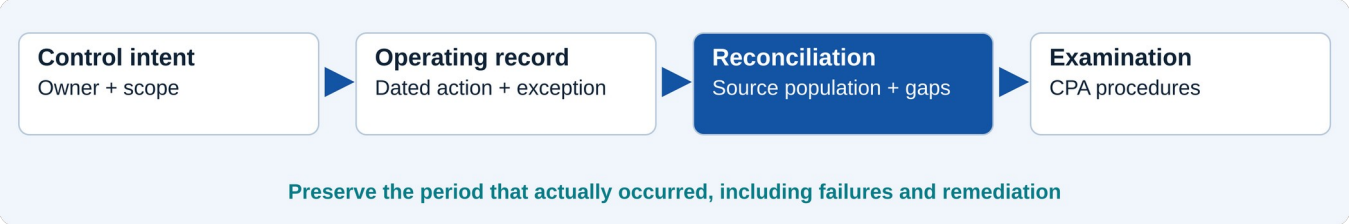
SOC 2 with AI systems in scope

Defining the service boundary and producing evidence that can be examined

Anwesh Rath, Founder and CTO, AI Standards Inc.



Conceptual evidence chain. Operating records support examination within the agreed scope.

Executive summary

Adding AI to a service can change what must be described, controlled and evidenced in a SOC 2 examination. The important boundary includes the data pipeline, retrieval layer, model service, tools and people who operate them. A report on the underlying hosting service does not establish that the complete AI service is covered.

We recommend starting with the service commitments and system description, then tracing the controls that support them. Collect evidence from real operation with known populations and dates. Keep model evaluation and control effectiveness distinct. Useful evidence should let an independent examiner understand what was supposed to happen, what actually happened and how exceptions were handled.

The scope follows the service

When a customer asks whether an AI assistant is covered by SOC 2, identify the service, report type, categories and responsibility boundary. A Type 1 report addresses the system description and control design as of a specified date. A Type 2 report also addresses operating effectiveness over a specified period [7]. A model vendor’s report cannot answer those questions for the entire workflow.

Define the service in operational terms. Include the interface customers use, the records the assistant retrieves, the outputs it creates and any systems it can change. Identify the people who approve access, change configurations, review exceptions and respond to incidents. A retrieval-index update can change the service’s behavior even when the model identifier stays the same.

Treat customer-controlled infrastructure with the same discipline. The contract and operating model determine who administers the environment and supports the commitments. Hardware location does not decide the examination boundary. Record the allocation of responsibilities before assuming that a control is inherited or outside scope.

Criteria and current proposal status

AICPA’s 2017 Trust Services Criteria, with revised points of focus from 2022, address security, availability, processing integrity, confidentiality and privacy [1]. Use the applicable Description Criteria to prepare the system description [5]. SOC 2 is an examination by a licensed, independent CPA [2]. AICPA’s October 2026 proposed criteria update adds AI considerations but remains a proposal [2]. Agree the applicable criteria and scope with the CPA.

An examination of controls does not certify a product or promise that every generated answer is true. Translate service commitments into controls over the relevant process, and evaluate model limitations in context.

Build evidence from the operating system

For each control, identify its owner, trigger, population and retained record. If the control reviews access grants, the population is all relevant grants during the period, not the examples someone remembered to save. Reconcile the evidence source to the underlying system so omissions can be detected. Include emergency changes and exceptions, not only the ordinary approval path.

An AI configuration record should identify the model version where available, retrieval sources, permissions, tool definitions and release decision. Where the provider supplies only an alias, retain the identifier, request time, relevant change notices and limits of reconstruction. Define how changes become visible. Preserve enough history to explain a sampled run without retaining unnecessary sensitive content or claiming exact replay.

Control mapping

The table illustrates engineering evidence for selected security outcomes [1]. The engagement team must determine the complete criteria mapping and examination procedures.

Requirement or outcome	Control in the AI service	Evidence available for examination
Authorized access	Approve scoped grants and periodically review current access.	Complete grant population; approvals; revocations; review outcomes.
Controlled changes	Assess changes to models, prompts, retrieval and tools before release.	Versioned change record; evaluation; approver; deployment receipt.
Detection and response	Monitor relevant events and assign incident ownership.	Alert population; triage decisions; escalation and recovery records.
Supplier responsibilities	Assess dependencies and document who performs each control.	Supplier review; applicable report scope; responsibility allocation.
Reliable evidence	Protect records and reconcile collection coverage.	Coverage checks; retention settings; access history; missing-record exceptions.

Keep supplier evidence within its boundary

Inspect a supplier report for the service used, covered period, relevant controls and exceptions. Identify complementary controls the supplier expects the customer to operate, and assign an owner and evidence source to each relevant control. State how subservice organizations are treated in the system description and how management evaluates dependencies outside the examination boundary [3].

Avoid copying restricted customer or supplier material into an unrestricted AI tool to summarize it. Apply the organization’s data-handling rules to the examination process itself. Evidence collection should improve the ability to assess controls without creating another uncontrolled repository.

