AEP-002

Evidence Docket Standard

Institutional Edition v1.1 · The portable proof package for machine work

Vincent Boucher · QUEBEC.AI & MONTREAL.AI · 2026-06-05

A model can answer. An agent can act. An institution must prove.

AEP-002 makes the proof portable.

1. Executive Summary

AEP-002 defines the Evidence Docket: the standard proof package for AI-agent work, machine work, governed AI-assisted workflows, institutional AI adoption, Proof Rooms, and Agent Control Planes.

AEP-001 defines the constitutional loop: Commit → Execute → Prove → Evolve. AEP-002 defines the portable proof object that makes the loop operational.

---

2. Canonical Law

No proof, no evolution.

No eval, no propagation.

No rollback, no release.

3. What an Evidence Docket Answers

- What was the machine asked to do?
- Who or what authorized it?
- What context, tools, policies, and constraints were in scope?
- What happened?
- What evidence exists?
- Which checks or evals passed?
- What failed or remains uncertain?
- What can honestly be claimed?
- What may be reused or promoted?
- What can be rolled back?

4. Relationship to AEP-001

AEP-001 defines the constitutional architecture: Artifact Vault, Run Fabric, Proof Ledger, and Selection Gate. AEP-002 defines the evidence object produced and consumed by that architecture.

5. Core Object Model

Docket Section	Purpose
Manifest	Identifies the docket, owner, protocol, status, jurisdiction, and confidentiality class.
Claims Matrix	Separates supported claims, evidence, confidence, and what is not claimed.
Commitment Record	States goal, success criteria, constraints, risk, tools, evals, approvals, and rollback obligations.
Execution Summary	Records what happened without exposing restricted traces.
Evidence Inventory	Lists evidence IDs, types, locations, hashes, access

	classes, retention, and related claims.
ProofPackets	Atomic evidence units for claims, traces, evals, policy, risk, and outputs.
Tool-Use Ledger	Records tool calls, permissions, approvals, risk flags, and rollback possibility.
Policy and Approval Ledger	Records governance decisions, actors, reasons, timestamps, and escalations.
Evaluation Results	Records how the output was checked, scored, and limited.
Cost and Latency Ledger	Records efficiency, cost, latency, and human-review burden.
Risk Ledger	Records known risks, severity, mitigation, residual risk, owners, and status.
Selection Certificate	Records whether a capability may be promoted, canaried, rejected, revised, archived, or rolled back.
Rollout / Canary Plan	Defines scope, monitoring, stop conditions, escalation, and rollback trigger.
Rollback Plan	Explains how to undo or stop the capability.
Public-Safe Report	The shareable accountability layer.
Private Appendix	Restricted prompts, traces, logs, security notes, and protected details.
Claim Boundary	What is claimed, what is not claimed, and what remains private/protected.

6. Public / Private / Protected Evidence Boundary

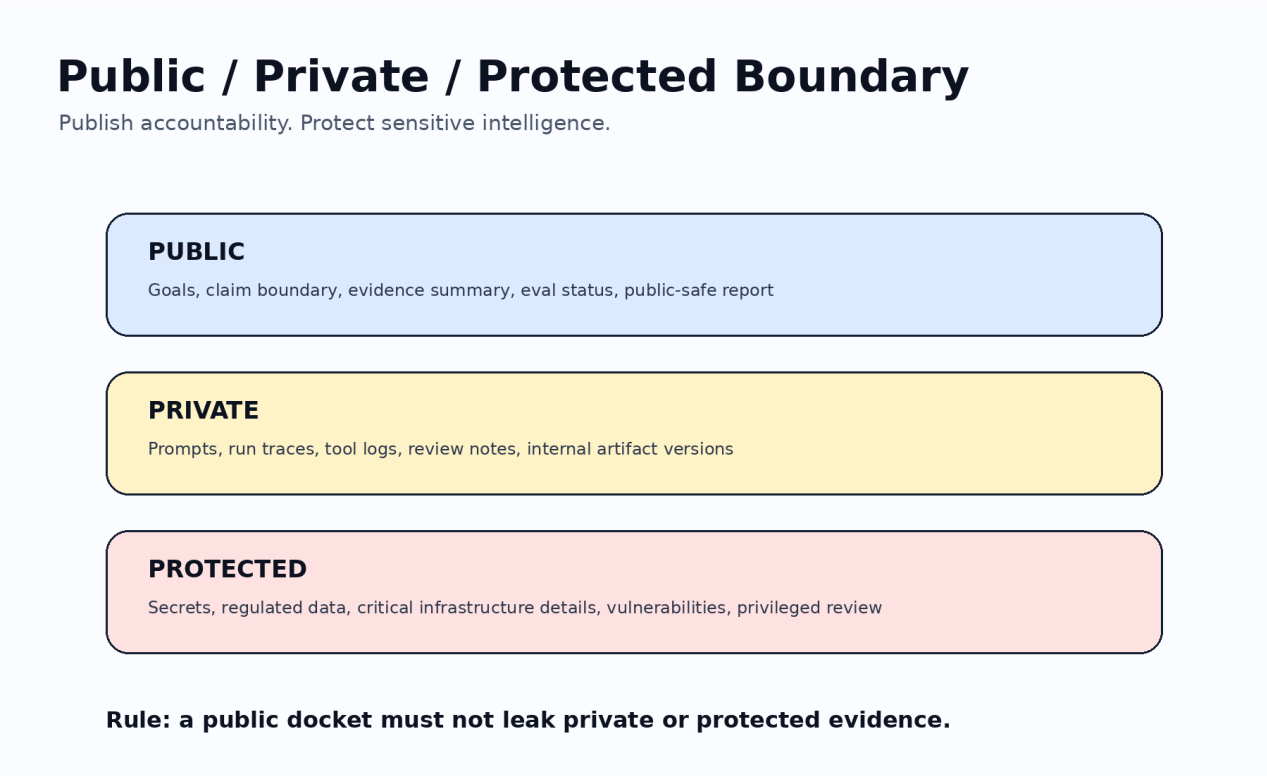

Class	May include	Publication rule
Public	High-level goals, evidence summaries, eval status, rollback status, public-safe reports.	May be shared externally after claim-boundary review.
Private	Prompts, traces, tool logs, review notes, operational cost details, internal artifact versions.	Internal only unless explicitly authorized.
Protected	Secrets, regulated data, critical infrastructure details, security vulnerabilities, privileged analysis.	Restricted to authorized roles and never published by default.

