sigon

WHITEPAPER · AI GOVERNANCE

AI GOVERNANCE IS BUILT, NOT WRITTEN

How Internal Audit and IT Compliance Help Leadership Lead AI

Tifany Espindola · Consulting Executive **Tyler Parenti** · COO November 2026

EXECUTIVE SUMMARY

AI HAS TO EARN ITS PLACE

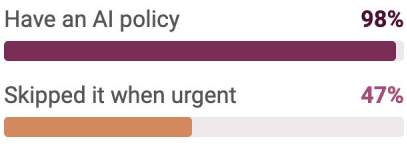
The board wants growth or savings from AI, without creating new compliance problems. At Sigon, governance means each use case is on track to bring back at least $2 for every $1 it costs, and problems are caught before they become an audit finding, an incident or a loss.

2:1
The Sigon bar: at least $2 back for every $1 an AI use case costs, counting people's time

7%
of 2,145 senior leaders placed their organization at established returns on AI (KPMG, 2026)[1]

25%
of AI initiatives had delivered the return CEOs expected (IBM CEO Study, 2025)[2]

Most companies are not there yet. The return leaks in ways the business case never shows: spending no one tracks, duplicate tools, pilots that never reach production, wrong answers people have to catch and fix, data leaks and regulatory penalties, and lost trust when something goes wrong in public.



The gap is not a missing policy. Most companies already have one; few have AI governance. In an EY survey published in September 2026, 98% of senior AI decision-makers at large US public companies reported a formal AI governance policy, yet 47% said their organization had previously skipped its governance process for urgent AI deployments.[3] **The policy existed. The control did not.**

AI THAT PAYS, WITH NO SURPRISES.

THE OUTCOME WE PROPOSE

- Each AI use case delivers the return its business case promised.
- Leadership knows about every high-risk AI use.
- No AI-related audit finding, incident or loss gets through that a control should have caught.

Section 6 shows how governance gets each use case to 2:1, and turns this into seven measures leadership can review each quarter.

© 2026 Sigon Group LLC. All rights reserved.

EXECUTIVE SUMMARY

CLOSING THE GAP COMES DOWN TO THREE IDEAS

01

Governance is built, not written.

It works when it starts from the outcome the business wants and the AI actually in use, not from principles written before anyone knows what that is.

02

AI is a new actor, not a new control set.

An AI agent goes through the access, change and review controls the company already runs. What is genuinely new is relying on its conclusions.

03

Human review is the default; removing it is a control change.

AI fails a control in three ways: accuracy, permissions and problems no one has seen yet. The reviewer comes out one control at a time, against all three. Where a decision rests on judgment, a person stays on it.

Management owns AI governance

IT compliance and internal audit play a critical role in making it work.

IT COMPLIANCE

Tracks, monitors and provides feedback

IT compliance (or the SOX program team) owns the inventory, intake, the controls and the testing required before go-live.

INTERNAL AUDIT

Advises first, assures later

Maps where AI is actually in use, pressure-tests control design while changes are still inexpensive, and later gives assurance that the controls work.

For mid-market companies. Many lack a dedicated IT compliance function. There, the SOX program lead or internal audit may cover both roles, which makes the independence safeguards in Section 6 even more important.

© 2026 Sigon Group LLC. All rights reserved.

SECTION 01

ADOPTION IS OUTRUNNING GOVERNANCE

AI is already inside the business, and most of it arrived without a governance decision. It came in through a vendor's product update, a team's new subscription or a process owner who built an agent over a weekend.

21%

of companies have a mature governance model for AI agents, though about three in four plan to deploy them (Deloitte, Jan 2026)[4]

5%

of about 400 internal auditors called their AI governance mature; about half had a policy with little enforcement (live poll, Oct 2026)[5]

6%

of executives said a typical AI use case paid back within a year; most said two to four years (Deloitte, 2025)[6]

The gap is not awareness. It is mechanics.

Policy-first governance does not close the gap. A policy written before anyone knows which use cases exist describes an organization that does not exist. Employees who need a tool this week route around an approval process that takes a quarter. Shadow AI is usually impatience, not defiance, and a faster intake process fixes more of it than a stricter policy does.

Governance is not project management. A project delivers one approved change and ends at go-live. Governance is the standing authority every project runs under: who may approve what, which risks the company accepts, which controls must exist and what leadership sees. The test: when a vendor switches on an AI feature in your ERP, there is no project, but governance still has to catch it.

© 2026 Sigon Group LLC. All rights reserved.

The rulebook is still being written

Jan 2023	**NIST AI RMF 1.0.** Remains the reference framework. NIST has said it is under revision and has published no date.[7]
Dec 2023	**ISO/IEC 42001:2023.** Offers a certifiable AI management system, but few mid-market companies are ready to certify.[8]
Jun 2024	**PCAOB technology-assisted analysis amendments.** Effective for fiscal years beginning on or after December 15, 2025, they expressly left AI aside. The PCAOB has issued no AI-specific auditing standard or requirement.[9]
Feb 2026	**COSO, *Achieving Effective Internal Control Over Generative AI*.** Applies the existing 17 principles of internal control to generative AI rather than creating a new framework.[10]
In progress	**PCAOB staff guidance on AI in financial reporting.** Until it arrives, external auditors evaluate controls that rely on AI under existing standards.[9] Meanwhile, COSO's evidence expectations for AI reliance are the best available benchmark.[10]

What this means for leadership: waiting for a settled standard means governing nothing in the meantime. Govern what is already in use, with frameworks that already exist, and adjust as the standards settle.

