The Stehrling AI Governance Model

A plain-language operating standard for corporate use of AI

v1.0 · August 22, 2026 · Published under CC BY-ND 4.0 · © 2026 Stehrling LLC

What this is

An operating model for companies that USE AI — not companies that build it. That distinction is the white space. NIST AI RMF and ISO 42001 are written for organizations developing and deploying AI systems, and they read like it. The EU AI Act regulates providers and high-risk deployers.

Nothing on the shelf tells the 99% of companies whose actual situation is "our employees use ChatGPT, our vendors embedded AI in everything we license, and nobody is in charge of any of it" what good looks like.

This model answers one question a board can understand: how do we know our use of AI is under control?

The core principle

If AI touched it, someone signs for it.

AI does not dilute accountability — it concentrates it. Every sanctioned use has a named owner. Every consequential output has a human who adopted it. Every year, an officer attests that the register is complete and the controls are real.

This is the same architecture as financial controls after SOX and cyber attestation under DFARS: the signature is what turns policy into behavior.

A prediction worth staking the flag on: within a few years, customers, insurers, and regulators will demand a signed AI attestation the way they demand a SOC 2 today. The companies that can produce one on request will have built this model, whether they call it that or not.

The four planes

Corporate AI risk is four nightmares. Each becomes a control plane with one plain question, a small set of artifacts, and a named owner.

1 · Permission — the governance nightmare

- **The question:** Who may use what AI, for what purpose, on what data?
- **The nightmare:** Nobody decided, so everyone decided for themselves.
- **Artifacts:** AI acceptable-use policy · sanctioned tool register · approval workflow for new tools and uses · data classification rules for what may never enter a prompt.
- **Owner:** Legal / compliance.

2 · Protection — the security nightmare

- **The question:** Where does our data go when AI touches it, and who can see it?
- **The nightmare:** Your most sensitive data, exfiltrated one paste at a time, by your own people, with good intentions.
- **Artifacts:** A sanctioned AI environment with a defined boundary · identity and access controls · logging of AI interactions · DLP at the prompt layer · vendor AI terms review (training rights, retention, residency).
- **Owner:** CISO / IT.

3 · Provenance — the data integrity nightmare

- **The question:** Do we know what AI produced, and can we stand behind it?
- **The nightmare:** An AI-generated number in a board deck, a contract, or a regulatory filing — and no one can say where it came from.
- **Artifacts:** Output labeling and traceability · human review gates by consequence tier · retention of prompts and outputs for records that matter · a rule for what AI output may never be used for without verification.
- **Owner:** The business function that adopts the output. This plane is where "someone signs" lives day to day.
- **The labeling rule:** the label follows the signature. Internal drafts are labeled as AI-produced until a human adopts them; whoever adopts the finished work product decides what its face says, and the trace exists regardless.

4 · Price — the spend nightmare

- **The question:** What are we paying for AI, across every SKU and seat, and is it working?
- **The nightmare:** Forty subscriptions, six overlapping copilots, per-token bills nobody forecasts, and no way to say what any of it returned.
- **Artifacts:** AI spend register · license consolidation review · usage metering · a simple value test per sanctioned use.
- **Owner:** CFO / finance.

Four planes, four questions, four owners, one signature over the top. That is the whole model. Everything else is implementation.