7. Lifecycle

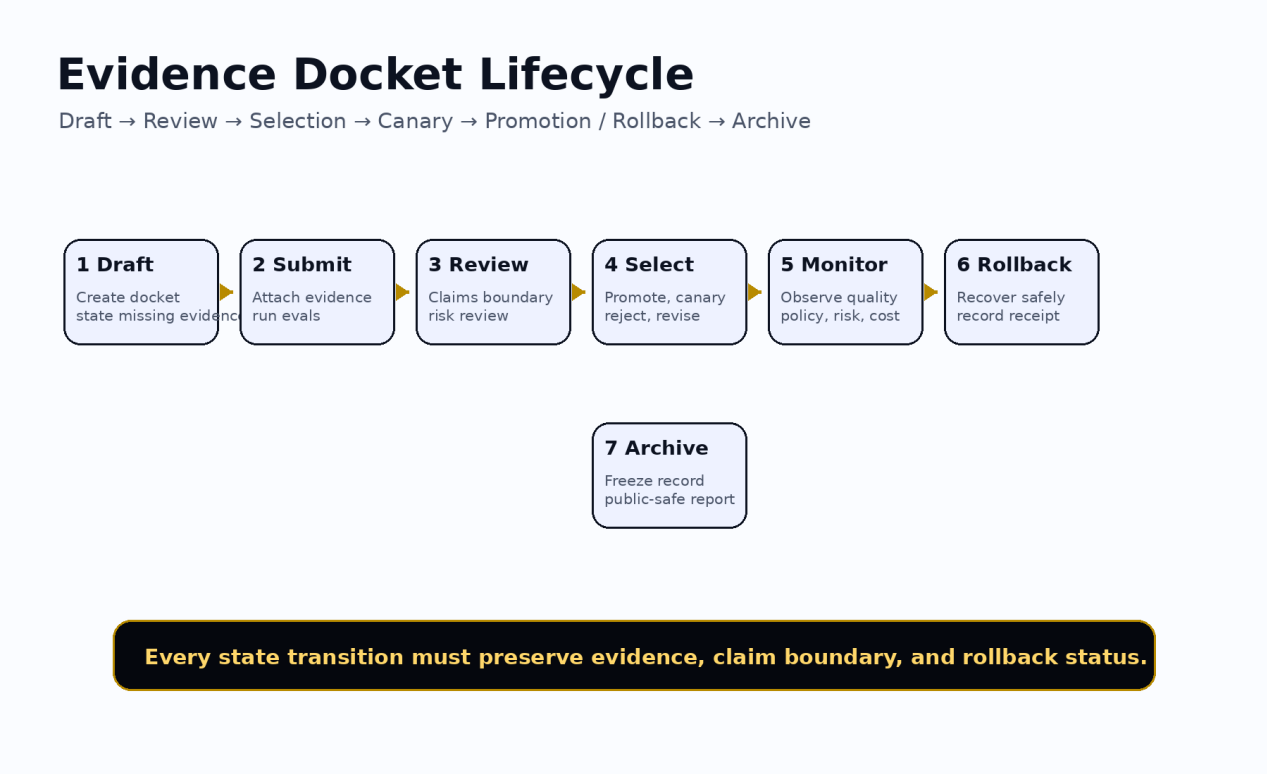

- Draft
- Submitted
- Under review
- Evidence accepted
- Selection decision issued
- Canary / monitor
- Promoted / rejected / revised / rolled back
- Archived
- Public-safe report published, if allowed

8. Conformance Levels

Level	Requirement
Level 0	Informal proof page with goal, output, evidence, checks, and claim boundary.
Level 1	Basic Evidence Docket with minimum viable sections and at least one evaluation or explicit evaluation gap.
Level 2	Operational Evidence Docket with tool-use, policy, cost, risk, selection, and rollback records.
Level 3	Institutional Evidence Docket with Selection Certificate, canary plan, monitoring, rollback receipt, and audit export.
Level 4	Sovereign / regulated docket with public/private/protected classification, retention, jurisdiction, authorized boundary,

evaluator attestations, and restricted appendix controls.

9. Implementation Guidance

- Add a docket template.
- Generate one docket per important AI workflow.
- Record claims separately from evidence.
- Store private evidence separately from public proof.
- Run at least one eval.
- Create a rollback plan.
- Require Selection Gate review before propagation.
- Publish a public-safe report only after claim-boundary review.

10. Security and Privacy Requirements

- Do not publish secrets.
- Do not publish private prompts unless authorized.
- Do not publish customer data unless authorized and safe.
- Do not publish protected operational traces.
- Do not publish critical infrastructure details.
- Do not imply more than the evidence supports.

11. Claim Boundary

AEP-002 does not claim achieved AGI, achieved ASI, perfect safety, legal compliance certification, financial or legal advice, guaranteed ROI, production readiness, government endorsement, or national-security readiness. AEP-002 defines a proof package standard.

AEP-002 makes the proof portable.