SECTION 02

THE LENS: TRANSFORMATION, TECHNOLOGY, RISK, CONTROLS

Every AI conversation should start with the outcome the business wants, not the model. The model is the least stable part of the picture; it will be replaced or updated several times in the life of a single use case. The outcome is what leadership approved, funded and is accountable for. Then ask, in order: what result do we want, what has to change in the process to get there, what technology supports that process, and only then where AI belongs.

© 2026 Sigon Group LLC. All rights reserved.

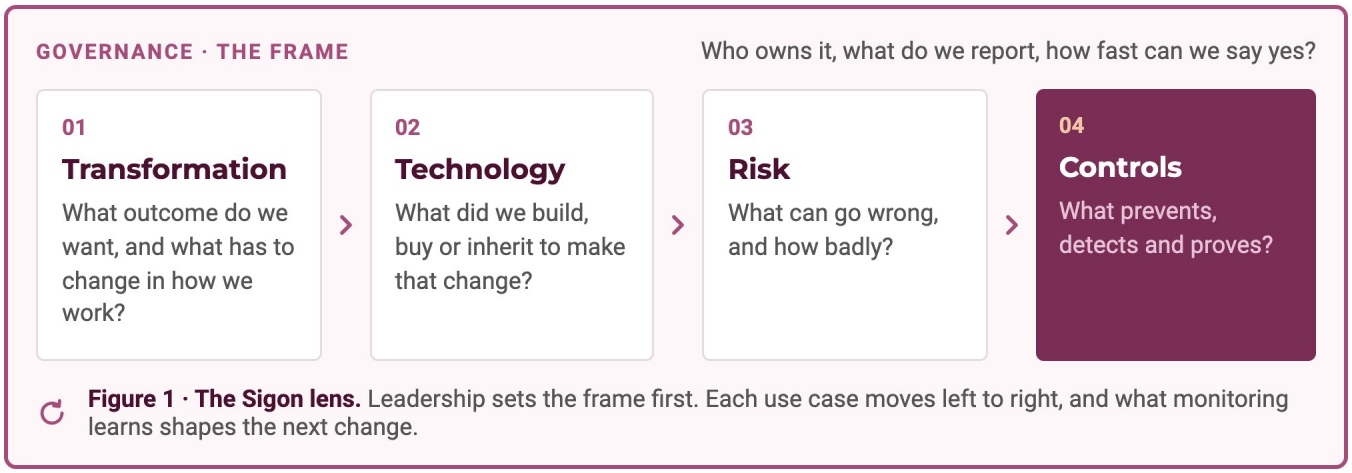


Figure 1 · The Sigon lens. Leadership sets the frame first. Each use case moves left to right, and what monitoring learns shapes the next change.

Governance is the frame, not the last step. Leadership sets ownership, risk appetite and the intake process before the first use case moves. What gets built bottom-up is the *content* of governance: the inventory, the risk tiers, the controls and the reporting. NIST treats its GOVERN function the same way, as cross-cutting rather than sequential.[7]

The loop matters more than the line. Monitoring one use case tells leadership what the next one needs. An organization that runs the loop three times has a governance program built from evidence. One that writes a policy first has a document.

Who does what at each step

STEP	QUESTION LEADERSHIP ASKS	IT COMPLIANCE · TRACKS, MONITORS	INTERNAL AUDIT · ADVISES, ASSURES
Transformation	What outcome do we want, and what has to change in how we work to get there?	Tracks the AI inventory and keeps it current	Maps where AI is actually in use, including what no one declared
Technology	What did we build, buy or inherit to make that change?	Brings each agent into existing access and change controls	Baselines the access, change and data controls it sits on
Risk	What can go wrong, and how badly?	Tiers each use case so control effort matches the risk	Challenges the ranking
Controls	What prevents, detects and proves?	Designs the controls, tests before go-live, reviews requests to remove the human	Independently tests that the controls work once they operate
Governance (the frame)	Who owns it, what do we report, and how fast can we say yes?	Runs intake and the quarterly readout; a voting committee seat	Design advice first, assurance later; a non-voting committee seat

© 2026 Sigon Group LLC. All rights reserved.

