

Compliance readiness coverage

What an AiSSAY™ evidence pack covers, framework by framework and industry by industry — and what it refuses to claim.

AiSSAY, Inc. · aissay.io · September 22, 2026 · engine V18.6 · generated from the engine's own requirement tables

What this document is

AiSSAY keeps sealed, hash-chained records of AI work on your own machines, and computes readiness packs from those records: for each framework below, a table of requirements with a status computed from your store at the moment you ask. This document lists the requirement tables the engine ships today, and it is generated from those tables. It describes what a pack covers, not what any organization has achieved.

It is not a certification, an audit opinion, or legal advice. A pack is evidence you hand to your auditor, your regulator or your own review — computed, dated, and honest about its limits.

How a status is earned

Evidenced	a deterministic check ran against your records and the counts support the row. A store holding no relevant records reads unevidenced — a true statement, not a passing one. Empty never counts as clean.
Delegated	the mechanism is enforced by a system you already trust — a store sealed by your own key manager delegates access control to it, and the pack says the enforcing log is your key manager's, not ours. The pack names the endpoint the key manager was reached at. An AWS endpoint other than AWS's own default for the region reads delegated to that endpoint, with a request that you confirm it is your key management service. One backend — AWS KMS — has run against the real AWS KMS service; no other backend has.
Test seam	sealed by an in-process test backend — a stub, or a stand-in for the AWS client — and the status names which. No key manager enforces access. The seam is built and tested, and the pack refuses to call that delegation.
Partially evidenced	the mechanism is built and some of your runs carry its evidence, but not all — usually because they were sealed before the mechanism existed. The pack says how many runs are in each group, and never counts the missing ones as passing.
Organizational	the requirement is a procedure only people can perform. The pack names it as yours instead of dressing a policy up as a mechanism.
Named gap	the pack cannot evidence it — either the mechanism is not built, or no software can supply it — and it says which. The gap is printed, because a reviewer should read it here rather than discover it.

No status is typed in by hand: every status is recomputed from the store each time a pack is produced, the requirement and mechanism text is written and reviewed as part of the engine, and the dependency evidence comes from our release scan. A failing mechanism can never be read as evidence that the control works.

Rows marked “computed” are the access-control rows. Their status is computed from the provider your store declares: a gap unless your key manager seals the store, “test seam” under a stub or stand-in backend, “delegated” only under a real one. Today they can never read evidenced, and marking them as a built mechanism would overstate them.

By industry

Every industry gets the general security and AI governance views. Where its own regulators set rules for records or for AI, it gets a view written to those rules too.

Industry	Views that apply
----------	------------------

Law	Legal ethics (ABA Formal Opinion 512); SOC 2, ISO/IEC 27001, ISO/IEC 42001, NIST AI RMF
Accounting and audit	Audit documentation (PCAOB AS 1215, AICPA AU-C 230); SOC 2, ISO/IEC 27001, ISO/IEC 42001, NIST AI RMF
Healthcare	HIPAA Security Rule; SOC 2, ISO/IEC 27001, ISO/IEC 42001, NIST AI RMF
Banking	US bank model risk guidance (SR 26-2); SOC 2, ISO/IEC 27001, ISO/IEC 42001, NIST AI RMF
Trading and broker-dealers	SEC Rule 17a-4; FINRA; SOC 2, ISO/IEC 27001, ISO/IEC 42001, NIST AI RMF
Education	FERPA; SOC 2, ISO/IEC 27001, ISO/IEC 42001, NIST AI RMF
Engineering	No engineering-specific view yet; SOC 2, ISO/IEC 27001, ISO/IEC 42001, NIST AI RMF
Any firm with EU high-risk AI	EU AI Act Articles 12 and 14, in addition to its industry's view

Financial services

For broker-dealers, trading desks, banks and model-risk functions: records that cannot be altered after the fact, supervision evidenced in writing, and model documentation an independent reviewer can challenge.

