

Compliance and governance

Where AiSSAY stands, 19 September 2026

Build V18.6. Every capability named here maps to a module and a test in the code as it stands today. Regulatory dates were checked against published sources on 19 September 2026 and are listed at the end.

The short version

AiSSAY is a record-keeping and evidence layer for AI agent work. It runs on the customer's own machines, under the customer's own key, and we keep no copy. It is built to help a firm answer the questions a regulator, an auditor or a court will ask about work an AI agent touched.

Where we stand:

- **The record layer is strong and checkable.** Tamper-evident sealing, access logging, retention, legal hold and deletion that leaves proof are all built, tested, and evidenced from the customer's own records.
- **Keys can be retired.** A recorded ceremony replaces a store's key and moves every record under the new one. The store can hand its key to the customer's own key manager — as a tested seam. An AWS KMS backend is built and unit-tested, and has not yet run against the real service.
- **We map to thirteen frameworks** with evidence packs, including views written to the rules of law, audit, healthcare, broker-dealers and education. Each shows the evidence from the customer's store, and says where a requirement is the customer's own policy or something we have not built.
- **Nothing has been independently checked.** No penetration test, no third-party review of the cryptography, no audit, no certification. Those need someone outside the company, and they are next on the list.
- **Nothing below has run in a customer's environment yet.** It has been built and tested by us, not by a customer, an auditor or a regulator.

What the product does for a compliance team

A firm using AI agents has three problems it cannot solve with the agent's own logs:

1. **Proving what happened.** The agent's log is written by the system being examined. AiSSAY keeps an independent record, sealed so that any change afterwards shows.
2. **Proving who looked.** Every read, export and deletion is itself recorded in a tamper-evident log.
3. **Answering the examiner in their own format.** One command turns the stored records into an evidence pack laid out in the framework's own structure.

Because it is self-hosted, the vendor is outside the trust boundary. We cannot be compelled to produce a customer's records, because we do not hold them. The only network calls on the recording path go to the model vendors the customer already uses, on the customer's own keys.

Where we stand, control by control

Built and evidenced from the customer's own records

Area	What exists	How it is shown
Integrity	Each record encrypted under its own key and sealed into a hash chain; external anchoring; verification names the altered record	Verification run on demand; chain and anchor checks in every pack
Access logging	Every read, export, hold and deletion logged in its own tamper-evident chain; reads can be cryptographically signed by the reader	Access-log depth and attribution in the pack
Retention	Per-record retention dates. Deleting before expiry is refused unless explicitly overridden with a reason, which the record of the deletion keeps	Retention spans in the pack
Legal hold	Holds block deletion and override retention	Holds counted in the pack
Deletion	Cryptographic erasure; a permanent tombstone proves the record existed	Tombstones counted in the pack
Erasure requests	A lineage map says what was made from what, so an erasure request can be scoped	Lineage report
Key rotation	A re-key ceremony: new key sealed, every record moved, old key destroyed only after every record reads back; old grants and custody shares stop working	Evidenced in the pack once a ceremony has run on that store
Irreversible actions	Payments, filings and sends are held mid-run until the firm's policy answers	Hold recorded as a decision step in the record
Injection	Instructions hidden in pasted documents are scanned for on three surfaces	Findings in the record

Built, and shown at release rather than in the customer's store

Area	What exists
Dependency scanning	Every Python package we ship is checked against the public vulnerability database at release, and the report ships with its hash. The first scan found five published advisories in our cryptography library; the pin moved the same day and the rescan is clean.
Software bill of materials	Generated from the environment the tests pass in, per release.
Pinned dependencies	Every test suite passes in a fresh install built only from the pinned list.
Supply-chain claims	51 written claims about what the code does and does not do, each checked against the code by a script that fails the build if one breaks.
Network audit	Generated from the code: every module that can reach the network, and what it sends. No telemetry.
Security contact	A published disclosure policy (aissay.io/security.html) and a security.txt file. The build fails 30 days before the file expires, so the contact cannot quietly go stale.

Built as a seam, not yet as a product

Area	Where it stands
Customer key manager	The store can hand its key to the customer's key manager (AWS, Azure, Google Cloud, Vault). The interface is two operations, failures are typed, and nothing falls back to a local key. An AWS KMS backend is built and unit-tested; it has not yet run against the real service , and an acceptance run against real AWS is written and waiting. Other key managers get a backend when a customer names one. Until then the packs mark access control as a "test seam", not as delegated.
Split custody	The key can be split so that K of N named people are needed to reconstruct it. Composing that into multi-party custody of the store is designed, not built.

Not built, and why

Gap	Why it is open	What closes it
Penetration test	Needs an outside firm	An engagement, scoped to the self-hosted build
Independent crypto review	Needs an outside reviewer	A review of the key handling, chain and receipt design
External audit or certification	None has been sought	See the roadmap below
Access control inside the store	By design, access is delegated to the customer's key manager rather than rebuilt by us	The first real key-manager backend
Vulnerability remediation, OS packages, continuous monitoring	The release scan covers identification of Python dependency issues only	Written procedure; scanning in CI; OS scanning is the customer's

We think naming these is a strength in a review, not a weakness. The packs name their own gaps before a reviewer finds them.