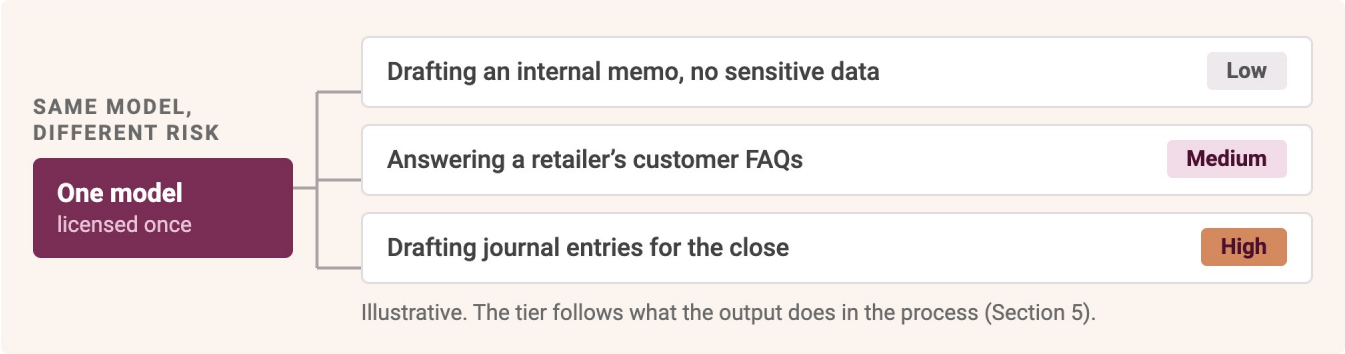
SECTION 03

TRANSFORMATION: START FROM THE OUTCOME

Fix the process before you add AI. Many steps exist only to check or move work by hand. Ask what the process must deliver, remove what does not serve that, make it consistent, and then place AI.

The first deliverable is an honest inventory, and it is almost always longer than leadership expects. Ask business leaders what AI they use and you get the sanctioned list. The real list sits in vendor contracts, expense reports and process owners' toolkits.

Inventory use cases, not tools. A use case is one job AI does inside one business process, such as matching invoices to purchase orders. Track it there, because that is where the risk lands: the same model can be harmless in one process and high-risk in the next. Each use case gets an inventory record, and so does each agent inside it (Section 4).



Illustrative. The tier follows what the output does in the process (Section 5).

Finding the AI no one declared

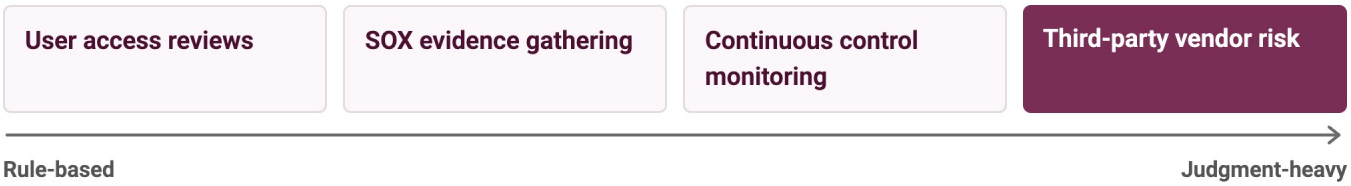
IT compliance builds the inventory. Internal audit maps what no one declared; this is advisory work, since there is nothing to test yet.

1. **Vendor inventory.** Compare it to the approved AI tool list. Vendors that added AI in a product update rarely trigger a new review.
2. **Spend data.** Mine purchasing card, accounts payable and expense data for AI subscriptions.
3. **SOX sub-certifications.** Add one AI question to the quarterly 302 sub-certifications process owners already sign.
4. **Ask about automation, not AI.** People describe agents and scripts as automation long before they call them AI.
5. **Usage logs.** A team that clearly works with AI but barely touches the approved tool is using a different one.

> **The first few real use cases teach more than any framework document.**

© 2026 Sigon Group LLC. All rights reserved.

Start with what is being built now, especially anything that touches financial reporting. In our own work, the use cases we have built first are user access reviews, gathering evidence for SOX control testing, continuous control monitoring and third-party vendor risk assessments. They run from rule-based work, where every answer can be checked, to judgment-heavy work like vendor risk, where a person stays on the call.



The output for leadership is a one-page view

Each use case, the process it touches, who owns it, which cost center pays for it, the business case it was approved on, and whether it affects financial reporting, customers or regulated data. Every AI tool and API key comes through this same intake, so every dollar of AI spend has an owner and a purpose.

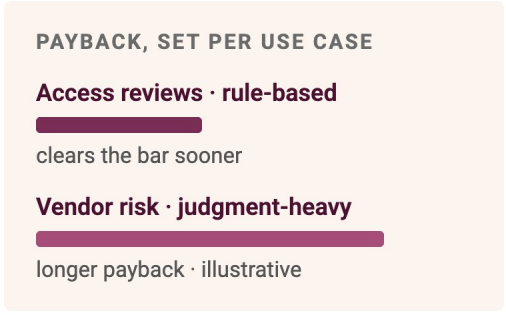
AI use case register · one record ILLUSTRATIVE

USE CASE	PROCESS	BUSINESS OWNER	COST CENTER
Matching invoices to purchase orders	Accounts payable	[Name, title]	[Cost center]

BUSINESS CASE	AGENTS	AFFECTS
$2 back per $1 by [date]	Each with its own record	Financial reporting · Customers · Regulated data

The business case is where the 2:1 bar lives

Each use case shows how it will bring back at least $2 for every $1 it costs, counting people's time, and by when. A rule-based use case like access reviews should clear the bar sooner than a judgment-heavy one like vendor risk, so the payback period is set for each use case, not once for all of them.



