AI DECISION EVIDENCE

The Evidence Layer for AI Decisions

Verifiable proof for every governed AI decision - an executive brief.

THE PROBLEM

AI now makes or shapes decisions that regulators, auditors, and counterparties will scrutinize - in lending, healthcare, insurance, and beyond. When an examiner asks "how do you know this decision happened the way you say?", a log or a dashboard is not an answer. Both ask the examiner to trust you.

WHAT EVE PROOF IS

EVE Proof turns every governed AI decision into a signed, independently verifiable certificate: the verdict, the reason code, the policy version, and a digest of the inputs - cryptographically signed (Ed25519) so anyone can confirm, offline, that it is authentic and unaltered. It is the evidence plane of the EVE control-plane stack: Governance decides, CoreGuard enforces, EVE Proof attests.

THREE PROPERTIES THAT MATTER TO AN AUDITOR

- Offline-verifiable: a certificate, the public key, and stock crypto are all that is required - no EVE service, no account, on an air-gapped machine, years later.
- Replay-verifiable: re-run the recorded inputs through the same policy version and the verdict reproduces - proof the outcome is what the engine would decide again.
- Tamper-evident: any altered field breaks the signature; removing or reordering a decision breaks the hash chain and Merkle proof.

LOGS AND DASHBOARDS VS. CRYPTOGRAPHIC EVIDENCE

A log can be edited and proves nothing about why. A dashboard shows a claim you must trust. A decision certificate can be checked by anyone, needs no vendor in the loop, and survives your company. The difference is who has to be trusted - with cryptographic evidence, the answer is no one.

WHAT IT CHANGES, BY ROLE

- Model risk & audit: confirm what a control did, when, and under which policy version - independently of the first line's attestations.
- Compliance & legal: hand a regulator a self-verifying evidence pack instead of a screenshot; answer "prove it" with a signature.
- Platform & ML engineering: evidence is emitted inline with the decision - one API call, no separate logging path to build or secure.

DESIGNED FOR THE OBLIGATIONS YOU ALREADY FACE

EVE Proof is built to support EU AI Act Article 12 and 15 record-keeping, SR 11-7 model-risk independent verification, ECOA/Reg B adverse-action evidence, HIPAA minimum-necessary decisions, and NIST AI RMF - as infrastructure for your program, not as legal advice.

HOW TO EVALUATE IT IN AN AFTERNOON

- Verify a real certificate in your browser at eveproof.com/verify - then press Tamper and watch it fail.
- Download the evaluation kit: a real signed certificate, an evidence pack, and a standalone verifier you run yourself.
- Use the open-source eve-verify package (Python / TypeScript / CLI) - no EVE dependency, no network.
- Read the verification specification and confirm the math against your own crypto library.

NEXT STEP

Book a working session that maps your highest-risk AI workflow to deterministic controls and the evidence your examiners will ask for. Start at eveproof.com/resources or contact sales@eveaicore.com.

See a real certificate verify - eveproof.com/verify

EVE Proof - the evidence layer for AI decisions. Part of the EVE control-plane stack.