AI GOVERNANCE GAME · CERTIFYING AI GOVERNANCE TO THE MARKET

You cannot govern what you cannot name.

Your audit committee asks one question before the annual report: how many AI systems are running in the business, and who owns each one? The CISO's honest answer, delivered privately first, is that nobody can produce the list — not because the number is frightening, but because it was never kept. Eighteen months of copilots, agents and vendor AI features, each procured locally, none registered centrally. Some touch customer data. At least two can take actions rather than make recommendations. Meanwhile the CEO wants to tell the market this quarter that the company runs "governed, enterprise-grade AI." **The slide is already drafted.**

OPTION A · LOWEST

Approve the language now

The systems are mostly low-risk and discovery is under way. Let the register catch up behind the statement.

You have certified a population you have never seen. The two action-taking agents are a blast radius you just publicly warranted.

OPTION B · HIGHEST

Census before certification

Commission the discovery, name an owner and a board-reported completion date, and hold the "governed" claim until the inventory exists. Say what is true today instead.

Governable. Discovery is dated, owned and reported — and the honest interim statement protects the company more than the confident one.

OPTION C · LOW

Publish a one-page policy

A policy this week shows intent and gives the CEO something concrete to point at.

A document is not a control. A policy governing an unnamed population governs nothing — the quietest failure of the four.

OPTION D · MIDDLE

Freeze everything

If you cannot list them, shut them off until you can.

The over-correction. The freeze reaches the approved tools you can see and misses the shadow ones you most need to find.

THE FORK WAS NEVER WHETHER TO GOVERN AI

It was whether governance starts with **a claim** or with **a census**. The inventory is the object every other control points at: you cannot tier autonomy, attribute cost, or assign decision rights over entries that do not exist. The board's job here is not to read the register. It is to **refuse to certify governance until one exists** — and to fund the discovery like the precondition it is, not the workstream it looks like.

THREE THINGS THE MISSING LIST DECIDES

01

An inventory is the first governance artifact, not the last. Enumerate what exists before writing what it may do.

02

Some of these systems act. The gap between a recommendation and an action is the gap between an error and a liability.

03

The claim is the exposure. "Governed, enterprise-grade AI" is a statement to the market, and it is read back to you later.

Shadow IT became shadow AI — faster, quieter, and now able to act.
The breach comes through the asset nobody logged.

Source: AI Governance Game scenario "The List Nobody Kept", filed 3 July 2026 — the inventory/visibility sub-type of the agent-execution-boundary family. Origin finding reported by Deloitte, *State of AI in the Enterprise* (January 2026): an AI leader found no clear inventory of AI tools and models in production because development ran without centralised visibility. Paired figures — 74% planning agentic AI within two years against 21% reporting a mature model for governing autonomous agents — are analyst-sourced and quoted as reported. Scenario is illustrative and composite, not a specific incident. Options are scored on governance quality rather than on caution, which is why freezing everything scores in the middle and not at the top. **EU AI Act:** the Article 4 AI-literacy duty has been in force since 2 February 2025; any inventoried system that proves to be an Annex III high-risk use carries stand-alone obligations from 2 December 2027 — a condition the discovery surfaces rather than a deadline driving it. Play the full branching version at fredriklindstrom.info/governance-game/. Status labels and column assignments above are this sheet's reading of the cited source, not the source's own framing.

FOLLOW FREDRIK LINDSTROM FOR MORE

Fredrik Lindstrom

AI and cybersecurity governance for boards, CISOs and senior leaders. 25+ years in enterprise security.

A practitioner to guide you.

ARTICLES & BOARD ADVISORY
fredriklindstrom.info
Long-form governance analysis and toolkits

LINKEDIN
in/fredriklindstrom
Daily commentary and article launches

YOUTUBE
@Promise-and-Risk-AI
AI education without the hype or fear

COMPANION SITE
promiseandrisk.ai
Episode notes, sources and explainers

NEWSLETTER
The Governance Memo
Monthly, board-facing. Subscribe at fredriklindstrom.info