FRAMEWORK TIE · NIST MAP 1 · COSO Principles 6 (suitable objectives) and 9 (identifying significant change) · ISO 42001 clause 4.1, Annex A.9.4

© 2026 Sigon Group LLC. All rights reserved.

SECTION 04

TECHNOLOGY: A NEW ACTOR, NOT A NEW CONTROL SET

Built, bought and embedded AI carry different risks, and control sits in different places.

Built

An internal agent that prepares account reconciliations

WHERE CONTROL SITS

Your development, testing and change process

Bought

A licensed AI tool for contract review

WHERE CONTROL SITS

Vendor due diligence, contract terms, your configuration

Embedded

An AI feature switched on in your ERP or ticketing system

WHERE CONTROL SITS

Often nowhere, until someone notices

Citizen-built

A process owner's agent built in a low-code tool

WHERE CONTROL SITS

The business, which usually has no change-control discipline

Citizen-built agents move IT general controls into the business. When a finance analyst builds an agent that pulls data and drafts journal entries, that analyst is now operating access, change and data-integrity controls, usually without knowing it.

SIDEBAR · VENDOR AI

The most common contract gap is **training rights**: whether the vendor may use your data to train its models. Also worth adding: notice before a model change, audit and testing rights, incident notification, and disclosure of the underlying model providers.

AI does not need a parallel set of controls

The instinct when AI arrives is to write a new list of AI controls next to the existing ones. That doubles the control matrix without reducing much risk. An AI agent is a new actor inside controls the company already runs.

Access

An agent is an identity. Provision it, review it and remove it like a person or a service account, with least-privilege access. Sharing an agent also shares whatever data it can reach.

Change

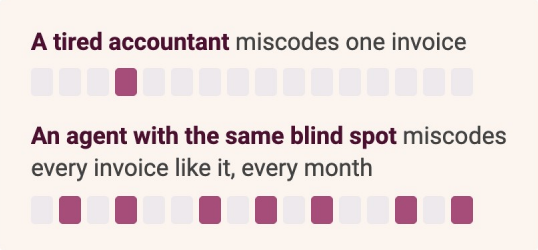
Prompts and model versions are configuration. A prompt edit can change a control's behavior as much as a code change, so it goes through the same change process: request, test, approve, deploy, record.

Scope

Grant actions, not systems. List what the agent may do, such as read the vendor master or draft a reconciliation, and block everything else. Job-role access is too coarse for an agent; NIST's attribute-based model, which decides each request on who is asking, for what and under what conditions, fits better.[11]

© 2026 Sigon Group LLC. All rights reserved.

What AI does change is speed and reach. Weak access, poor data quality, missing segregation of duties and loose change control were risks before; AI makes each one faster and harder to see. Scale turns small errors into big ones. And if people only see the items the agent chooses to raise, nobody sees what it misses.



Don't rely on the agent to police itself.

Build the hard stops into the systems around it: access rights, approval steps, posting rules. They hold whatever the agent decides. That matters most when the agent reads outside content: an invoice, email or vendor document can carry text written to steer it. The hard stops hold even when it follows that text.

Every agent is a joiner, a mover and a leaver

Run agents through the joiner, mover and leaver process you already use for people. Each agent has a named owner, and that owner answers to the business owner of the use case. Its inventory record answers five questions.

Agent record — Joiner › Mover › Leaver

1. What may it decide, and what must it hand to a person?
2. What data and systems can it reach?
3. Who checks its work, and how often?
4. When was it last tested?
5. What is it known to get wrong, and what should its reviewer watch for?

When the agent is no longer needed, or its owner changes jobs, its access comes off that day and its record closes. **A forgotten agent is an inventory gap and an open access risk.**

IT compliance brings each agent into the access and change controls. Internal audit baselines those controls first: **AI amplifies weak access, data or change controls, so confirm the foundation before assessing the AI.**

FRAMEWORK TIE · NIST MAP 2–4 · COSO Principle 11 (general controls over technology) · ISO 42001 Annex A.4, A.6, A.10

© 2026 Sigon Group LLC. All rights reserved.

SECTION 05

RISK AND CONTROLS: THREE WAYS AI FAILS A CONTROL

Where AI does control work, it needs the same evidence a person would. An agent that prepares a reconciliation, tests a sample or drafts a journal entry is performing a control. When management relies on that output as evidence the control was designed or operated effectively (COSO calls this AI reliance[10]), the external auditor will ask who approved it, what data it used and whether it would produce the same answer again. If no one can say, the control fails, however good the model is.

Rank use cases first, so effort matches risk

Tier by what the AI's output does, not by who sees it. Three tiers are enough to start; the right column is what each tier needs.