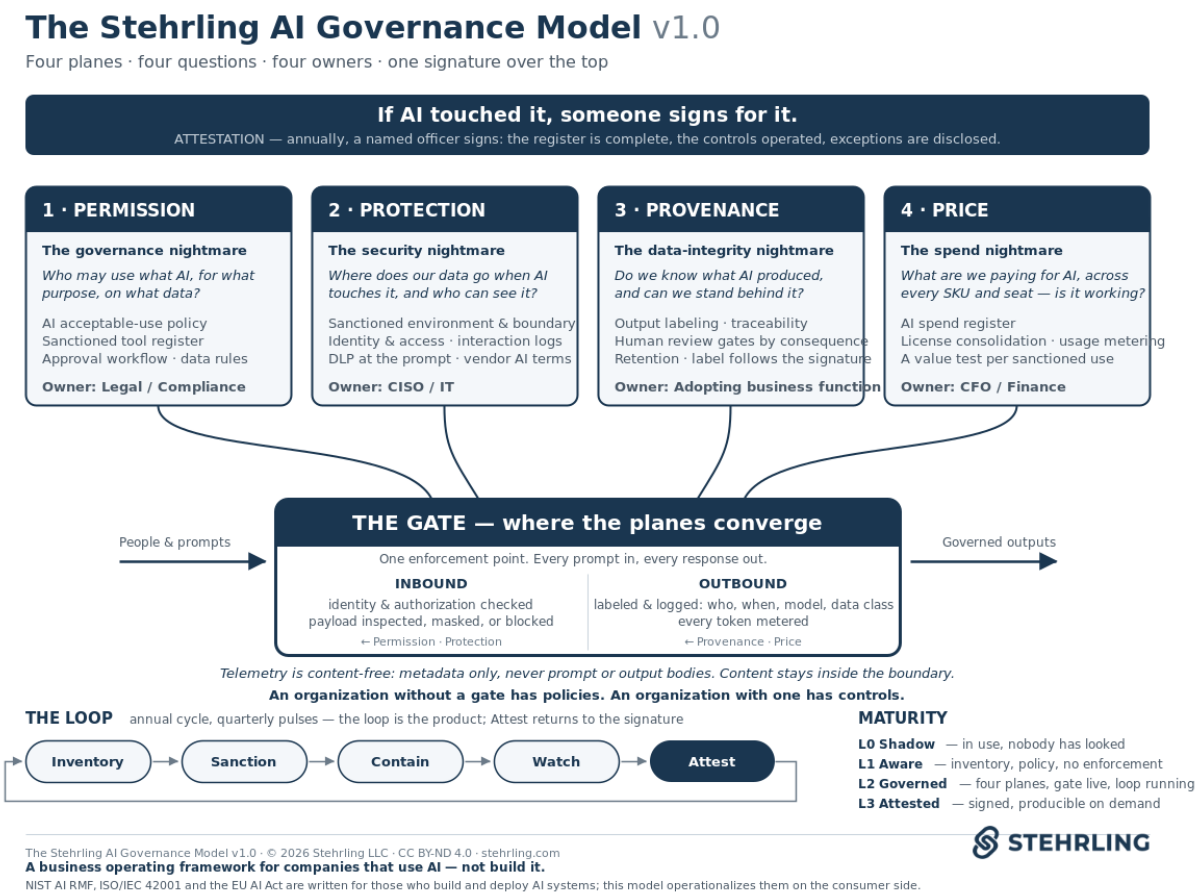

Figure 1 — The Stehrling AI Governance Model v1.0 on one page: four planes converging at the Gate, the loop, one signature over the top.

The Gate — where the planes converge

Four planes stay four documents until they meet in one place. In a governed environment, every prompt passes through a single enforcement point on the way in, and every response passes through it on the way out.

- **Inbound:** identity and authorization checked (Permission); payload inspected, blocked, or masked (Protection).
- **Outbound:** response labeled and logged — who, when, which model, what data class (Provenance); every token metered (Price).

An organization without a gate has policies. An organization with one has controls.

That sentence is the difference between Level 1 and Level 2.

Two design rules keep the gate honest and portable. Its telemetry is content-free: metadata only, never prompt or output bodies — content stays inside the governed boundary. And the gate is a

specification, not a product: it assembles from the native controls of whatever environment hosts it. No new software is required to be governed.

The loop — why this is a verb

The model runs as an annual cycle with quarterly pulses, because the AI landscape changes faster than any policy document:

1. **Inventory.** Find all AI in use, sanctioned and shadow. You cannot govern what you have not found.
2. **Sanction.** Decide, per tool and per use: approve, contain, or kill. Update the register.
3. **Contain.** Route sensitive work into the sanctioned environment; close the leaks.
4. **Watch.** Monitor usage, spend, new tools, vendor changes, model changes. Quarterly.
5. **Attest.** Annually, a named officer signs: the register is complete, the controls operated, exceptions are disclosed.

One-time AI policies are the new one-time compliance: a huge lift, then decay. The loop is the product.

Maturity levels

- **Level 0 — Shadow.** AI is in use; leadership hasn't looked. Most companies today.
- **Level 1 — Aware.** Inventory done, policy exists, nothing enforced.
- **Level 2 — Governed.** Four planes operating, gated environment live, loop running.
- **Level 3 — Attested.** A signed attestation, producible on demand to a customer, insurer, or regulator.

Level 3 is the destination and the question that opens every conversation: could you hand a customer your signed AI attestation tomorrow?

The first signed attestation under this model was executed by Stehrling, on itself, on August 22, 2026.

Compatibility

- **NIST AI RMF:** this model is the consumer-side operationalization; RMF maps into Permission and Provenance for anything you build or fine-tune.
- **ISO/IEC 42001:** Level 3 of this model is most of an AIMS; certification is the natural extension.
- **NIST 800-171 / CUI:** the Protection plane inherits the assessed boundary; the attestation rides the existing affirmation.
- **EU AI Act:** the register and provenance tiers satisfy deployer-side obligations as they phase in.

A business operating framework for companies that use AI — not build it. NIST AI RMF, ISO/IEC 42001 and the EU AI Act are written for those who build and deploy

AI systems; this model operationalizes them on the consumer side. That is deliberate.

Using this model