

Data Lakehouse Implementation

One governed platform for every workload — open table formats, medallion layers, and the pipelines and controls that keep your data trustworthy. Delivered near-shore by senior architects across the US and Latin America, on your cloud or on-prem.

A lake nobody trusts is a liability.

Data lakes were supposed to end the silos. Most created a new one: cheap storage nobody can vouch for, ringed by warehouses that cost more every quarter and pipelines only one person understands. A lakehouse fixes the architecture — open storage with warehouse guarantees — but only if the governance, quality and operating model are built with it. That is the part we do.

What we implement

Greenfield Lakehouse Build

From nothing to production: platform, medallion layers, ingestion, orchestration and the governance model — designed as one system rather than assembled later.

Migration & Modernization


Move off legacy warehouses, Hadoop and sprawling ETL — workload assessment, phased migration, parallel run and a cutover your business does not feel.

Governance & Quality Layer

Catalog, lineage, access control and data quality wired into an existing platform, so analytics and AI run on data you can stand behind.

Streaming & Real-Time Ingestion

Change data capture from operational systems and continuous streams — including plant telemetry from the industrial edge — landing continuously and reliably.



Medallion layers on open table formats, with governance spanning every layer rather than bolted on at the end.

HOW WE DELIVER

From scattered data to one governed platform.

Understand

We start with the workloads you actually run and the data behind them — never the technology first.

Architect

Senior architects design the target lakehouse end to end: layers, formats, governance and cost model.

Deliver

We build ingestion, transformation and quality pipelines with engineering discipline and tests.

Operate

Managed services keep the platform reliable, governed and cost-efficient long after go-live.

We deliver in increments rather than one big bang: a first governed layer serving a real business workload, then expansion. Migrations run in phases with old and new in parallel, so cutover happens per workload and the business never depends on a single switch.

WHAT YOU GET

Deliverables

- Target lakehouse architecture and platform decision
- Medallion layer design on open table formats (Delta / Iceberg)
- Ingestion, transformation and data quality pipelines
- Governance model — catalog, lineage, access control, data quality
- Runbooks and handover, or an ongoing managed service

TRACK RECORD

Fortune 100 operators. Delivered for Fortune 100 oil & gas operators and enterprise clients across Latin America and the United States.

20+ years. Two decades architecting, building and operating enterprise data platforms.

Certified in-house. AWS, Azure, Databricks, Snowflake and Inductive Automation certifications held by our own architects.

Cloud or on-prem. Your data, your rules — delivered on your cloud or in your own facilities.