The security-and-cost control plane is one surface. Staffing it is the gap.

Two of the four pillars in the cost keynote — Security and Cost — converge on a single control plane. I have built and measured the compute-in-place compression architecture that sits underneath it. This is a point of view and a track record, not a licensing proposal.

THE COST SURFACE, IN YOUR OWN ARCHITECTURE

The keynote metered four things. Each is a place where data is stored, moved, or decoded and re-encoded:

- Storage footprint across hot, warm, and cold tiers.
- The compute-to-storage round trip, paid on every query.
- Database branching in Lakebase, which multiplies stored copies.
- Ingest bandwidth and landing-zone cost in Lakeflow.

Unity AI Gateway already fuses spend control with identity and audit on every model call. Security and cost are not two roadmaps. They are one plane, and that plane needs an owner.

THE APPROACH — TIER-AWARE COMPUTE-IN-PLACE COMPRESSION

Tier	Optimizes for	Mode	What it buys
Hot · Lakebase / Postgres	Low-latency reads, writes, point lookups	Light real-time mode, append-in-place	Fast mutation and seek; no heavy re-encode on write
Warm · Lakehouse / Delta · Iceberg	Analytical scans, columnar throughput	Scan-friendly mode, region and predicate ops	Predicate and region work without full decode
Cold · object archive	Maximum density, rare access	High-ratio archive mode, offline batch	Smallest footprint where the storage bill concentrates

This is not one byte-layout pretending to be optimal everywhere — those tiers exist because they optimize for opposite things. It is one interface with tier-tuned encodings, and it must compose with Delta Lake and Iceberg, not replace them. That is consistent with the no-lock-in pillar, not in tension with it.

WHAT RUNS WITHOUT DECODING THE DATA

- Predicate gating: skip blocks that cannot match, before any decode.
- In-place signature search: find records by compact signature, no rehydration.
- Region decode: return only the requested slice, not the whole object.
- Header-only health and metadata: integrity and statistics without reading the body.

Each removes a decode-then-recompute step that the platform currently meters. This is the structural case: a faster point codec does not reproduce it without re-introducing the seams.

EVIDENCE — BUILT AND MEASURED, ANONYMIZED

The architecture above is not theoretical. I built and benchmarked the full hot-to-cold chain on real raw industrial sensor data from a Fortune 500 industrial-equipment manufacturer, under NDA, with bit-exact lossless verification on every row.

Capability	FHC	FFV1	x265	zstd	lz4
Bit-exact lossless	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**
Fits real-time decode budget	**Yes**	no	no	**Yes**	**Yes**
Compressed-domain predicate gating	**Yes**	—	—	—	—
In-place signature search	**Yes**	—	—	—	—
Region / tile decode	**Yes**	—	—	—	—
Health from header only	**Yes**	—	—	—	—

bit-exact
Lossless on every benchmarked row of real raw sensor data

24–26 ms
Decode against a 33 ms real-time frame budget

~15%
Lossless storage reduction vs the strongest lossless baseline (FFV1) on the benchmark set

4 ops
Compressed-domain operations no general-purpose codec offers

Customer-data result: 34.1× lossless compression ratio, measured on a Fortune 500 industrial-equipment manufacturer's own measurement (TDMS) files — structured telemetry of exactly the kind that lands in a lakehouse. A distinct, stronger work stream from the raw-Bayer benchmark above, and the most platform-relevant evidence on this page.

THE HONEST RESIDUALS

- Cost figures are workload-modeled and become committed only after a run on the target platform's own data.
- The benchmark sample is small; a larger corpus tightens the confidence bounds.
- High-ratio archive encode is offline-batch today; a known optimization improves its economics, not its feasibility.
- The advantage shifts with data shape. The deal-grade number is the one measured on your workload.

THE ROLE THIS DESCRIBES

Principal / Staff Architect — Storage Efficiency and Security Control Plane. Field-facing, San Francisco–based.

First 6–12 months:

- Define the compute-in-place compression interface spanning Lakebase, Lakehouse, and cold archive that composes with Delta and Iceberg.
- Partner with the field org to shorten Fortune 500 security reviews — the gate where seven-figure deals stall.
- Own the unified security-and-cost narrative for strategic enterprise accounts.

Why me:

- 109 patent assets across encode, compute, storage, and transmission (three families, nine jurisdictions); inventor of Fully Homomorphic Compression.
- 8+ years enterprise security architecture for a Fortune 100 private bank and a leading SIEM, with zero security incidents.
- Built for constraints harder than a datacenter: defense-grade systems and a lunar payload, where you cannot decode in the clear and bandwidth is scarce.

An architecture point of view, prepared independently. Fully Homomorphic Compression is cited here as evidence of relevant prior work, not offered as a product or license. · C. McElveen · chelsea.anne.mcelveen@gmail.com