SEC Rule 17a-4 — books and records (electronic storage posture)

Records preserved in a manner that prevents alteration (WORM / audit-trail alternative)	mechanism
Retention for the required period, readily accessible	mechanism
An accurate index of records	mechanism
Audit system providing accountability for access	mechanism
Ability to furnish records promptly	mechanism

FINRA rules applied to generative AI and AI agents — broker-dealers

Regulatory Notice 24-09 (27 June 2024) says existing rules apply to generative AI and creates no new requirements. The 2026 Annual Regulatory Oversight Report's AI-agent practices are guidance, not rules. Notice 26-14 (July 2026), on supervising retail communications including AI-generated ones, was a proposal as of September 2026. SEC Rule 17a-4 has its own view.

Rule 4511 — make and keep books and records; six years where no period is specified	mechanism
Rule 3110(a) — a supervisory system reasonably designed for compliance	organizational
Rule 3110(b)(4) — principal review of correspondence, evidenced in writing with the reviewer and the date	mechanism
Rule 2210 — communications, including AI-generated ones: approval where required, fair and balanced content, retained with dates	mechanism
Regulatory Notice 24-09 — protect customer information when using generative AI	mechanism
2026 Oversight Report (guidance) — for AI agents, keep a human in the loop, track actions and decisions, keep prompt and output logs	mechanism
2026 Oversight Report (guidance) — track which model version an agent used	mechanism
Reliance on this tool — independent evaluation of its controls	named gap

US bank model risk guidance — SR 26-2, which superseded SR 11-7 on 17 April 2026 — model risk evidence

SR 26-2 (Federal Reserve, OCC and FDIC, April 2026) replaced SR 11-7 and places generative and agentic AI models outside its scope for now; the agencies said a request for information on AI would follow. This view maps agent records to the model-risk principles both letters share: monitoring, effective challenge, documentation and records for audit. For agentic systems it is voluntary alignment, not a supervisory requirement.

Ongoing monitoring of model performance	mechanism
Effective challenge, independent of the model	mechanism
Documentation sufficient for parties unfamiliar with the model	mechanism
Records supporting validation and audit	mechanism

Law

For firms whose lawyers use AI on client work: the Model Rules as ABA Formal Opinion 512 applies them, and the client file.

US legal ethics and generative AI — ABA Formal Opinion 512 (29 July 2024) and the Model Rules it applies

ABA opinions are advisory. They bind a lawyer only as the lawyer's own state adopts the rule or position. Client-file retention periods are set state by state; Model Rule 1.15 brackets five years for trust records as a placeholder. No later ABA formal opinion on AI had issued by September 2026.

Model Rule 1.1 — competence: understand what the AI tool can and cannot do, and keep that understanding current	organizational
Model Rule 1.6 — protect client information used with AI tools	mechanism
Model Rule 1.6 — informed client consent before confidential information goes into a self-learning tool	organizational
Model Rules 3.1, 3.3 and 8.4(c) — AI output checked before it reaches a tribunal	mechanism
Model Rules 5.1 and 5.3 — supervision of lawyers, staff and vendors using AI	mechanism
Model Rules 1.4 and 1.5 — telling clients about AI use, and reasonable fees for AI-assisted work	organizational
Model Rule 1.16(d) — surrender the papers and property the client is entitled to when the representation ends	mechanism
Formal Opinion 520 (January 2026) — respond to information requests from a former client or successor counsel, which can include information that was never written down	organizational
Client-file and trust-record retention (state rules)	mechanism
Reliance on this tool's records as evidence — independent evaluation of its controls	named gap

Accounting and audit

For audit firms: documentation that shows who did the work and who reviewed it, kept for the required period and never silently changed after completion.

Audit documentation — PCAOB AS 1215 (issuer audits) and AICPA AU-C 230 (non-issuer audits)

AS 1215 as amended by AS 1000: the documentation completion date is 14 days after the report release date, phased in for fiscal years beginning on or after 15 December 2024 or 2025 depending on firm size (it was 45 days). QC 1000 takes effect 15 December 2026; amendments adopted 9 September 2026 await SEC approval. The technology-assisted analysis amendments to AS 1105 and AS 2301 apply to fiscal years beginning on or after 15 December 2025.

