

INSIGHTS & ASSURANCE

KOMPLEYE
Assurance Your Way

EXECUTIVE WHITE PAPER / 2026

AI in SOC 1 and SOC 2 *Examinations*

A Leadership and Operational Readiness Whitepaper

GOVERNANCE · RISK · EVIDENCE · TRUST

EXECUTIVE SUMMARY

When AI becomes *SOC-relevant*

Artificial intelligence becomes relevant to a SOC 1 or SOC 2 examination when the organization relies on it to deliver a covered service, execute or support a control, make a consequential decision, or manage risks that affect report users. The practical issue is not the presence of AI alone, but the degree of reliance, the adequacy of human review, and the potential effect on financial reporting, service commitments, system requirements, and customer trust.

01 / 03



Reliance

AI used to deliver a covered service or support a control

02 / 03

Human review

Independent validation, exception handling and approval

03 / 03

Report users

Effects on financial reporting, commitments and trust

THE DECIDING FACTOR

Reliance determines *relevance.*



EXECUTIVE BRIEF

What leaders need to *know*



EXECUTIVE OVERSIGHT

Decisions grounded in *risk and trust*

01

Scope must follow reliance and risk

SOC 1 scoping should consider effects on internal control over financial reporting. SOC 2 scoping should consider the applicable trust services categories and whether the selected scope meets intended-user needs.

02

Existing frameworks remain usable

AI does not automatically require a separate control framework, but it may change risk assessments, control design, disclosures, testing, and evidence expectations.

03

Human review must be substantive

Review is more persuasive when personnel can independently validate or reperform outputs, identify exceptions, and document approval rather than merely accept AI recommendations.

04

Third-party AI remains management's responsibility

Providers and subservice organizations require due diligence, contractual safeguards, ongoing monitoring, complementary controls, resilience planning, and clear disclosure.

EXECUTIVE ACTIONS

Five immediate *priorities*

01

Establish accountability

Assign executive oversight and operational ownership for AI governance, risk acceptance, deployment, monitoring, incidents, and retirement.

02

Create enterprise visibility

Inventory approved, embedded, internally developed, third-party, and unsanctioned AI use across services, processes, and controls.

03

Define reliance standards

Specify permitted uses, required human review, independent testing thresholds, escalation criteria, and prohibited applications.

04

Align SOC scope and disclosures

Reassess SOC 1 relevance, SOC 2 categories, system descriptions, customer responsibilities, and subservice organization treatment.

05

Build evidence readiness

Operate and retain evidence for data governance, model validation, security, privacy, monitoring, change management, exceptions, and vendor oversight throughout the examination period.

SECTION I / EXECUTIVE PERSPECTIVE

Why this *matters*

AI changes the SOC conversation when it is integral to a service, relied upon within a business process or control, or introduces or mitigates risk to outcomes covered by the examination. The central question is not whether AI is present, but whether the organization depends on it and whether that dependence could affect financial reporting, service commitments, system requirements, or the information needs of report users.



HUMAN JUDGMENT

Independent review *is essential*

> *The central question is not whether AI is present, but whether the organization depends on it.*
>
> EXECUTIVE PERSPECTIVE

Human involvement is a key factor. AI-assisted work that is fully reviewed, corrected, tested, and independently approved may present less examination relevance than AI output that cannot be completely reperformed or is accepted with only limited review. Materiality remains qualitative as well as quantitative: pervasive effects across customers or critical processes may make AI-related controls and governance important to the report.

EXAMINATION SCOPE

SOC 1 *versus* SOC 2

DIMENSION	SOC 1	SOC 2
PRIMARY LENS	Relevance to user entities' internal control over financial reporting.	Relevance to selected trust services categories and related service commitments and system requirements.
SCOPING TRIGGER	AI affects transaction flows, accounting records, financial reporting processes, the relevant IT environment, or controls supporting those areas.	AI affects the service or its security, availability, processing integrity, confidentiality, or privacy.
CONTROL DESIGN	Existing control objectives will often remain appropriate, but AI may alter risks and the controls needed to achieve them.	Existing trust services criteria remain usable, while AI-specific points of focus, commitments, or system requirements may be needed.
SCOPE CONCERN	Assess reliance, human review, materiality, and impact on financial statement assertions.	Assess whether selected categories meet intended-user needs; a security-only scope may be insufficient where processing integrity is central.

