Independent Verification for Insurers and Auditors

Intentrax Whitepaper | Execution Assurance
Published 2026-07-09 - Intentrax Labs - www.intentrax.com/labs

How signed, replayable execution proof gives adjudicators and examiners evidence they can check themselves - without trusting the AI agent's own account.

Executive summary

Insurers pricing and adjudicating AI-agent risk, and auditors examining AI-agent activity, face the same structural obstacle: the evidence they are handed is usually produced by the party under examination. Logs and transcripts are self-attested, so relying on them requires trusting the entity whose conduct is in question. This paper describes how an execution-assurance primitive - a signed, replay-verifiable Autonomous Execution Proof (AEP) that a third party can check offline - changes what an adjudicator or examiner can independently establish. The practical claim is narrow and deliberate. An AEP lets you confirm, without runtime access and without trusting the executor's logs, that a specific execution occurred, under a specific policy version, bound to a specific authority, and that the artifact you were given is unaltered and re-derivable. This document explains what that supports in adjudication and audit workflows, and - equally important - what it does not.

The adjudication and audit gap

When an AI agent takes an action that later becomes the subject of a claim, a dispute, or an examination, the reconstructing party typically receives application logs and a narrative from the operator. Those records may be accurate. But they are authored by the operator's own systems, which means the examiner cannot distinguish a faithful record from a curated one without independent corroboration. For an insurer this complicates causation and loss adjudication; for an auditor it undermines the assertion that recorded activity is complete and unaltered. Richer monitoring does not resolve this. Observability improves the operator's own visibility, but it produces more self-attested telemetry, not independently checkable evidence. What adjudicators and examiners need is an artifact whose integrity does not rest on trusting the operator, and whose correctness they can re-establish themselves.

What independent verification means here

Independence is operational, not rhetorical. Intentrax provides a standalone verifier that runs on the evidence bundle you are given, using only local files and SHA-256 recomputation. It requires no access to the operator's runtime, no network connection back to the vendor, and no private key material. It recomputes the artifact's hashes, re-derives the governing decision from the recorded inputs and the referenced policy version, and fails closed on any drift or

inconsistency. For an examiner, this inverts the usual dependency. Rather than asking the operator to attest that its logs are true, you take the portable proof and confirm it yourself: the signed proof hash matches the canonical content, the decision re-derives, the authority binding is intact, and nothing has been altered since signing. A pass is reproducible by you; it does not depend on the vendor's systems being online or honest at the moment you check.

What the proof lets you establish - and what it does not

An AEP supports specific, bounded assertions. You can establish that an execution event occurred as recorded; which policy version governed it; which principal, delegation, and constraints authorized it; that both the execution and authority gates permitted the action under a fail-closed model (an ALLOW required both gates; otherwise it would have been a DENY); and that the artifact is tamper-evident and replayable. For an insurer, this sharpens questions of whether an action was within authorized scope and policy at the moment it occurred. For an auditor, it supports completeness and integrity assertions grounded in an append-only registry rather than mutable logs. Equally important is what it does not do. An AEP does not certify that the business decision was correct or prudent, does not attest to real-world outcomes, and is not a compliance certification. Intentrax makes no claim to guarantee compliance, eliminate risk, or provide total security. External evidence you may already hold can be ingested as an execution passport bound to the proof, but it never replaces the AEP. Treat the proof as evidence about execution and authority, not a verdict on judgment or law.

Fitting proof into existing workflows

The evidence is designed to travel. Each AEP can be packaged as a portable envelope containing the machine-authoritative canonical JSON, an AEP export (AEP-JSON) and PDF for human review, and human-readable companions - a certificate, a verification report, and a bundle README - alongside an Agent bill-of-materials and Execution bill-of-materials that reference model, policy, authority, and delegation without requiring disclosure of raw prompts. That envelope is intended to drop into an existing claims file, audit workpaper, or evidence-management system. A typical examiner flow: receive the portable bundle from the operator; run the standalone verifier locally; confirm the pass and read the human-readable certificate and verification report; retain the canonical artifact and its hash in the file. Because verification is offline and reproducible, the same check can be repeated later, by a different reviewer, with the same result - a property that matters for defensible adjudication and for audit re-performance.

Current stage and market context

We are candid about maturity because that is what an assurance conversation requires. Intentrax is pre-revenue and in production-beta. The engine has recently begun signing real proofs and the standalone offline verifier is operational. Assurance Levels 1-4 and a broader trust network are planned and labeled as such. Live production traffic and mutation are currently blocked pending controlled unlock, and production-facing evidence bundles are deliberately verified as integrity passes rather than production go-live claims. The motivation for this work is market demand - from insurers, auditors, and enterprise buyers who need to evidence AI-agent activity - not a regulatory mandate. AEVN is an open verification protocol; Intentrax is its reference implementation, not its owner, which means the verification model is not proprietary lock-in. We engage with insurers and auditors as design partners and will always describe the current state,

including the blocks above, rather than an idealized platform.

How to engage

The most useful first step for an adjudicator or examiner is hands-on: obtain a sample evidence bundle and run the standalone verifier yourself, confirming that it recomputes hashes and re-derives the decision with no runtime access, no network, and no keys. From there, we can walk through how the portable AEP export maps to your claims file or audit workpaper, and where its bounded assertions do and do not carry weight in your process. For design-partner engagements we can discuss the planned Assurance Levels, the append-only registry model, and how ingested external evidence is bound as a passport without displacing the canonical proof. We do not offer customer references or performance metrics at this stage because there are none to offer honestly; we offer the primitive, the verifier, and a precise account of what each establishes.