AS 1215.06 — documentation shows who performed the work and who reviewed it, and when	mechanism
AS 1215.14 — retain audit documentation for seven years from the report release date	mechanism
AS 1215.15 — assemble the complete file by the documentation completion date	mechanism
AS 1215.16 — after completion nothing is deleted or discarded; anything added records when, by whom and why	mechanism
AS 1105 and AS 2301 — evaluate the reliability of electronic information and follow up items the analysis flags	mechanism
AU-C 230 — retain documentation at least five years, and record who changed what after assembly, when and why	mechanism
QC 1000 — the firm's system of quality control, and its documentation	organizational
Reliance on this tool's records as audit evidence — independent evaluation of its controls	named gap

Healthcare

For covered entities and business associates: the HIPAA Security Rule's technical safeguards over systems that hold ePHI.

HIPAA Security Rule — technical safeguards for systems holding ePHI (45 CFR Part 164, Subpart C)

This view maps the rule in force. The Security Rule update proposed in December 2024 was not final as of September 2026. A business associate agreement, the risk analysis and workforce training are the covered entity's. Records stay on the covered entity's own machines; we hold no copy.

164.312(a)(1) — access control: only authorized persons or software reach ePHI	computed
164.312(a)(2)(iv) — encryption and decryption of ePHI (addressable)	mechanism
164.312(b) — audit controls: record and examine activity in systems containing ePHI	mechanism
164.312(c)(1) — integrity: protect ePHI from improper alteration or destruction	mechanism
164.312(d) — person or entity authentication	mechanism
164.308(a)(1)(ii)(D) — regularly review information system activity	organizational
164.308(a)(6)(ii) — identify, respond to and document security incidents	mechanism
164.316(b)(2)(i) — retain required documentation six years	mechanism
Reliance on this tool — independent evaluation of its controls	named gap

Education

For schools and the providers they engage: FERPA's conditions on school officials and its record of disclosures.

FERPA — student education records (34 CFR Part 99)

No FERPA regulation specific to AI had issued by September 2026. The Department of Education's July 2025 letter says AI use must comply with FERPA. Access by a provider acting as a school official is not a disclosure that 99.32 requires the school to record; this store logs it anyway.

99.31(a)(1)(i)(B) — an outside provider acts as a school official only under the school's direct control over use and maintenance of records	organizational
99.31(a)(1)(ii) — reasonable methods so officials reach only records they have a legitimate educational interest in	computed
99.32(a) — record each request for and disclosure of personally identifiable information, with the requester and their interest	mechanism
99.32(b)–(c) — record redisclosures; parents, eligible students and auditors may inspect the record of disclosures	mechanism
99.33 — limits on redisclosure by recipients	organizational
Reliance on this tool — independent evaluation of its controls	named gap

EU-regulated and high-risk AI

Articles 12 and 14 place their duties mainly on the provider of a high-risk AI system: logging capability built in, oversight measures built in. Deployers carry their own duties under Article 26 — keeping the logs the system produces, assigning competent oversight. These tables map what the records evidence for whichever role holds the duty.

EU AI Act Article 12 — record-keeping for high-risk AI systems

Automatic recording of events over the system lifetime	mechanism
Logging appropriate to the intended purpose, enabling identification of risk situations	mechanism
Records kept for an appropriate period	mechanism
Traceability of the system's functioning	mechanism

EU AI Act Article 14 — human oversight of high-risk AI systems

Oversight measures built into the system, enabling a natural person to oversee its operation	mechanism
The person can decide not to use the output, or to disregard, override or reverse it	mechanism
The person can intervene or interrupt the system	mechanism
The person is enabled to remain aware of automation bias	mechanism
The person can correctly interpret the system's output	mechanism
Oversight is exercised by a natural person, not another automated system	mechanism
Article 14(1) — oversight measures commensurate with the risks of the system	mechanism
Article 14(4) — the overseeing person has the necessary competence, training and authority	mechanism
Hold and override counts per run	mechanism

Security and audit

For SOC 2 and ISO/IEC 27001 programs: integrity, access accountability, key custody and dependency hygiene, with the organizational clauses named as organizational.

SOC 2 Trust Services — Security (Common Criteria) readiness view

SOC 2 is an attestation of ORGANIZATIONAL controls operating over time, issued by a CPA firm. Code cannot be SOC 2; this view shows which criteria the shipped mechanisms evidence and which require organizational controls no software supplies.