The thirteen evidence packs

Framework	Rows	Gaps we name	Rows that are the customer's policy
SEC Rule 17a-4 (books and records)	5	0	0
EU AI Act Article 12 (record-keeping)	4	0	0
EU AI Act Article 14 (human oversight)	9	2	0
US bank model-risk guidance (SR 26-2)	4	0	0
NIST AI Risk Management Framework	11	1	3
SOC 2 security criteria	12	1, plus access control unless a real key manager seals the store	3
ISO/IEC 27001:2022 Annex A	12	1, plus access control on the same terms	4
ISO/IEC 42001 (AI management)	9	1, plus access control on the same terms	2

Framework	Rows	Gaps we name	Rows that are the customer's policy
Law: ABA Formal Opinion 512 and the Model Rules it applies	10	1	4
Audit documentation: PCAOB AS 1215 and AICPA AU-C 230	8	1	1
Healthcare: HIPAA Security Rule	9	2, plus access control on the same terms	1
Broker-dealers: FINRA rules applied to AI	8	1	1
Education: FERPA	6	1, plus access control on the same terms	2

Every pack carries the same statement on page one: it is an informational mapping prepared by the operator's own tooling, not a legal opinion, not a certification and not a substitute for counsel. Every status is computed from the store, never asserted, and a requirement with no records behind it yet reads as "unevidenced", not as met.

Things worth knowing about these frameworks:

- **SR 11-7 no longer exists as current guidance.** The banking agencies replaced it on 17 April 2026 with revised guidance (SR 26-2), which puts generative and agentic AI models outside its scope for now and promises a request for information on AI. Our pack now says so. For agentic systems it is voluntary alignment, not a supervisory requirement.
- **NIST AI RMF has no certification.** The only accurate claim is "aligned", self-assessed. Our pack says that in its own text.
- **SOC 2 and ISO certify an organisation, not a product.** A customer's auditor can use our packs as evidence. The packs cannot make anyone compliant.
- **ABA opinions are advisory.** Opinion 512 binds a lawyer only as their own state adopts it, and client-file retention periods are set state by state.
- **Some of the rules are moving.** The HIPAA Security Rule update proposed in December 2024 was not final by September 2026, and the law and FINRA views say which of their rows are guidance rather than rules.
- **One gap is the same in every view.** No one outside the company has evaluated this software yet, so each industry view names that first.

The regulatory calendar that matters

Date	What	Why it matters to us
27 Jul 2026	The EU "Digital Omnibus on AI" entered into force	It moved the high-risk dates below
2 Aug 2026	EU AI Act transparency duties (Article 50) applied	Not our core use, but the first live EU AI obligations
15 Dec 2026	PCAOB QC 1000 (quality control) takes effect	Audit firms' quality systems; our audit-documentation view names it as the firm's
1 Jan 2027	Colorado's replacement AI law takes effect	The original Colorado AI Act was repealed and re-enacted in May 2026, with narrower deployer duties

Date	What	Why it matters to us
2 Dec 2027	EU AI Act high-risk obligations for stand-alone (Annex III) systems	Logging (Art. 12), human oversight (Art. 14) and deployer duties (Art. 26) start to bite. This is our main EU timing window
2 Aug 2028	High-risk obligations for AI in regulated products (Annex I)	Later and narrower
Pending	US banking agencies' request for information on AI	Where agentic AI in banks will be addressed after SR 26-2

The EU delay gives firms fourteen more months before Articles 12, 14 and 26 apply to stand-alone high-risk systems. Records kept from now on will be the ones an examiner can ask for then.

What comes next

Now

1. **A threat model and a scoped brief for an outside review**, so a penetration test and a cryptography review start from our own map of the system rather than from scratch.
2. **A written policy set** — information security, incident response, access review, vendor management, secure development. These are the "organizational" rows in our own packs, applied to ourselves.
3. **A standard security-questionnaire answer pack** for the questionnaires banks and law firms send, built from the security, supply-chain and deployment documents that ship with the code.
4. **The dependency check and the full test suite on every change**, not only at release.
5. **Signed releases**, so a customer can verify the code they install is the code we shipped.

With our first deployments

6. **A customer's own key manager in production**: the AWS KMS backend once its run against the live service is recorded, or a backend for the key manager the customer names, tested the same way. A store sealed by a real key manager reads "delegated" on its access-control rows.
7. **An independent penetration test** of the self-hosted build in a realistic deployment.
8. **An independent review of the cryptography**: key handling, the chain, the re-key ceremony and the selective-disclosure receipts.
9. **Multi-party custody of the store**, so no single person can open it. The design is written, and the ceremony format already has room for it.

Later

10. **SOC 2 Type I, then Type II, for AiSSAY as a company**. Because the product is self-hosted, the scope is small: how we build, release and support it.
11. **ISO/IEC 27001, and alignment with ISO/IEC 42001** for our own use of AI in development.
12. **Mapping the packs to the EU harmonised standards** for high-risk logging and oversight as they are published.
13. **Taking part in the US banking agencies' request for information on AI** when it is issued.

Sources

- Digital Omnibus on AI in force, 27 July 2026: [Lewis Silkin](#); [European Commission](#)
- New high-risk dates and Article 50 timing: [Gibson Dunn](#); [Jones Walker](#)
- SR 26-2 replacing SR 11-7, AI excluded from scope: [Orrick](#); [OCC Bulletin 2026-13](#)
- Colorado repeal and re-enactment: [Crowell & Moring](#); [Seyfarth](#)

This document describes software in development. It is not legal, audit or compliance advice, and nothing in it says any organization is or will be compliant with anything. A framework-by-framework table of what each evidence pack covers is published alongside it.