



AEVN: An Open Verification Protocol for Autonomous Execution Evidence

Intentrax Labs Research Paper | Verification Theory

Published 2026-07-09 - Intentrax Labs - www.intentrax.com/labs

The evidence model, layering, and authority boundaries by which autonomous execution can be verified independently of the agent that produced it.

Problem statement

When an autonomous agent acts, the record of what it did is normally written by the agent's own runtime. That record is convenient and non-authoritative: a party relying on it must trust the executor's logs, its clock, and its willingness to report faithfully. AEVN, the Autonomous Execution Verification specification, addresses a narrower and more durable question. It defines the evidence model required to determine, without trusting the executor, what execution evidence was normalized, which principal and authority were bound, which policy governed the event, which decision was produced, and whether that evidence can be re-checked outside the runtime that generated it. AEVN is a verification specification, not a monitoring, observability, or orchestration layer. It does not watch agents; it defines how the record of a completed execution becomes independently checkable.

Governance posture: open specification, reference implementation

AEVN is the open verification specification for autonomous execution evidence. Intentrax is the reference commercial implementation aligned with AEVN and does not position itself as the protocol owner. This separation is deliberate and load-bearing. A settlement layer for autonomous work is only credible if the rules of verification are legible to, and re-implementable by, parties who do not run Intentrax. AEVN therefore fixes the evidence model; Intentrax supplies deterministic runtime admission and proof pipelines, Autonomous Execution Proof (AEP) generation, append-only registry storage, portable proof export, and offline and external verifier workflows. Any conformant implementation that produces different canonical serialization or different hashes for the same inputs is, by definition, non-conformant.

What AEVN defines

AEVN specifies seven surfaces: execution event normalization; policy and authority binding; deterministic canonicalization rules; proof artifact requirements; cryptographic integrity rules; external verifier expectations; and registry and replay boundaries. Each surface exists to remove a specific dependency on the executor. Normalization removes dependence on the agent's native log format. Policy and authority binding removes ambiguity about which rules and which delegated identity governed the action. Canonicalization and integrity rules remove dependence on the executor's serialization and make alteration detectable. The verifier and replay

boundaries remove dependence on the executor's continued cooperation, because the evidence can be checked after the fact by a third party.

The layering model

AEVN sits inside a layered stack with explicit ownership at each level. Above it, an external passport or guardrail system may describe an agent's identity, declared capabilities, and pre-action decisions. Below it, Intentrax admission deterministically normalizes proposed agent activity; the Autonomous Execution Proof is the canonical artifact produced under AEVN; the append-only registry persists proofs and provides lookup; and the independent verifier performs replay and proof verification. The boundary rule is strict: external passport or guardrail evidence may be ingested as evidence, but it never replaces AEVN canonicalization, AEP construction, AEP signing, registry append, or verifier replay. When external evidence conflicts with AEVN evidence, AEVN and the locked contracts win.

Interoperability without dilution

Customers will encounter multiple agent passport and guardrail systems. AEVN's stance toward them is compatibility without surrender of authority. An external passport can answer who the agent is, what capabilities it declared, and whether a proposed tool call was allowed by an external guardrail. AEVN answers a different set of questions: what execution evidence was normalized, which policy and authority governed it, how the decision was derived, whether the proof was altered, whether an independent verifier can replay it, and whether it is anchored in the registry. The differentiating property is that AEVN turns those external decisions into replayable, registry-backed evidence rather than accepting them as final. This is not a governance-workflow or compliance-documentation product; it is the evidence substrate underneath such products.

Current state and planned scope

The engine has recently begun signing real proofs, and a standalone offline verifier recomputes hashes and fails closed on authority drift without requiring runtime access, network access, or private key material. Onboarding is available through a one-command SDK initializer across six adapters, and the pipeline is deterministic: the same execution yields the same proof hash. The following are planned and are labeled as such: graded Assurance Levels 1 through 4, and an Intentrax Ledger and trust-network. Live production traffic and live mutation are currently blocked pending independently verified production authorization. Intentrax is at a pre-revenue, production-beta stage; AEVN makes no claim of customers, logos, or operational metrics, and it is not a regulatory mandate.