SENTINELSQUARED

AI GOVERNANCE SELF-ASSESSMENT

7 Questions Every Executive Should Be Asking About AI Governance

A practical self-assessment for organizations deploying AI in environments where a wrong answer isn't an inconvenience — it's a liability event.

ISSUED	DOCUMENT TYPE	AUDIENCE
April 2026	Self-Assessment	C-Suite / Board / Risk

> Most organizations discover their AI exposure **after the fact** — after a data incident, after a compliance inquiry, after a client discovers their work product was generated without oversight. These seven questions surface that exposure **before it becomes a liability event**. Work through each honestly. Any hesitation is a signal worth acting on.

QUESTIONS 01 – 03 · VISIBILITY & EXPOSURE

01 Do you know which AI tools your employees use daily?

> If you can't answer this, your AI footprint is ungoverned. Shadow AI use is the most common source of data leakage and compliance exposure — and the last thing an organization discovers in a breach investigation.

My answer: ☐ Yes — documented ☐ No / Unsure ☐ Partial visibility

Notes: _____

02 Are your contracts and proposals being drafted with AI?

> AI-assisted drafting without oversight creates professional liability risk. Errors, hallucinations, and inadvertent disclosures in client-facing documents are direct exposure. Most organizations have no process for this — and no awareness it's happening.

My answer: ☐ Yes — with oversight ☐ No / Unsure ☐ Yes — without process

Notes: _____

03 Is customer data being entered into public LLM interfaces?

> Customer PII, confidential business data, and privileged information entered into public AI systems may be used for model training, subject to foreign jurisdiction, and constitute a data breach under GDPR, CCPA, or HIPAA depending on your industry. If you don't know the answer, assume yes.

My answer: ☐ No — policy enforced ☐ Yes / Unsure ☐ No policy exists

Notes: _____

QUESTIONS 04 – 05 · OVERSIGHT & AUDIT

04 Who reviews AI-generated decisions before they go to clients?

> Unreviewed AI output carries the full liability of the organization that delivered it. If the answer is "nobody" or "whoever wrote the prompt," you have no oversight layer. That gap is what regulators, plaintiffs, and insurance carriers will ask about first.

My answer: ☐ Defined role / process ☐ No defined process ☐ Varies by team

Notes: _____

05 Can you produce an audit trail of AI-assisted work product?

> When litigation, regulatory inquiry, or a client dispute occurs, you will be asked to demonstrate what was AI-generated, what was reviewed, and what controls were in place. If the answer is "we'd have to reconstruct it from memory," that is not a defensible posture. Audit trails are the minimum bar for enterprise AI use.

My answer: ☐ Yes — system-generated ☐ No capability ☐ Manual only

Notes: _____

QUESTIONS 06 – 07 · POLICY & REGULATORY READINESS

06 **Do you have a written AI acceptable-use policy?**

> Without a written policy, you cannot hold employees accountable, demonstrate intent to regulators, or invoke it in litigation. A policy doesn't need to be 80 pages — it needs to exist, be signed, and be enforced. Most organizations don't have one. That gap is increasingly indefensible.

My answer: ☐ Yes — signed & current ☐ No ☐ Draft / In progress

Notes: _____

07 **If a regulator asked tomorrow, could you demonstrate governance?**

> This is the governing question. NIST AI RMF, ISO/IEC 42001, EU AI Act, and SOC 2 are all converging on a single requirement: documented, demonstrable governance. "We're working on it" is not a defensible answer when the inquiry arrives. The question is not whether you'll be asked — it's whether you'll be ready when you are.

My answer: ☐ Yes — fully documented ☐ No ☐ Partially

Notes: _____

INTERPRETING YOUR RESULTS

What Your Answers Mean

`0–1 gaps` **Governance-ready posture.** Your organization has the foundational controls in place. The next step is hardening: incident response planning, board-level reporting cadence, and framework alignment (NIST AI RMF / ISO 42001).

`2–4 gaps` **Significant exposure.** You have meaningful AI activity without sufficient governance infrastructure. Each gap is a documented risk. A structured framework engagement — typically 4–8 weeks — closes these before they surface as incidents.

`5–7 gaps` **High-priority remediation needed.** Your AI footprint is material, ungoverned, and likely expanding. This is not a "future initiative" — it is an active liability. The risk briefing below is the right starting point.

Why These Questions Matter Now

The EU AI Act has entered its enforcement phase. NIST AI RMF adoption is accelerating across regulated industries. State-level AI legislation is expanding. **Cyber insurers are beginning to underwrite AI risk separately.** The window to establish a defensible governance posture — before a regulatory inquiry, a client dispute, or a board question — is measured in months, not years. Organizations that treat governance as a design constraint rather than a checkbox are the ones that scale AI without creating existential exposure.

Book Your AI Governance Risk Briefing

No pitch decks. A direct conversation about your AI governance situation and what it would take to get to a defensible posture. Thirty minutes. We'll tell you exactly where you stand.

`sentinelsquared.com`

Use the contact form to request your Risk Briefing · Reference: 7-Questions Self-Assessment