High	The output decides or materially shapes a consequential decision about a person; is used in financial reporting or a SOX control; processes sensitive regulated data; gives customers information about regulated products (rates, terms, coverage, eligibility); or acts toward customers without review of each action.	**A person in the loop, documented controls, testing before go-live, periodic internal audit review**
Medium	Customer-facing information with disclosure and accuracy controls, or internal decisions with meaningful human review.	**A named owner, an inventory entry and logging**
Low	Drafting and summarizing with no sensitive data.	**Approved tools and acceptable-use training**

Human review lowers a tier only when no High rule applies. A retail FAQ chatbot is Medium; a bank chatbot quoting loan terms is High.

Code does the math, AI does the words.

Our rule: if a number can be worked out by fixed rules, compute it in ordinary code, which your IT controls already cover. Let AI write the narrative around it, and review that like any draft.

© 2026 Sigon Group LLC. All rights reserved.

For high-risk uses, AI fails a control in three ways

Each has a control, an owner and a rhythm.

1 Accuracy

Wrong conclusions or data; different answers to the same input; confident output that is wrong

THE CONTROL

Define every major scenario the AI must handle and test each at volume (hundreds of cases, not the traditional sample of 20). Run the agent alongside the current process and compare, and keep the results. Break results out by the groups of people or transactions affected; a good average can hide uneven errors. Include hostile inputs, such as instructions planted in a document the agent reads.

OWNER

Use-case owner runs the tests; someone who did not build the agent signs them off

FREQUENCY

Before go-live; after any change to the model, prompt, data source or process; when incoming data shifts or a vendor updates its model; scheduled re-runs; a person re-performs a sample of live decisions

2 Permissions

The AI does more than it was designed to do

THE CONTROL

Manage agent access like a person's (Section 4): least privilege, scoped by task, changed only through change management.

OWNER

Agent owner; IT security runs the review

FREQUENCY

Every access review cycle, and whenever the agent's role changes

3 Unknowns

Failures no one has seen yet

THE CONTROL

Log 100% of activity: inputs, tool calls, model and prompt version, output and the reasoning given. Keep source evidence unchanged, in logs the agent cannot alter. That also answers the fraud question: people are deterred from hiding what they did by codes of conduct and consequences, and an agent needs the equivalent built in. If a second AI checks the first, that checker is a control too, with its own owner, test and log.

OWNER

Agent owner reviews; alerts go to a named person with authority to suspend the agent

FREQUENCY

Logging always on; log review monthly, feeding the next change

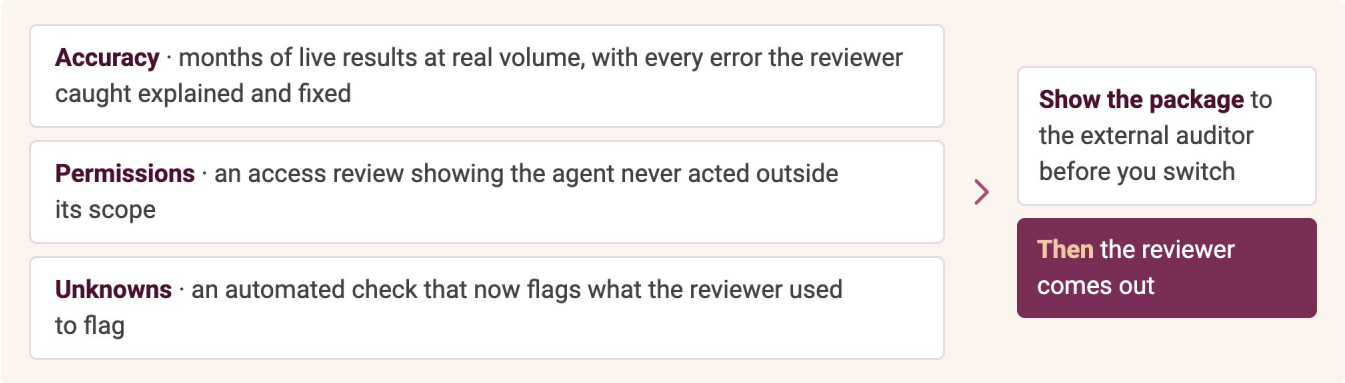
© 2026 Sigon Group LLC. All rights reserved.

Keep a human in the loop by default

For high-risk uses, a named person approves the output before it is used. That reviewer is accountable, and the review has to be real. It counts only if it is detailed enough to catch the kind of error the AI makes; a glance at a plausible summary does not qualify. People tend to check AI output less carefully than human work, so audit the decision, not the model. Choose a reviewer who knows the process well enough to spot an answer that looks right but isn't.

