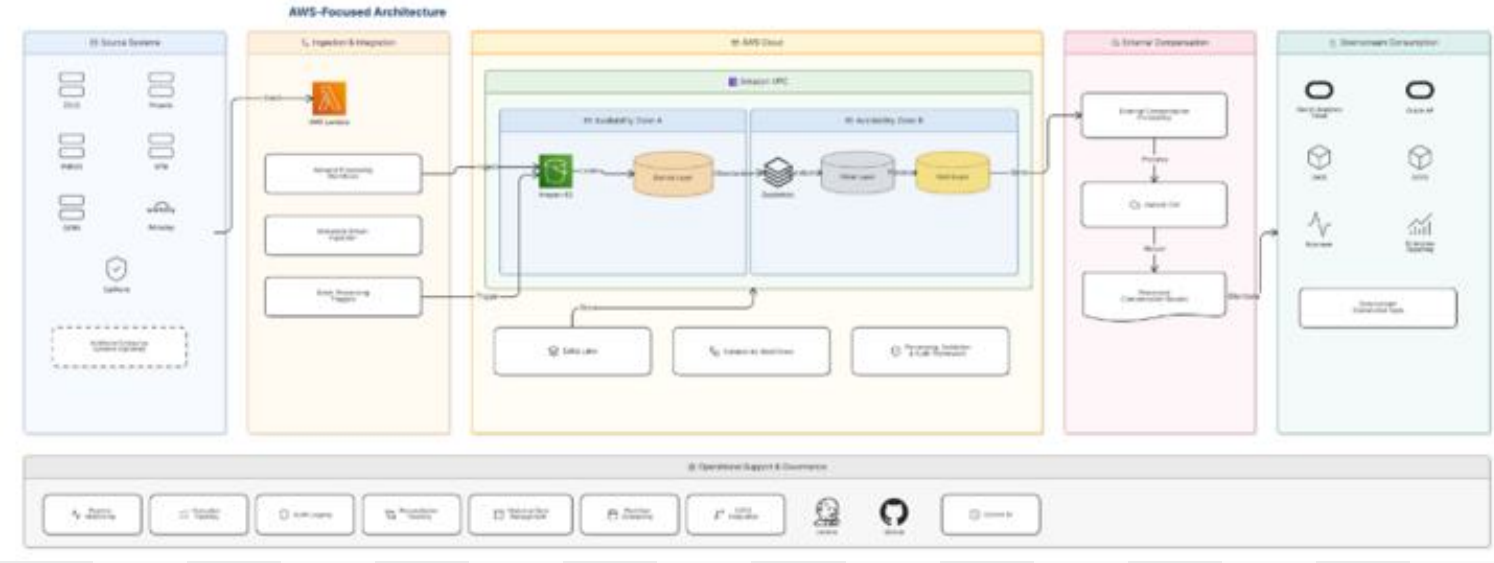

These services support automated data ingestion, cloud-native processing, pipeline automation, deployment workflows, application monitoring, and operational support for software delivery.

AWS Cloud Operations Competency Services:

- Amazon S3
- Amazon EC2 (Databricks Clusters)
- Amazon VPC
- Amazon CloudWatch
- AWS CloudTrail

Optional (where implemented)

- Amazon EC2 (Jenkins Server Hosting)

Compensation Data Hub (CDH)
Centralized Compensation Data Processing & Cloud Operations Platform (AWS-Focused Architecure)