GOVERNANCE PRIORITIES

Six executive *decisions*

01	**Establish accountable AI governance**	Assign board or executive oversight, name operational owners, define approval authority, and maintain appropriate competencies.
02	**Demand enterprise visibility**	Maintain an inventory of sanctioned and unsanctioned AI use, including embedded AI, internally developed models, third-party services, and AI used in controls.
03	**Set a risk-based reliance standard**	Define when AI output may be relied upon, the depth of human review required, and when independent testing or re-performance is mandatory.
04	**Align report scope with customer risk**	Reassess SOC 1 relevance and SOC 2 trust services categories as AI use changes; ensure the report is not misleading by omission.
05	**Invest in evidence-ready controls**	Fund data governance, model testing, monitoring, security, incident response, change management, and vendor oversight before the examination period.
06	**Make third-party AI visible**	Require vendors and subservice organizations to disclose AI use, accept contractual safeguards, and provide monitoring evidence.

SECTION II / OPERATIONS PLAYBOOK

II

Inventory AI *before controls*

Create a complete inventory of AI use cases across production services, software development, cybersecurity, finance, customer support, compliance, and control execution. For each use case, document the business purpose, AI type, model or provider, data sources, users, outputs, downstream systems, decision rights, human review, and customer impact.

CLASSIFY EACH USE CASE

01 Is AI the service, a key enabler of the service, or embedded within a process or control?

02 Does the organization rely on AI output, or can personnel fully verify and reperform the result using independent evidence?

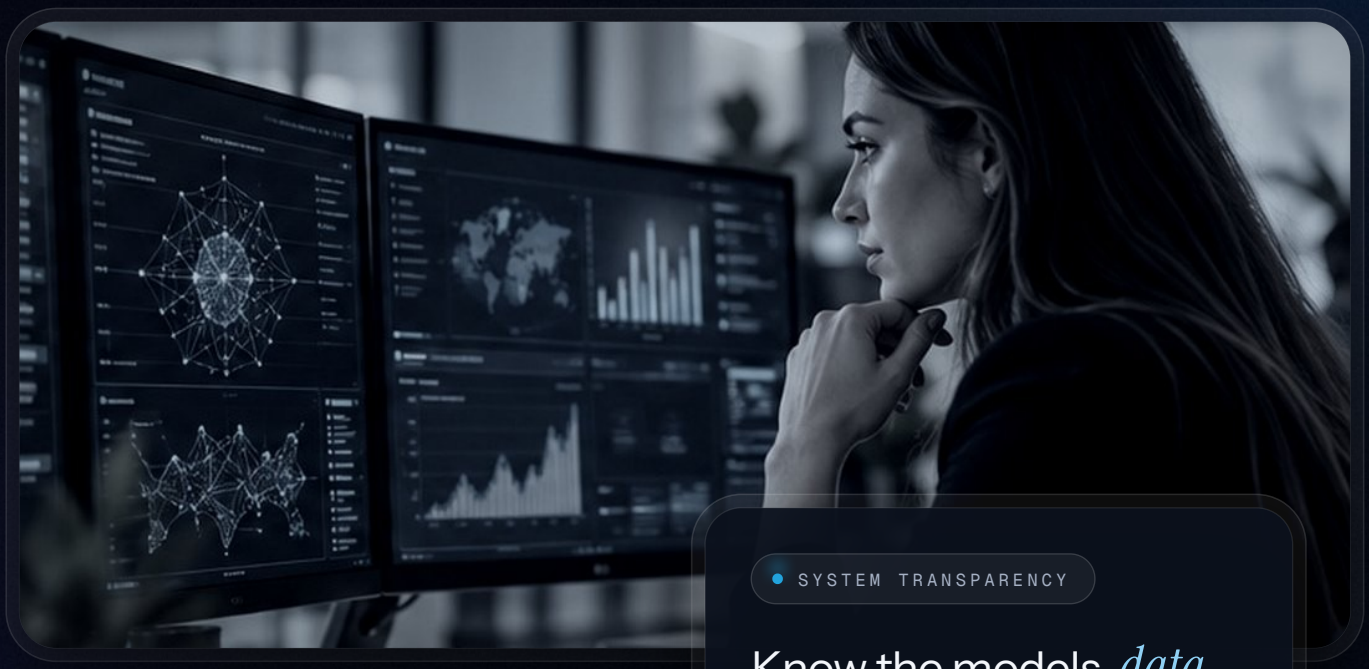
03 Could the AI affect user entities' financial reporting, a stated control objective, a service commitment, or a system requirement?

04 Could the effect be pervasive or significant to intended report users?

05 Is AI operated by a subservice organization, vendor, or foundation-model provider?

SYSTEM DESCRIPTION

Describe the system *accurately*



SYSTEM TRANSPARENCY

Know the models, *data and controls*

Document relevant AI dependencies

When AI is relevant, the system description should explain how AI is used to deliver the covered services; the types of AI employed; relevant models, algorithms, inputs, assumptions, and important data characteristics; AI governance and risk management; human-in-the-loop controls; and applicable complementary user entity controls and complementary subservice organization controls.

