The Missing Layer in Ten Pages

A free briefing on AI governance: what it is, what it isn't, and how to prove it

Lindsay Hiebert, CISSP

Trust is not the absence of doubt. It is the presence of evidence.

A free briefing from the book *The Missing Layer: What AI Governance Is, What It Isn’t, and How to Prove It*. Copyright 2026 Lindsay Hiebert, PIGENAI LLC. You are welcome to share this briefing in full. It is for education and is not legal advice.

1. The question

Sooner or later, someone will ask your organization a simple question: are we governing AI? It may come from a board member, an insurance broker, a large customer or a lawyer. The words change, but the question underneath stays the same. Can you show what AI you use, who agreed to it, on what terms, and when?

Most organizations can answer part of that out loud. Very few can hand over a document. That gap is not a gap in talent or effort. AI simply arrived faster than the record keeping around it. This briefing gives you the whole picture in ten pages so you can start closing the gap this week.

AI governance, in one sentence: the decisions an organization makes about how AI may be used, the people who own those decisions, and the dated records that show both.

2. The missing layer

Every organization that uses AI already does two things. It observes, and it enforces. The third job, governing, is the one that usually has no home.

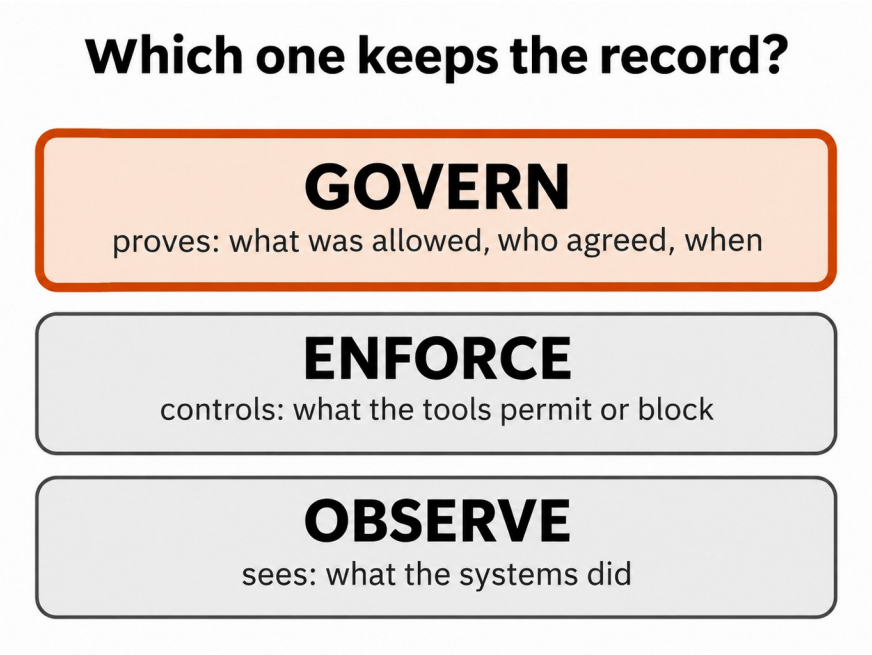

Three layers: Govern, Enforce, Observe

- **Observe.** Logs and monitoring show what the systems did.
- **Enforce.** Firewalls, access controls and guardrails permit some things and block others.
- **Govern.** Dated records show what was allowed, who agreed, and when.

The three layers work best together. Governance tells enforcement what to enforce. Observation tells governance whether reality still matches the record. The tools you already own are valuable, and nothing here replaces them. Governance is the layer you add above them, and it is made of documents that people write, sign and keep current.

A quick check: pick one log entry, one security rule and one approval. Can anyone show how all three relate to the same use of AI?

3. Due diligence and due care

Two phrases appear in almost every governance conversation, and they are worth separating.