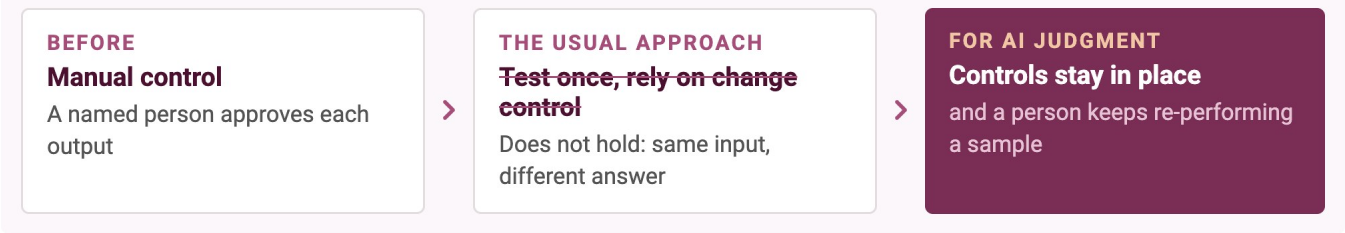
Take the human out one control at a time

A control owner who wants to remove human review asks for one specific use case, and IT compliance reviews the request. The reviewer comes out only when the evidence covers all three failure modes. **When every case is different, or the call is a judgment call, keep the person.**



When a reviewer comes out, the control changes class

What was a manual control becomes an automated one. The usual approach to automated controls (test once, then rely on change control) does not hold for AI judgment, because the same input can produce a different answer. Controls stay in place, and a person keeps re-performing a sample. After a handful of these decisions, the organization knows what its AI control framework actually needs; writing it before then is guessing.



© 2026 Sigon Group LLC. All rights reserved.

Ask control owners for four things, not a matrix

1 The name of the owner and of the person who can stop the agent

2 The test results for the situations it will meet

3 Its access list

4 A sample of its log

IT compliance builds the AI risk and control matrix during design; internal audit advises on it, then tests it. Builders don't sign off their own agents: test results are approved by someone outside the build. Changing a control during design costs an afternoon. Testing comes once the controls have operated long enough to produce evidence.

> **IN PRACTICE AT SIGON**
> One written security policy governs both our people and our AI agents. A security-review agent checks AI-assisted work against that policy whenever client data, credentials, new tools or access changes are involved, and the work stops until a person acknowledges any high-risk finding.

Where the Responsible AI principles land

Many organizations start from a list of principles. Each one lands on a control in this paper: the principles describe what good looks like, and the three failure modes are how you prove it.

Fairness and robustness	Accuracy tests, broken out by group, run against hostile inputs and repeated when the data shifts
Privacy and safety	Permissions: what the agent can reach and what it may do
Transparency	The activity log and the agent's record, including what it is known to get wrong
Accountability	The named owner and the person who can stop the agent
Human oversight	The reviewer, removed only one control at a time

FRAMEWORK TIE · NIST MAP 5, MEASURE 1–2, MANAGE 1–3 · NIST AI 600-1 · COSO Principles 7–8 and 10–12 · ISO 42001 clauses 6.1.2, 6.1.4, 8.1–8.4, Annex A.5–A.7, A.9

© 2026 Sigon Group LLC. All rights reserved.

SECTION 06

YOUR AI GOVERNANCE OUTCOME

AI governance should steer toward a result you can check: **AI that pays, with no surprises.** SOX governance steers toward a clean audit opinion; security governance toward no data loss.

THE DEFINITION

Every AI use is known, owned, paying back what it promised and controlled in proportion to its risk, and leadership can prove it.

Built from the steps above, governance comes down to four things

01 Named ownership

One executive owns AI risk overall, and each use case has a business owner. “It depends” is the answer that signals the gap.

02 A policy that matches practice

Check it against what the inventory found, not against a template. Common gaps: no prohibited uses, no risk tiers, nothing on vendor AI, no exception process.

03 Intake fast enough to beat shadow AI

Low-risk requests should be approved in days. If the official path is slower than a credit card, people will use the credit card. A committee that meets once a quarter guarantees an out-of-date inventory; pre-approve common use cases so it spends its time on new risk.

04 A regular readout

Each quarter, leadership and the audit committee see the seven measures, what changed, and whether each use case is meeting the 2:1 bar on cost, accuracy, risk and speed. Put people’s time in the cost: a tool that saves a license seat but adds two days of checking at every close is not saving money.

© 2026 Sigon Group LLC. All rights reserved.

SECTION 06

HOW GOVERNANCE GETS YOU TO 2:1

Governance is often treated as the cost of using AI. Designed well, it is how the 2:1 bar gets met: it raises what each use case returns, lowers what it costs and protects what it earns.

↑ RAISE THE RETURN

Approve on a business case. Spend follows the value each use case promises.

Start where payback is fastest. Rule-based work clears the bar first.

Fix the process first. Otherwise AI automates the waste.

Reuse what works. One team's success becomes the next team's template.

↓ LOWER THE COST

One intake. No duplicate tools or untracked subscriptions.

Effort matches risk. Tiering keeps low-risk uses light.

Code does the math. Rule-based numbers need no AI review.

Design controls early. An afternoon in design, not a remediation project later.

PROTECT THE RETURN

Catch errors before they scale. An agent repeats its blind spot on every transaction like it.

Avoid the losses that erase gains. Audit findings, incidents, rework and lost trust.

Price the exposure. The SOX sidebar that follows shows how.