Descriptions should be specific enough for users to understand the risk without disclosing unnecessary proprietary detail. Ensure terminology, diagrams, process narratives, control matrices, policies, and vendor inventories tell a consistent story.

RISK-CONTROL MATRIX / PART 1

AI risks, controls *and evidence*

RISK	CONTROLS	AUDIT EVIDENCE
R01 **Data quality and drift**	Source approval, lineage, validation, change detection, drift thresholds, and remediation.	Data inventories, validation results, monitoring logs, exceptions, and approvals.
R02 **Model error and concept drift**	Predeployment testing, acceptance criteria, version control, periodic validation, rollback, and change management.	Test plans, results, model cards, release records, monitoring reports, and change tickets.
R03 **Hallucination and reliability**	Grounding, output validation, restricted use cases, human review, escalation, and independent corroboration.	Review records, exception samples, prompt and output logs, quality metrics, and escalation tickets.
R04 **Bias and overreliance**	Representative testing, defined fairness or accuracy criteria, diverse reviewers where appropriate, and training against automation bias.	Bias assessments, test populations, reviewer signoffs, training records, and remediation plans.

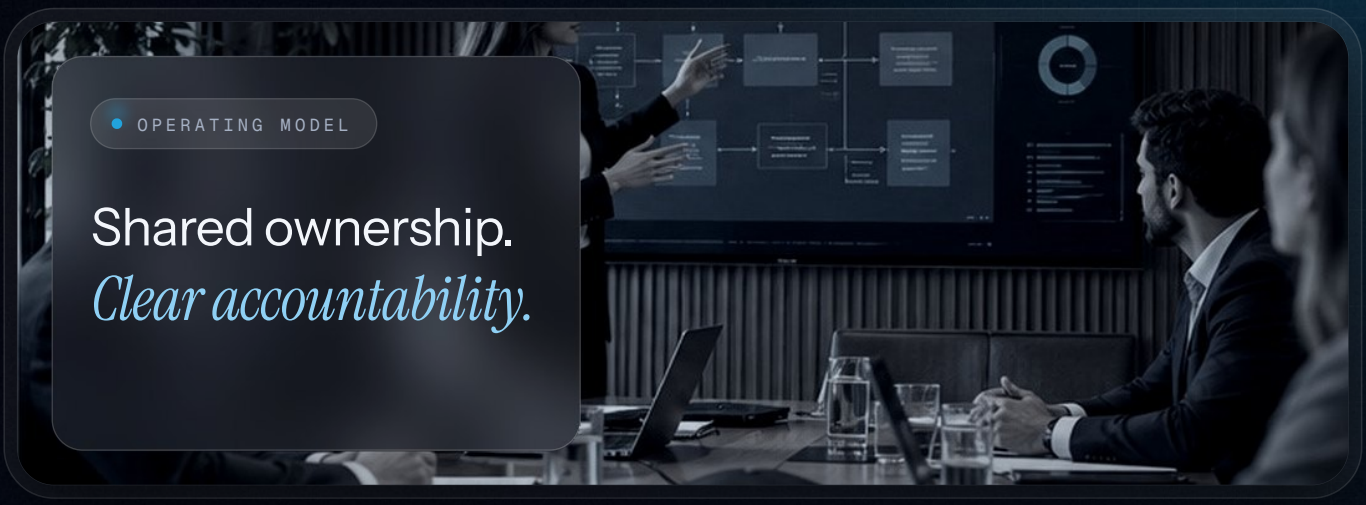
RISK-CONTROL MATRIX / PART 2

AI risks, controls *and evidence*

RISK	CONTROLS	AUDIT EVIDENCE
R05 **Explainability**	Document intended use, limitations, decision rationale, re-performance expectations, and user guidance.	Model documentation, decision records, user procedures, and validation reports.
R06 **Security**	Least privilege, data and model integrity, prompt-injection defenses, secure development, secrets protection, logging, and incident response.	Access reviews, security tests, monitoring alerts, incident records, and remediation evidence.
R07 **Confidentiality and privacy**	Approved data use, minimization, retention, disposal, consent or legal basis, provider restrictions, and leakage prevention.	Data assessments, contracts, configurations, deletion records, and privacy reviews.
R08 **Third-party dependency**	Due diligence, contractual AI terms, subservice monitoring, service changes, resilience, and exit planning.	SOC reports, questionnaires, contracts, monitoring reviews, issue logs, and contingency tests.

