1.1 Architectural Context and Motivation

Return to 1. Introduction

Regulated data interpretation systems operate under conditions that differ materially from typical information systems. They support statutory and supervisory mandates, influence enforcement and policy outcomes, and may support legal, audit, and cross-jurisdictional activities. Interpretive behavior in such systems is not merely operational. It is institutional and often consequential.

These systems frequently persist for decades. Over time, reporting definitions evolve, validation criteria change, analytical techniques mature, and supervisory expectations shift. Organizations replace technology platforms, change vendors, and restructure responsibilities. Yet prior interpretations must remain reconstructable and defensible under the semantic and contextual assumptions in force when the systems produced them.

In many environments, systems implicitly embed semantic interpretation in code, configuration, documentation, or tooling conventions. Interfaces reflect structural exchange rather than governed meaning. Organizations treat version changes as software releases rather than as shifts in interpretive authority. As systems evolve, these implicit semantics accumulate, making cross-period comparison difficult, complicating inter-agency coordination, and increasing maintenance and audit burdens.

Regulated reporting regimes are also frequently bespoke. Differences in terminology, architecture, and implementation practice across agencies and jurisdictions make interoperability expensive and fragile. Even when organizations adopt common data formats, those formats do not guarantee shared meaning. Without explicit architectural treatment of semantic governance, comparable data may conceal divergent interpretations.

The motivation for this Reference Architecture arises from this structural gap. It provides a platform-independent framework that treats meaning, version context, and interpretive authority as first-class architectural concerns. The architecture separates evidence custody, observation generation, semantic canonicalisation, interpretation, validation, and analytical consumption into explicit responsibilities with governed boundaries. This separation supports long-term interpretive stability while allowing implementation diversity and controlled evolution.