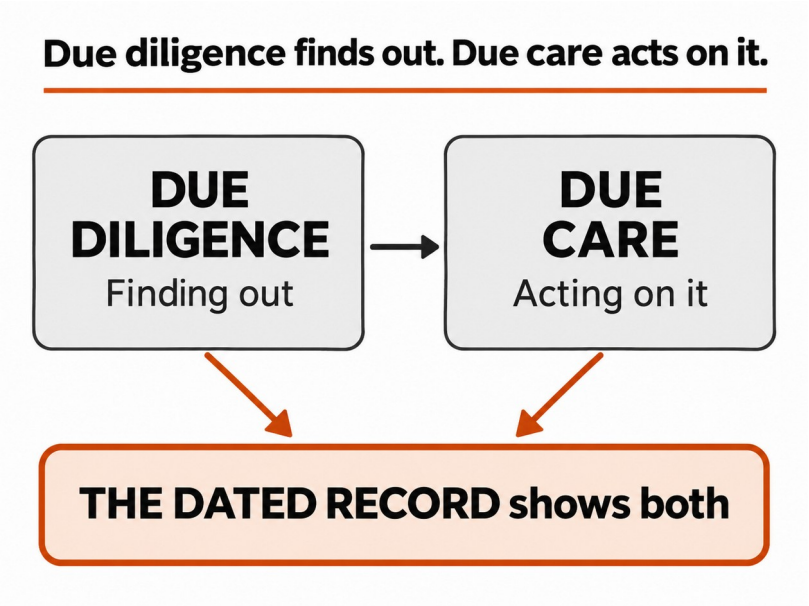

Due diligence finds out. Due care acts on it.

Due diligence is finding out. What AI is in use, what it touches, and what could go wrong. **Due care is acting on what you found.** Rules, owners, training, and risks that are either reduced or accepted on purpose. The dated record is how you show both.

Due diligence is a rhythm, not a single event. AI changes every month, so the finding out repeats, and each round leaves a new dated record. Good judgment that leaves no record cannot be shown to anyone later.

4. The surface you cannot see

AI in use that nobody has written down is often called shadow AI. It is rarely a sign of wrongdoing. People found a tool that helps them work. The task is to see it, decide about it and record the decision.

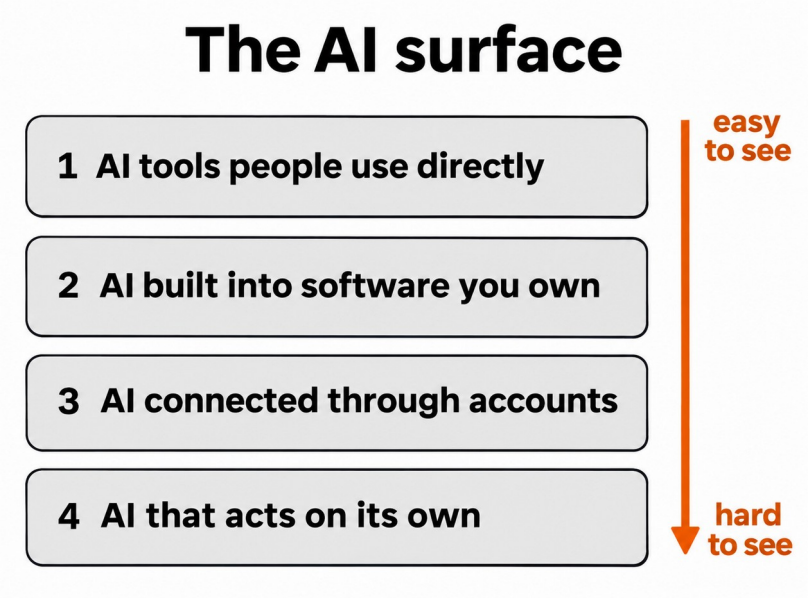

The four layers of the AI surface

1. **AI tools people use directly,** in a browser or an app.
2. **AI built into software you already own,** often switched on in an ordinary update.
3. **AI connected through accounts and integrations,** sometimes with one click.
4. **AI that acts on its own,** holding access and taking actions.

Each layer is harder to see than the one above it, so discovery works best when it compares two things: what people report, and what your systems record. Look in four places: money (what is paid for), access (what is connected to work accounts), people (what staff say they use) and network (what the logs show). Label each item by how you know it: observed, reported or stated.

5. How governance runs: five stages

Stage	What happens	What it leaves behind
Discover	Find the AI already in use	A dated inventory
Assess	Look at it across twelve areas	Ranked findings with owners
Establish	Write the rules and record the decisions	A signed policy and decisions
Prove	Date the records so others can check them	Evidence a third party can confirm
Sustain	Repeat on a rhythm and when things change	A series of dated records

Every AI use that comes up for review should leave with one of five decisions, each with a name and a date: approve, approve with conditions, defer with a date, reject, or retire.

6. Twelve areas worth an answer

A complete review covers twelve areas. Each is written as a statement you would like to be able to make. For each one, ask: can we say yes, and can we show it?

1. We know which AI we use.
2. Work accounts, not personal ones.
3. Agreed tools on the right terms.
4. We know what information goes in.
5. Clear rules everyone has agreed to.
6. Clear ownership and decisions.
7. AI that acts on its own is controlled.
8. Fair decisions about people.
9. Open and checked with customers.
10. People know how to use AI safely.
11. Work continues if AI stops or changes.
12. Ready if something goes wrong.