Make exceptions usable. If an approval was missing, record the event, exposure, corrective action and owner. A newly installed control cannot manufacture evidence that it operated earlier. Remediation history can be valuable evidence; replacing the history with a clean screenshot destroys that distinction.

Prepare the examination without manufacturing history

Begin with a scoping session between engineering, the service owner, security and the CPA. Describe the AI service and its dependencies in language that matches actual operation. Identify the commitments the selected controls support, the evidence that exists and any claim the available records cannot substantiate.

Next, test the collection process. Select records from the underlying change and access populations and follow them into the evidence repository. Check timestamps, approver identity and deployment state. Investigate gaps before relying on a dashboard count. For an examination covering a period, retain evidence throughout that period; a late collection effort cannot substitute for missing operating records.

Before fieldwork, reconcile the system description to the deployed service. Review material changes and unresolved exceptions with the examiner. Keep management responsible for its assertions and remediation. Tooling can organize evidence, but it cannot supply an independent opinion on its own.

Illustrative retrieval change

This is an invented preparation exercise, not a customer examination. A document assistant adds a new repository to its retrieval index. The change introduces records with narrower access permissions than the original collection. The model and customer interface remain unchanged, so a release process focused only on model versions could miss it.

The expected evidence includes the change request, data classification, permission tests, release approval and deployment record. Test permitted and denied retrieval, a recent access revocation, and any cache or shared index that could preserve earlier access. Retain failed tests and remediation records. A later passing test does not establish that the original release was correctly controlled.

Then follow a sampled output back to its configuration and relevant sources. Confirm which records exist and which do not. A lineage entry may help identify an inference event, while an access approval or change decision lives in a different record system. The examination needs the applicable evidence for each control, not a single universal audit artifact.

Limits and next action

Control evidence can be incomplete, and service commitments can exceed what the architecture supports. A clean sample does not establish that the full population was collected. Model performance may also vary after a release. Start by reconciling one material change and one access decision end to end. Use the findings to improve the control and its evidence before expanding the scope of any assurance claim.



Illustrative change evidence. Separate records support different control assertions.

AI Standards engineering approach

AI Standards' 25 August 2026 engineering snapshot documents scoped operator grants, server-enforced gates in manual deployment workflows, and a governed inference path that requires lineage recording [4]. These are concrete subjects for control design and evidence collection. The applicable examination still needs a defined system, service commitments and evidence from the configuration and period in scope.

Our filed provisional patent applications address tamper-evident AI decision provenance and continuity of evidence across multi-agent workflows [6]. The design work connects recorded execution steps to agent identity and model-integrity evidence, with the aim of making model-state changes visible while preserving the workflow's history. These designs focus on traceability and change detection. Their implementation and evidence coverage must be tested in the deployment being examined.

For an AIS deployment, agree the service boundary, responsible operators and evidence access. Reconcile access and change populations, record collection gaps and retain operating evidence for the applicable period. Provenance records should link to the approvals, supplier records and incident evidence needed for each control. No SOC 2 report, certification or assurance opinion for AIS or a customer deployment is asserted in this paper.

References

Public sources checked 5 November 2026. Company sources are identified separately below.

[1] AICPA, 2017 Trust Services Criteria with revised points of focus 2022; resource dated 30 September 2023. https://www.aicpa-cima.com/resources/article/2017-trust-services-criteria-with-revised-points-of-focus-2022

[2] AICPA, proposed SOC 2 criteria update announcement, 28 October 2026; proposal status. https://www.aicpa-cima.com/news/article/aicpa-proposes-update-of-soc-2-r-trust-services-criteria-to-enhance-the-relevance-usability-and-value-of-cpas-reports

[3] AICPA, SOC 2 reporting guide, October 2022; public guide description. https://www.aicpa-cima.com/cpe-learning/publication/soc-2-reporting-on-an-examination-of-controls-at-a-service-organization-relevant-to-security-availability-processing-integrity-confidentiality-or-privacy

[4] AI Standards Inc., engineering architecture documentation, snapshot 25 August 2026; scoped access, deployment gates, governed inference and lineage. Internal engineering source.

[5] AICPA, 2018 SOC 2 Description Criteria, with revised implementation guidance 2022. https://www.aicpa-cima.com/resources/download/get-description-criteria-for-your-organizations-soc-2-r-report

[6] AI Standards Inc., filed provisional patent application records on AI decision provenance and continuity of evidence across multi-agent workflows. Company records.

[7] AICPA, SOC 2 reporting guide, 2018 excerpt, paragraphs 1.05 and 1.08. https://catalogimages.wiley.com/images/db/pdf/9781945498602.excerpt.pdf

About the author Anwesh Rath, Founder and CTO, AI Standards Inc.

Disclaimer Informational, not legal advice. Applicability depends on jurisdiction, sector, role and use. Obtain qualified advice and validate controls for the intended deployment. Examples are illustrative; no certification, compliance or performance guarantee is offered.