CC3.x — risk assessment process	organizational
CC6.1(c) — attribution of access to identified individuals	mechanism
CC6.1 — logical access restricted to authorized users	computed
CC6.1(b) — access authorization and periodic review	organizational
CC6.6 — protection of data at rest	mechanism
CC7.2 — monitoring for anomalies and security events	mechanism
CC7.3/7.4 — evaluation and response to security events	mechanism
CC8.1 — change management	mechanism
A1.2 (Availability) — backup and recovery	organizational
CC6.7 — cryptographic key management over time	mechanism
CC4.1 — independent evaluation of controls	named gap
C1.1 (Confidentiality) — confidential information protected	mechanism

ISO/IEC 27001:2022 Annex A — control readiness view

ISO 27001 certifies an information security management system, audited by an accredited body. This view maps shipped mechanisms to Annex A themes and names the ISMS gaps plainly.

A.5.1 — information security policies	organizational
A.5.19–5.22 — supplier relationships	organizational
A.6 — people controls (screening, training, NDAs)	organizational
A.8.15 — logging	mechanism
A.8.16 — monitoring activities	mechanism
A.8.24 — use of cryptography	mechanism
A.8.13 — information backup	organizational
A.8.32 — change management	mechanism
A.5.15/A.8.3 — access control and information access restriction	computed
A.8.24 — use of cryptography, including key lifecycle	mechanism
A.8.8 (identification) — known vulnerabilities in Python dependencies	mechanism
A.8.8/A.8.29 — remediation, monitoring and security testing	named gap

AI governance

For AI risk and management-system programs: mapping AiSSAY's records onto NIST AI RMF and ISO/IEC 42001, where an evidence store speaks and where only management can.

NIST AI Risk Management Framework 1.0 — alignment mapping

NIST AI RMF is a voluntary framework. No NIST certification exists and none can be claimed; the accurate public statement is that the operator ALIGNS with the framework, self-assessed, with this mapping regenerable from the evidence store at any time.

GOVERN 1.1–1.4 — policies, accountability structures, and risk tolerance documented across the AI lifecycle	organizational
GOVERN 2.1 — roles and responsibilities for AI risk are defined	organizational
GOVERN 4.3 — incidents and errors are shared and documented	mechanism
MAP 1.1/2.3 — context, purpose and assumptions of AI use documented	mechanism
MEASURE 2.5 — AI system performance monitored with metrics	mechanism
MEASURE 2.7(a) — adversarial evaluation of the AI system	mechanism
MEASURE 2.7(b) — security and resilience of the deployed system	named gap
MEASURE 2.8 — transparency and accountability records maintained	mechanism
MEASURE 3.1–3.3 — risks tracked over time; measurement itself is assessed	mechanism
MANAGE 2.3/2.4 — mechanisms to respond to and recover from incidents	organizational
MANAGE 4.1 — post-deployment monitoring with human oversight	mechanism

ISO/IEC 42001:2023 — AI management system readiness view

ISO 42001 certifies an AI management system. The operational clauses are where an evidence store speaks; the management clauses are organizational and say so below.

Cl. 4–5 — organizational context, leadership, AI policy	organizational
Cl. 6 / A.5 — AI risk assessment and treatment	mechanism
A.6 — AI system lifecycle: verification and validation records	mechanism
A.7 — data for AI systems: provenance and quality records	mechanism
A.8 — information for interested parties: incident reporting	mechanism
A.9 — human oversight of AI systems	mechanism
Cl. 9–10 — performance evaluation and improvement	organizational
A.4.2 — resources: access control for AI system data	computed
A.10 — third-party and independent assurance	named gap

More on demand

A framework here is a requirement table: each row maps a requirement to a mechanism in the engine, an organizational duty, a computed access-control status, or a named gap. Adding one your reviewers ask about — a state bar's rules, an engineering licensing board, an internal model-governance policy — is adding a table, and every new row gets the same treatment: a status computed from your store, or an honest statement of whose duty it is. If a framework you need is not listed, ask.

What we never claim

- We do not certify compliance. Certification belongs to auditors and regulators; we compute the evidence they read.

- We never mark an empty store as protected. A control with nothing to protect is unevidenced by definition.
- We never present a procedure as a mechanism. If people have to do it, the pack says so.
- Named gaps stay printed until the mechanism ships and its tests pass. A capability sentence in our materials maps to a module and a test, or it is removed.
- Your records stay on your machines. Packs are computed where the records live; we hold no vendor-side copy.