Most organizations can show two or three the first time. That is a normal place to begin. Watch for small findings that combine: a shared account, broad access, no training and no sign-in alerts may each be acceptable alone and form one serious risk together.

7. The five records

At the end of a first round, an organization should be able to produce five records.

1. **An Executive Risk Report.** The findings, ranked, with owners and a plan.
2. **An AI Acceptable Use Policy.** The rules, signed and acknowledged.
3. **A Board Memo.** One page that puts AI governance on the board's record.
4. **A dated proof that the assessment happened.** Something a third party can check without seeing the contents.
5. **Role-based training with a record of completion.** Who learned what, for which role, and when.

One point deserves care. Proof that a record was issued on a certain date is valuable, and it is not the same as proof that everything inside it is true. A good record is clear about both, and it states what is still open. A report that says "nine of eleven reviewed, two due December 15" earns more trust than one that is green everywhere.

8. When AI acts

An agent is AI that holds some access of its own, chooses steps and carries them out. Deploying one is delegation, and delegation has always come with three questions: what may it do, who answers for it, and how is the authority taken back? For each agent, record five things.

1. **What it may do,** and what it may not.
2. **Who owns it,** by name.
3. **What access it holds,** issued by whom, and when it expires.
4. **Where a person steps in** before an action happens.
5. **Who can stop it, and how.**

Every action an agent takes should trace back through its credential to an accountable person.

9. Your first 90 days

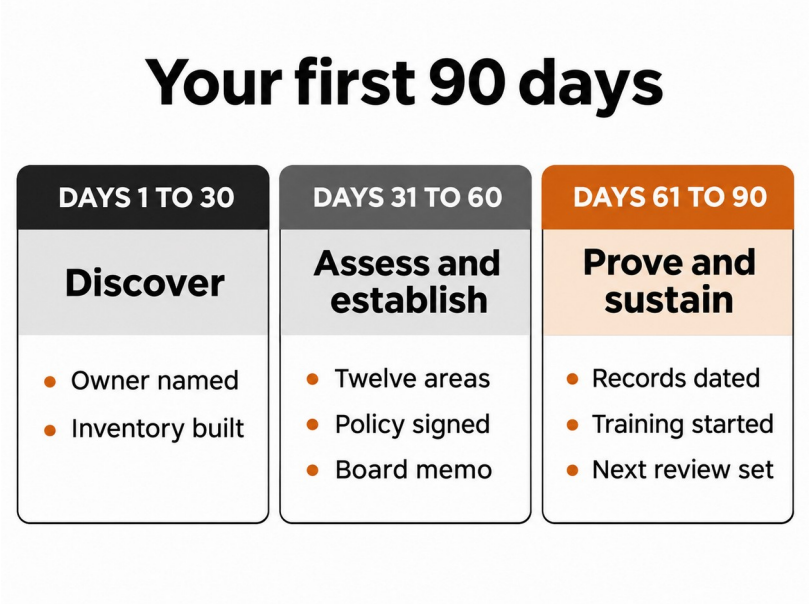

Your first 90 days

Days 1 to 30, discover. Name one accountable owner. Gather evidence from money, access, people and network. Build the first inventory.

Days 31 to 60, assess and establish. Work through the twelve areas. Rank the findings and name an owner for each. Sign the policy. Fix or formally accept each finding. Send the one-page memo to the board.

Days 61 to 90, prove and sustain. Date every record. Collect acknowledgments and begin training by role. Run one short tabletop exercise. Put the next review on the calendar.

Take one AI system all the way through first. One complete example teaches the organization more than a dozen half-finished ones.

10. Five questions for the card in your wallet

Whenever a new AI tool, model, vendor or agent arrives, come back to these.

1. What AI do we use?
2. What may it do?
3. Who agreed?
4. What evidence supports the decision?
5. What would make us look again, or stop it?

Where to go next

Read the full book. *The Missing Layer* follows one company, Calder Ridge Supply, from a board member's question to a complete, dated record. It adds the chapters on proof, regulation and insurers, agents, roles and training, a twelve-area self-check, a board conversation guide, record templates and a glossary. It is available on Amazon in Kindle and paperback: https://sanctumshield.com/the-missing-layer

See where you stand. The free Shadow AI Check takes a few minutes in your browser: https://sanctumshield.com/check

Train by role. The open syllabus for SanctumShield Academy: https://academy.sanctumshield.com/curriculum

Lindsay Hiebert, CISSP, is Founder and CEO of PIGENAI LLC, which builds SanctumShield and SanctumShield Academy.