The 2:1 gate

Scale it | Fix it | Stop it

At each use case's payback date, leadership makes one of three decisions: scale it, fix it or stop it. That turns governance from a cost into portfolio management for AI.

THE EVIDENCE

3× Companies with clear accountability for AI outcomes are more than three times as likely to report established returns (14% vs 4%).[1]

5× Companies with full visibility into AI operating costs are five times as likely to report established returns (15% vs 3%).[1]

© 2026 Sigon Group LLC. All rights reserved.

Seven measures make it reviewable each quarter

MEASURE	WHAT IT SHOWS	TARGET (ILLUSTRATIVE)
High-risk AI uses found outside the inventory	Whether the inventory reflects reality	**Zero at each sweep**
High-risk uses with designed and tested controls	Whether risk is actually controlled	**100%**
Days from an AI request to a decision (low risk)	Whether the official path beats the workaround	**Under 5**
AI spend against budget, by use case	Whether every dollar has an owner and a purpose	**Within budget**
Use cases at or above 2:1 against their business case	Whether AI spend is paying back	**100% past payback date, or a fix-or-stop decision**
AI incidents and time to close	Whether problems are caught and fixed	**Tracked; trending down**
Hours from request to approved controls, per new use case	Whether reusable tiers and controls are cutting the work	**Lower each quarter**

SIDEBAR

Start where you already own the outcome

A SOX-focused internal audit function already has a clear outcome: no external audit findings. Every AI use that enters a financial process is a new way to miss it, which makes SOX the natural first part of AI governance to fund, because the exposure can be priced.

AI uses in SOX scope from the inventory × **Likelihood of a finding** what the controls cut × **Cost of a finding** remediation, retesting, audit work = **Exposure** the business case

~43%
audit-fee premium in the year a material weakness is disclosed[12]

21–32%
premium still paid three to four years after remediation[12]

$1–2M
at stake for 20 AI uses in SOX scope, if each produced a significant deficiency

For a mid-market company we estimate a significant deficiency at $50,000–$100,000 (remediation, retesting and added audit work) and a material weakness at $250,000 or more. The controls in Section 5 cut the likelihood, and the drop in exposure is the business case. **The same 2:1 bar applies: the governance work should cost no more than half the exposure it removes.** Once the model works for SOX, extend it to customer-facing, privacy and HR uses.

© 2026 Sigon Group LLC. All rights reserved.

IT compliance runs it; internal audit advises first and assures later

IT COMPLIANCE · VOTING SEAT

Runs intake and the quarterly readout and holds a voting seat on the AI governance committee.

INTERNAL AUDIT · NON-VOTING SEAT

Early on, advises on ownership, tiers and intake: it asks questions and offers insight, while management designs, decides and owns the result.[13]

This holds with four safeguards

1. Internal audit's advisory role is written into the audit charter.
2. The audit committee is told about the advisory work and guides how far it goes.
3. Before any later assurance on the same area, the chief audit executive confirms the advisory work did not impair objectivity.
4. That assurance is staffed by people who did not do the advisory work, or comes at least 12 months later.

Then internal audit shifts to assurance. Once controls have operated, it reports to the audit committee, including a value-realization review: did each use case deliver what its business case promised, and would the audit committee fund it again?

FRAMEWORK TIE · NIST GOVERN 1–6, MEASURE 3–4, MANAGE 4 · COSO Principles 1–5 and 13–17 · IIA Three Lines Model; Global Internal Audit Standards 2.2, 7.1 · ISO 42001 clauses 5–7, 9–10, Annex A.2–A.3

© 2026 Sigon Group LLC. All rights reserved.

SECTION 07

THE FIRST 30 DAYS

	ACTION	OWNER	OUTPUT
WEEK 1	**Name the executive owner of AI risk; agree the three risk tiers; make human review the default for high-risk uses**	Leadership	Owner named; tiers and default set
WEEK 1–2	**Inventory AI in use, including vendor-embedded and citizen-built, with a record for each agent**	IT compliance with business owners; internal audit finds what no one declared	One-page AI landscape
WEEK 2–3	**Tier the inventory; pick one high-value, high-risk use case**	IT compliance, internal audit advising	Ranked inventory; first use case chosen
WEEK 3–4	**Test that use case against accuracy, permissions and logging; build its risk and control matrix; walk the external auditor through the controls you designed and the evidence you plan to show them, before go-live**	Use-case owner and IT compliance, internal audit advising	Controls designed, tested and owned; auditor aligned
WEEK 4	**Stand up a lightweight intake and schedule the first quarterly readout**	Leadership	Intake live; readout on the calendar

That is a governance program built from evidence.

At the end of 30 days the organization knows what AI it has, which uses matter most, how one of them is controlled and what it has to return to earn its place. Each later use case runs through the same loop faster.

HOW SIGON CAN HELP

We can run these 30 days with you

Our AI Governance Framework Builder makes this a fast reality check. We help you build your AI inventory, including what no one declared, tier it and pick the first use case, then build its risk and control matrix, mapped to the main frameworks, with the test steps and evidence your external auditor will expect. **You come away with a prioritized roadmap.**

sigongroup.com

© 2026 Sigon Group LLC. All rights reserved.

