



Cloud Data Architecture: Serverless Data Lakehouse Migration

LOGYQ Use Case: Serverless Data Lakehouse Migration

Focus: Migrating from Legacy Systems to Scalable, Consumption-Based Cloud Architecture

Executive Summary: The Weight of the Monolith

Maintaining aging, on-premise data warehouses and legacy infrastructure is a fundamental barrier to modern business growth. This infrastructure demands high, inflexible **Capital Expenditure (CapEx)** for hardware, struggles with scaling during critical peak loads, and is fundamentally incapable of integrating easily with modern workloads like Machine Learning and Generative AI. The result is a high **Total Cost of Ownership (TCO)**, slow time-to-market for new features, and stifled innovation.

The LOGYQ goal is to eliminate CapEx, maximize elasticity, and achieve significant TCO reduction by architecting a modern, serverless cloud environment.

LOGYQ's Solution: Phased ELT Cloud Transformation

We execute a structured, risk-mitigated migration to a consumption-based Cloud Data Lakehouse architecture (tailored for AWS, Azure, or GCP).

1. Strategic Assessment and Blueprint

- **Data Estate Audit:** A comprehensive audit of all existing data assets, dependencies, data flow complexity, and regulatory compliance requirements.
- **Target State Design:** We create a **Target State Architecture (TSA)** blueprint using Infrastructure-as-Code (IaC) templates (e.g., Terraform). This design prioritizes **decoupling storage from compute** and maximizing serverless components.
- **Risk Mitigation:** Define a phased, iterative migration strategy (e.g., "lift-and-shift" for some components, "re-architect" for complex ones) to ensure zero downtime.

2. Serverless ELT Pipeline Implementation

- **Ingestion:** We implement modern Extract, Load, and Transform (ELT) pipelines, using cloud-native, serverless services (e.g., AWS Glue, Azure Data Factory, GCP Dataflow). This eliminates the cost and management overhead of maintaining constantly-running ETL servers.
- **Data Lakehouse Deployment:** The architecture is centered around scalable object storage (e.g., S3, ADLS) for the Data Lake layer, paired with a cloud-native Data Warehouse (e.g., Snowflake, Databricks) for high-performance analytical queries.



Cloud Data Architecture: Serverless Data Lakehouse Migration

- **Cost Governance:** We implement technical controls for resource tagging and automated tiering of data from "hot" to "cold" storage, realizing immediate savings.

Technical Implementation & Architecture

The solution is built entirely on modern, auto-scaling cloud services.

- **laC:** The entire environment is version-controlled and deployed via Infrastructure-as-Code (Terraform), ensuring the architecture is auditable and repeatable.
- **Scalability:** The serverless nature guarantees resources **instantly scale up** to handle high-volume data ingestion or complex queries and **scale down to zero** during idle times.
- **Security:** Multi-layered defense, including private networking, full data encryption at rest and in transit, and role-based access control (RBAC) synchronized with enterprise directories.

Measurable Outcomes & ROI

Metric / Impact Area	Detail	LOGYQ Migration Outcome
Total Cost of Ownership (TCO)	Capital and operational expenses	20-35% reduction in TCO within the first year.
Time-to-Insight	Provisioning a new analytical cluster	Reduced from weeks to minutes due to laC and instant elasticity.
Innovation Readiness	Integration with advanced services	100% cloud-native integration with ML/GenAI toolsets.

LOGYQ Differentiator: We focus on the **TCO-optimized blueprint**. Our use of serverless architecture and laC, combined with detailed cost-monitoring dashboards, ensures that the financial benefits of cloud migration are immediately measurable and sustainable.