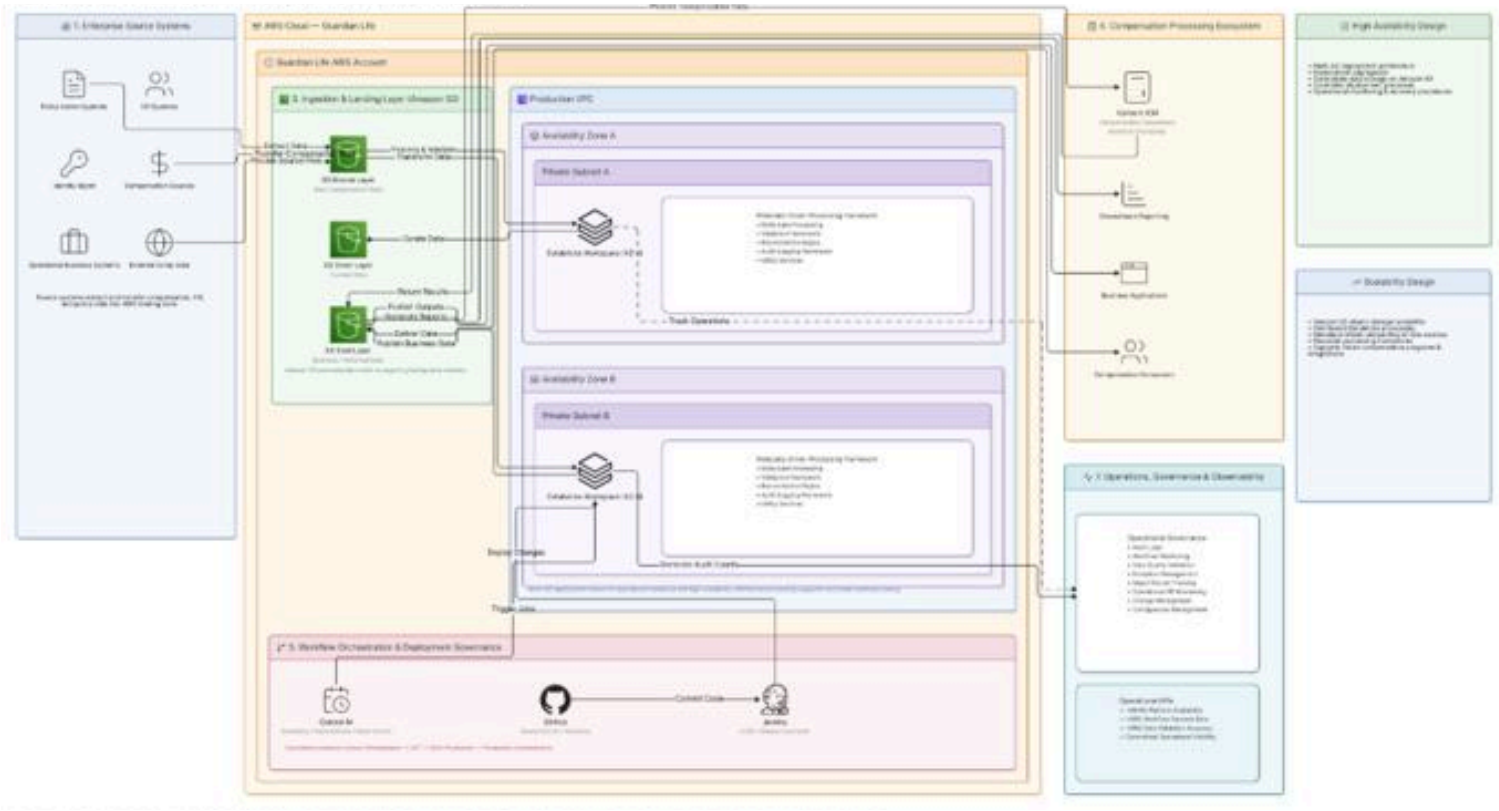

These services provide centralized monitoring, operational governance, audit logging, scalable infrastructure, resilient storage, and secure cloud operations supporting the production compensation platform.

Results and Benefits

The implementation of the Group Compensation Data Hub (GCDH) established a centralized platform for compensation data processing, enabling the customer to modernize its compensation distribution ecosystem while improving both software delivery practices and cloud operations.

By leveraging Generative AI and automated Terraform workflows, the solution reduced the manual effort associated with infrastructure request handling, code generation, validation, and deployment activities. Application teams gained the ability to submit infrastructure requirements through structured workflows and automatically generate standardized Terraform code aligned with enterprise engineering standards.

The implementation strengthened governance and operational traceability by integrating validation workflows, policy checks, pull request automation, audit tracking, and centralized reporting into the infrastructure provisioning lifecycle. Automated validation and sandbox testing improved deployment consistency and reduced the operational complexity associated with manual infrastructure verification processes.

Key business outcomes included:

- Established a single source of truth for compensation data by consolidating more than 300 data elements from multiple legacy systems into a centralized platform.
- Automated end-to-end compensation data processing, including generation of inbound files for Varicent ICM and standardized outbound extracts for downstream business applications.
- Reduced manual effort and processing cycle times through metadata-driven data pipelines, reusable processing frameworks, and standardized transformation workflows.
- Lowered technical debt by consolidating legacy compensation systems into a unified cloud-based architecture, simplifying ongoing platform maintenance and future enhancements.

- Improved transparency and traceability through audit logging, reconciliation processes, workflow monitoring, and standardized data validation across compensation processing workflows.
- Enhanced software delivery by implementing version-controlled development with GitHub, automated CI/CD pipelines using Jenkins, and scheduled workflow orchestration through Control-M.
- Strengthened operational governance through centralized monitoring, operational KPI tracking, configuration management, and audit capabilities supporting reliable day-to-day platform operations.
- Built a scalable and reusable data platform capable of supporting future compensation models, additional business lines, and new enterprise data sources with minimal architectural changes.

About the Partner

DataEconomy is an AWS consulting partner specializing in cloud modernization, data platforms, DevOps transformation, cloud operations, and generative AI solutions. The company helps enterprises modernize business-critical applications and data ecosystems by combining cloud-native architectures, automation, and AWS best practices to deliver scalable, secure, and operationally efficient platforms.

DataEconomy's expertise spans data engineering, application modernization, DevOps automation, Infrastructure as Code, cloud operations, observability, governance, analytics, and enterprise data platforms. By leveraging AWS-native services together with modern engineering practices, DataEconomy enables organizations to improve software delivery, strengthen operational resilience, simplify cloud management, and build scalable data platforms that support long-term business growth and measurable operational outcomes.