APPENDIX

CROSSWALK

How each step maps to the frameworks your auditors, regulators and board already reference. The mapping is our synthesis.

SIGON STEP	NIST AI RMF 1.0	COSO GENAI (2013 PRINCIPLES)	IIA	ISO/IEC 42001
Governance (frame)	GOVERN 1–6	Control environment, Principles 1–5	Three Lines Model: management owns; internal audit provides independent assurance and advice	Clauses 5–7; A.2, A.3
Transformation	MAP 1	Principle 6 (suitable objectives); Principle 9 (significant change)	Advisory services (e.g., mapping AI use and risks)	Clause 4.1; A.9.4
Technology	MAP 2–4	Principle 11: general controls over technology	IT general controls (IIA GTAG)	A.4, A.6, A.10
Risk	MAP 5; MEASURE 1–2	Principles 7–8: identifies risk, assesses fraud risk	Risk-based audit planning (Standard 9.4)	Clauses 6.1.2, 6.1.4, 8.2, 8.4; A.5
Controls	MANAGE 1–3	Principles 10–12: control activities	Assurance: control testing	Clauses 8.1, 8.3; A.6, A.7, A.9
Loop (monitoring)	MEASURE 3–4; MANAGE 4	Principles 13–17: information, communication, monitoring	Continuous auditing (GTAG 3); reporting to the audit committee	Clauses 9.1–9.3, 10

SOURCES

1 KPMG International, *Global AI Pulse Q2 2026* (2,145 senior leaders in 20 countries; fielded April 28–May 25, 2026).

2 IBM Institute for Business Value, *2025 CEO Study*, press release, May 6, 2025 (2,000 CEOs in 33 countries, surveyed February–April 2025).

3 EY, "EY survey finds that autonomous AI implementation outpaces oversight, yielding an AI governance gap," press release, September 15, 2026 (202 senior AI decision-makers at US public companies with $1B+ revenue; fielded May 28–June 15, 2026).

4 Deloitte, *State of AI in the Enterprise 2026*, press release, January 21, 2026 (3,235 leaders in 24 countries, surveyed August–September 2025).

5 Internal Audit Collective, AI governance webcast, live audience poll (about 400 respondents), October 8, 2026. Unpublished; self-selected audience; a poll, not a survey.

6 Deloitte, "AI ROI: The paradox of rising investment and elusive returns," October 22, 2025 (1,854 executives in 14 countries in Europe and the Middle East, surveyed August–September 2025).

7 NIST, *AI Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, January 2023; revision status per nist.gov. NIST, AI 600-1, *Generative Artificial Intelligence Profile*, July 2024.

8 ISO/IEC 42001:2023, *Artificial intelligence — Management system*.

9 PCAOB, Release No. 2024-007, Amendments Related to Aspects of Designing and Performing Audit Procedures that Involve Technology-Assisted Analysis of Information in Electronic Form (June 2024); Staff Spotlight on generative AI (July 2024); Standard-Setting, Research, Rulemaking, and Related Activities, pcaobus.org (accessed October 2026).

10 COSO, *Achieving Effective Internal Control Over Generative AI*, February 2026.

11 NIST, SP 800-162, *Guide to Attribute Based Access Control (ABAC) Definition and Considerations*.

12 K. Raghunandan and D. Rama, "SOX Section 404 Material Weakness Disclosures and Audit Fees," *Auditing: A Journal of Practice & Theory* 25(1), 2006; K. Munsif, K. Raghunandan, D. Rama and M. Singhvi, "Audit Fees after Remediation of Internal Control Weaknesses," *Accounting Horizons* 25(1), 2011.

13 The Institute of Internal Auditors: *Global Internal Audit Standards* (2024), Standards 2.2, 7.1 and 9.4; Statement of Position on the Three Lines Model (July 2026); Internal Audit's Role in Governing Body/Executive Committees (2019); AI Auditing Framework (updated September 2024); GTAG on IT general controls; GTAG 3, Continuous Auditing.

© 2026 Sigon Group LLC. All rights reserved.

sigon

EXPERT LED. OUTCOME DRIVEN.

Expert-led, AI-native technology risk advisory across AI, Controls and Audit.

ABOUT THE AUTHORS

Tifany Espindola

CONSULTING EXECUTIVE

Tifany helps global organizations turn responsible AI principles into practical frameworks, controls and oversight. She brings more than 18 years in technology risk and controls across Brazil, the United States and Europe, most recently as a Managing Director at BDO USA and before that at PwC.

tifany@sigongroup.com

Tyler Parenti

COO

Tyler helps organizations modernize IT governance, controls and internal audit with practical automation and AI. A former PwC Senior Manager with more than 10 years in technology risk, SOX compliance, IT controls and SOC reporting, he has since built technology that modernizes controls testing.

tyler@sigongroup.com

sigongroup.com
Miami, Florida

© 2026 Sigon Group LLC. This paper is for general information and is not audit, legal or regulatory advice.