OPERATING MODEL

Governance and *third parties*



04 Strengthen governance

Define an AI governance structure proportionate to complexity and reliance. Establish policies for approved use, risk acceptance, model and provider onboarding, testing, deployment, monitoring, incident management, and retirement. Assign owners for each AI component and control, and require periodic reporting to oversight bodies. Compliance, operations, internal audit, security, engineering, privacy, legal, and finance personnel need enough AI fluency to identify risks and evaluate evidence. Relevant competencies include AI concepts, data quality, model behavior, bias, explainability, governance frameworks, security threats, monitoring techniques, controls, professional skepticism, and communication with technical and nontechnical stakeholders.

05 Manage subservice organizations

For inclusive subservice organizations, identify and test relevant AI-enabled processes and controls. Under the carve-out method, define the complementary subservice organization controls needed to address AI risk. Vendor monitoring should determine how the provider uses AI, whether contractual terms govern that use, whether management's risk assessment captures the dependency, and whether monitoring controls operate effectively. Where achieving control objectives or service commitments depends on customers, clearly identify complementary user entity controls. Examples may include customer review of AI-generated outputs, restrictions on submitted data, approval of high-impact decisions, timely notification of erroneous results, and customer access administration.

OPERATING CONTROLS

Evidence-ready *by design*

01

Retain evidence at the frequency stated in the control, not only before fieldwork.

02

Preserve model, prompt, configuration, data, and policy versions sufficient to explain tested activity.

03

Link exceptions to investigation, risk assessment, approval, remediation, and closure evidence.

04

Use population-complete records where AI performs a continuous or automated control.

05

Validate that dashboards and monitoring reports are complete, accurate, and traceable to source records.

06

Reassess scope after material model, provider, data, process, or service changes.

VALIDATION QUESTIONS

Examination readiness *checklist*

Q1

Do we know every relevant AI use case?

Approved inventory reconciled to technology, vendor, and process records.

Q2

Have we documented reliance and human review?

Clear rationale supported by procedures and operating evidence.

Q3

Is the SOC scope still appropriate?

Documented SOC 1 relevance and SOC 2 category assessment.

Q4

Does the system description match operations?

Consistent narratives, diagrams, control matrices, and third-party disclosures.

Q5

Are AI risks mapped to controls?

Risk-control matrix covering data, model, human, security, privacy, and vendor risks.

Q6

Can controls be tested?

Named owners, defined frequency, testable criteria, complete populations, and retained evidence.

Q7

Are exceptions governed?

Thresholds, escalation, impact analysis, remediation, and closure records.

Q8

Are third parties covered?

Due diligence, contractual requirements, monitoring, CSOCs, and contingency plans.

READINESS ROADMAP

A practical *90-day roadmap*



DAY 1 — DAY 30 — DAY 60 — DAY 90

PHASE 01

Discover and assess

DAYS 1-30

Form the cross-functional working group; inventory AI; identify owners, providers, data, reliance, and downstream effects; and perform preliminary SOC 1 and SOC 2 scoping.

PHASE 02

Design and document

DAYS 31-60

Update risk assessments, system descriptions, policies, control matrices, service commitments, system requirements, complementary controls, vendor terms, and evidence standards.

PHASE 03

Operate and validate

DAYS 61-90

Run controls, collect evidence, test monitoring completeness and accuracy, remediate exceptions, perform a readiness review, and confirm that examination scope and organizational competence remain appropriate.

CONCLUSION

The principle remains *clear*



AUDIT ASSURANCE

Evidence must *stand up to review.*

VISIBLE · GOVERNED · TRACEABLE

The AICPA guidance reinforces a practical principle: SOC examinations remain focused on the service organization's system, control objectives, service commitments, system requirements, and controls. AI does not automatically require a separate control framework, but it can materially change risk, reliance, disclosure, governance, testing, and evidence. The strongest response is to make AI use visible, evaluate it through the correct SOC lens, and operate controls that are proportionate, traceable, and demonstrably effective.

- Visible
- Governed
- Traceable

Source basis: AICPA Q&A Section 9561, "SOC Examinations: Effect of the Service Organization's Use of AI on SOC 1 and SOC 2 Examinations," issue date September 2026. This whitepaper is an operational interpretation and is not legal or accounting advice.



NEXT STEP

Ready to strengthen your *SOC readiness?*

Make AI use visible, manage reliance, and build the evidence needed for a defensible examination.

Talk to an expert

Schedule your FREE consultation

KOMPLEYE
Assurance Your Way

kompleye.com