

Regulatory Audit Replay With Historical Policy Versions

by [Nick Clark](#) | Published April 25, 2026

What Regulatory Audit Replay Requires

Emerging AI and autonomy regulations specify audit requirements that go beyond data preservation. The EU AI Act requires post-market monitoring with the ability to reconstruct system behavior in support of investigation. FDA's AI/ML SaMD framework requires predetermined change control plans with the ability to demonstrate compliance with the plan over time. NHTSA's emerging autonomous-vehicle safety reporting requires reconstruction of decision context.

The common thread is policy-aware replay. Regulators don't just want to know what the data was; they want to know what the system was supposed to do with the data under the rules in force. Reconstruction without policy versioning produces incomplete answers that inadequately support regulatory determinations.

Why Regulators Are Converging on Policy-Aware Audit

Regulators have learned that data-only audit produces inadequate answers in autonomous-system contexts. A self-driving system that operated correctly under one policy version may have produced an outcome that under a later policy would be evaluated differently. A medical-decision-support system that classified a patient

correctly under one set of clinical rules may have been re-tuned subsequently. A predictive-policing system that operated under one bias-mitigation policy may have been replaced with a system under a different policy.

Each case requires audit that includes policy provenance. The regulator's question — 'was this system operating compliantly with its applicable rules' — has no good answer if the rules in force at the audited time aren't preserved structurally.

How Architectural Policy Versioning Composes With Audit

Each governance policy is a credentialed observation in the lineage. Audit replay walks the lineage to retrieve the relevant policy version at any target time. The replay can reconstruct: what data the system saw, what policy was in force, what processing the policy specified, what decisions the system reached, what the system's audit-grade lineage records of those decisions show.

The replay is governance-credentialed. The regulator's audit authority is itself a credentialed authority within the architecture. The replay produces credentialed audit observations that the regulator consumes through its own admissibility framework. The architecture supports the regulator-audit pattern structurally rather than through ad-hoc per-system audit tooling.

What This Enables for Regulated AI/Autonomy Deployment

Operators deploying AI/autonomy under EU AI Act, FDA AI/ML SaMD, and emerging US/Asian regulatory frameworks gain audit-grade architecture that maps directly to compliance requirements. The compliance pathway becomes architectural rather than per-deployment custom-built.

Cross-jurisdictional compliance — increasingly important as AI/autonomy operates across regulatory boundaries — gains the structural support that current custom-built compliance does not provide. The patent positions the primitive at the layer where regulated AI/autonomy is converging architecturally.