

Why financial-services data platforms fail without domain context

For CIOs and CDOs in insurance and banking. 4-minute read.

The problem

A lakehouse, a medallion architecture and a Databricks or Microsoft Fabric workspace are not, by themselves, a financial-services data platform. They are a substrate. When the engineering is done by teams who do not know what a policy file, a treaty, a reserve cell, an IFRS 9 stage or an SCR component actually is, the platform turns out technically correct and operationally useless. Actuarial, risk and finance keep rebuilding the same numbers in Excel because nothing on the platform speaks their language.

Five specific failure modes — and the domain-aware fix

- Generic entities, no business shape. Tables called `fact_transaction` and `dim_customer` cannot answer "what is the in-force at valuation date?". Treat policy, treaty, claim, reserve cell and exposure as first-class entities in the model. Validate against the statutory actuary, not just the data architect.
- No reporting-basis awareness. The same policy attribute has different values under IFRS 17, SAM and management view. Encode basis explicitly. Without it, every downstream consumer rebuilds the logic, and they will not agree.
- Lineage stops at the lakehouse. Lineage that ends at the gold layer is half a system. Trace through the actuarial run, the overlays and the disclosure. If the chief actuary cannot point at a CSM number and walk back to the source extract, you do not have lineage; you have metadata theatre.
- Reconciliation as afterthought. Build reconciliation pipelines as products, not as scripts. Actuarial-to-finance, sub-ledger-to-GL, model-to-disclosure: each is a domain-aware artefact, instrumented and owned.
- No path off the platform. Power BI and Tableau access actuarial outputs via cached extracts because nobody trusts the lakehouse latency. Publish governed BI data products with SLAs, not a generic semantic layer.

What good looks like

- Policy, treaty, exposure and reserve cell exist as named, versioned data products with owners, contracts and an SLA — not as tables somebody curated last quarter.
- The chief actuary, head of credit risk and finance director use the same numbers in their packs because the platform produces them once.
- New AI and analytics use cases start with a data product, not a six-week wrangling project.