CONSULTING GROUP

EMPOWER — DATA & AI

Data & AI Implementation

Governed lakehouse platforms and a semantic model of what your business means — so analytics and enterprise AI both run on one trustworthy source. Delivered near-shore by senior architects across the US and Latin America, on your cloud or on-prem.

Most AI never leaves the pilot.

The gap between a demo and a dependable system is data: governed, well-modeled, and connected to how your organization actually works. A lakehouse holds rows; a vector store holds embeddings. Neither holds meaning — what an entity is, how it relates to everything else, whether it is true, and what may legally be done to it. That was survivable while a person supplied the missing meaning on the way to a decision. An AI agent cannot. We build the foundation first, then put models to work against it with the controls, monitoring and cost discipline production demands.

What we deliver

Semantic Modeling & Enablement

A governed model of what your business means — entities, relationships, classifications, concepts and units — that analytics and AI both read from.

Lakehouse Platforms

Open, governed lakehouse platforms that unify analytics and AI on one copy of your data — no warehouse-versus-lake trade-off.

Data Engineering & Pipelines

Reliable change-data-capture and transformation pipelines from source systems into the platform — monitored, tested and production-grade.

Enterprise AI

RAG over documents and grounded query over governed records, with evaluation, guardrails and MLOps behind both.

Analytics, BI & Reporting

Trustworthy dashboards and a governed semantic layer that turn the platform into self-service insight your teams actually use.

Data Governance & Quality

Catalog, lineage, access control and quality that keep data trustworthy, secure and compliant — data you can stand behind.

REFERENCE ARCHITECTURE

Meaning in the middle, not bolted on the end.

The semantic layer sits between the systems that produce data and everything that consumes it — so quality is enforced as data is written, and an agent asking a question is answered through the same governed path a person would be.

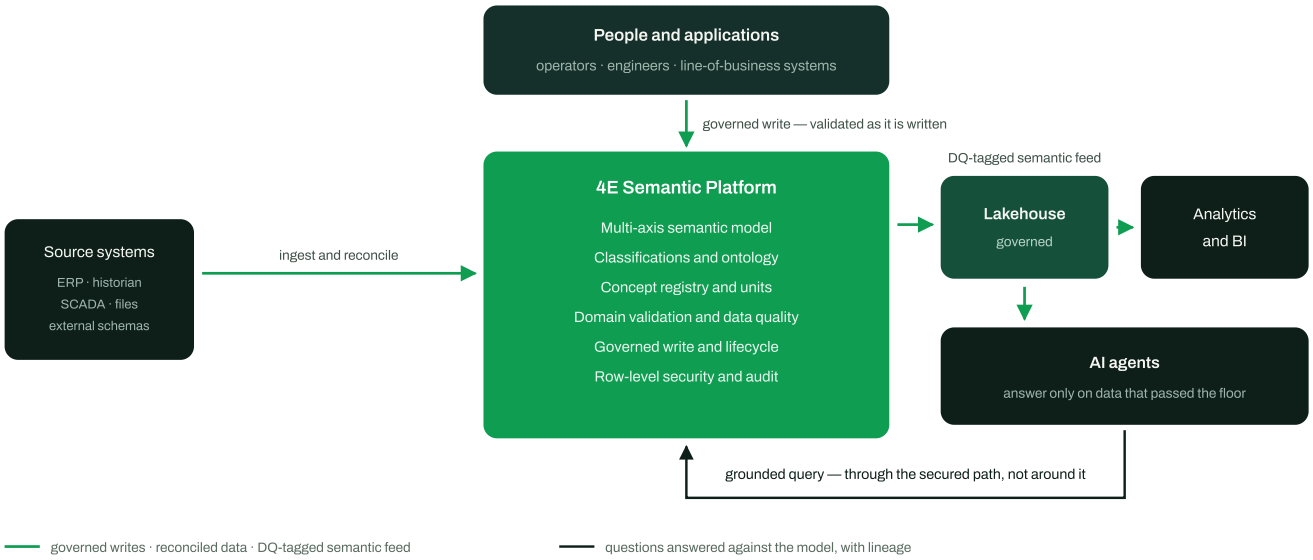

Quality is enforced at the write, not scored afterwards. The feed carries a per-record compliance signal your lakehouse cannot derive on its own, because the domain rules live in the model rather than in the lake — which is what lets an agent be held to data that passed.

BUILT BY 4E

The 4E Semantic Enablement Solution.

The industry's answer to the meaning gap has been to bolt a semantic layer onto the analytical stack. Those layers describe data at rest — a glossary, not a nervous system. Ours is active: every axis of the model does work at runtime. A classification drives query execution rather than documenting it; a domain rule validates at the moment data is written rather than scoring it afterwards; a policy governs what may be written at all. The model is multi-axis by design, because real operations are — a pump belongs to a physical hierarchy, an electrical topology, a process flow and an owning organization at once, and meaning that supports only one tree loses the other three. Objects are declared rather than hand-coded, so the platform generates the schema, the governed API and the validation from the declaration. Quality is enforced where it still can be: each field is validated against its bound domain, value list or concept as it is written, and every record carries the resulting compliance signal outbound — a signal your lakehouse cannot compute for itself, which makes it a gate an agent can be held to. The platform is our own technology; the solution is that platform plus the senior architects who model your business onto it, and the team that supports both long after go-live.

HOW WE DELIVER

From business goal to operating platform.

Understand

We start with your business goals and data reality — never the technology first.

Architect

Senior architects design the semantic model and the governed lakehouse and AI target architecture end to end.

Deliver

We build pipelines, models and analytics with DevOps and MLOps discipline.

Operate

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

WHAT YOU GET

Deliverables

- A governed semantic model of your business — entities, relationships, classifications
- Target lakehouse and AI architecture, cloud or on-prem
- Ingestion, transformation and data quality pipelines
- Enterprise AI grounded in that model, with evaluation and guardrails
- 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 and Snowflake certifications held by our own architects.