



Governance Is Not Cybersecurity

Why People-first, Human-in-the-Loop AI Governance Wins

By

Michelle Strasburger, CEO
AI Rebooted Incorporated

Executive summary

Cybersecurity protects your systems from AI. Governance makes sure AI does the right work, the right way, with a person accountable for the outcome. Most organizations have invested in the first and assumed it covers the second. It does not.

New research makes the cost of that assumption clear. The 2026 Co-Intelligent Workforce Report by Lever Talent, Inc. surveyed 484 Director-and-above leaders at mid-market companies. Companies that invest across technology, programs, and people reach extreme value from AI 50% of the time. Technology-only adopters get there 5.7% of the time.

The 2026 Co-Intelligent Workforce Report by Lever Talent, Inc. also found that 61% of companies using agentic AI can't say a person reviews every autonomous workflow before the work moves on. That is not a security gap. It is a governance gap, and it is where liability now lives.

This is not just a mid-market problem. The EY AI Risk and Governance Survey of 202 senior AI executives at companies with \$1 billion or more in revenue found that 98% have formal AI governance policies, yet 47% have skipped their own governance process for urgent deployments. Among those using agentic AI, 85% say at least some of those systems take action without real-time human involvement.

AI is not a technology rollout. It is a workforce transformation driven by technology. This paper explains the difference between governance and cybersecurity, what the data says about people-first adoption, and how AI Rebooted keeps humans in the loop where the work actually happens.

Governance and cybersecurity are not the same thing

We get this question often. A leadership team has locked down data access, approved a vendor list, and blocked a few consumer tools. They believe AI is governed. What they have is a secure perimeter around AI that nobody is supervising.

Cybersecurity asks: can someone break in, and can data get out? Governance asks a bigger set of questions. Is the AI output accurate? Is it fair? Who approved this use? Who reviews the work before it reaches a customer, an employee, or a regulator? What happens when it goes wrong?

	Cybersecurity	AI governance
Core question	Is the system protected?	Is the work right, and who is accountable?
Primary threat	Attackers, data leakage, unauthorized access	Hallucination, bias, drift, bad decisions at scale
Typical owner	IT and security	A cross-functional governance committee, with HR at the table
Where it operates	Networks, endpoints, identity	Chats, agents, and vibe-coded tools where work gets done

	Cybersecurity	AI governance
Measure of success	No breach	Trusted output, documented oversight, adoption that sticks
Role of people	Users to protect against	Accountable owners of every AI-assisted outcome

Cybersecurity belongs inside governance. Prompt injection and jailbreaking are real risks, and a governance program has to address them. But a perfectly secure agent can still invent a policy, deny a qualified candidate, or send a customer the wrong answer. Security keeps the instruments safe. Governance makes sure someone is conducting the orchestra.

The EY survey backs this up. Cybersecurity was the most common AI risk respondents encountered last year (52%), but human risk (47%) and shadow AI (46%) were close behind. And when organizations ran formal AI assurance reviews, the most common problems were data quality (57%), model drift (48%), and shadow AI (39%). None of those are stopped by a firewall. All of them are caught by governance with a person watching.

The data shows where the gap comes from. According to the 2026 Co-Intelligent Workforce Report by Lever Talent, Inc., the ownership seat for how people and AI work together sits in IT 48% of the time. HR holds it less than 10% of the time. When AI is owned as a technology, it gets governed like one.

What the research shows: people-first wins

Companies that treat AI as a workforce transformation get dramatically better results than companies that treat it as a software install. The 2026 Co-Intelligent Workforce Report by Lever Talent, Inc. surveyed 484 Director-and-above decision-makers at US and Canadian companies with \$10M to \$500M in revenue. Findings are reported as associations, not causation, and segment differences only were significant at 95% confidence.

Finding	The number	What it means for governance
The Adoption Illusion	97% have adopted AI and 92% are raising spend, yet only 29% call it extremely valuable	Buying AI is not the same as getting value from it
The Tech-Only Dead End	Technology-only adopters reach extreme value 5.7% of the time; companies investing in technology, programs, and people reach it 50% of the time	People and programs are the difference, not the tools
The Accountability Gap	45% have no single owner for how people and AI work together; only 21% have someone whose main job it is	Governance without an owner is a document, not a practice

Finding	The number	What it means for governance
The Orchestrator's Dividend	A primary owner makes extreme value 2.4x as likely, and 4.4x versus companies where no one owns it	Accountability is a value driver, not overhead
Design Decides	Clear human-AI division of labor plus clear error accountability makes extreme value 5.3x as likely; half of companies haven't done this work	Deciding who reviews what is the highest-leverage step

Read those together and the message is hard to miss. The organizations getting value from AI decided who owns it, who checks it, and who answers for it. That is governance. None of it is solved by a firewall.

This matches what we see with clients. When a rollout leads with people, employees understand the guardrails, trust the tools, and use them. When it leads with technology, adoption stalls and shadow AI fills the gap. Employees find their own tools because the sanctioned path doesn't fit how they work.

Source: 2026 Co-Intelligent Workforce Report by Lever Talent, Inc. (fielded September 2026; research design by Haus of Bold, fielding by Researchscape). Used with attribution under the expert contributor terms.

The oversight gap: autonomy is outrunning accountability

61% of companies using AI that completes multi-step work on its own can't say a person reviews every such workflow before the work moves on. Agents are drafting, deciding, and acting, and in most organizations nobody is checking the handoff.

"This number should alarm every leadership team. AI agents can hallucinate, carry bias, and act on bad data, and when no one reviews their work, those errors flow straight into the business. Regulators and courts are making it clear that when an agent goes rogue, the liability falls on the organization that deployed it, not the model or the vendor that built it. Running agents without governance that includes human oversight is playing Russian roulette, and the next class action is a matter of when, not if."

Michelle Strasburger, Co-Founder and CEO, AI Rebooted

An agent with no reviewer is not a productivity gain. It is an unsupervised employee with access to your systems and no manager. Every error it makes compounds as the work moves downstream.

This is exactly where security-only governance falls short. A secure agent is still an agent. Encryption and access controls do nothing to stop a confident, wrong answer from reaching a customer.

Two studies, one conclusion

Two independent surveys, fielded months apart with very different companies, point to the same gap.

	2026 Co-Intelligent Workforce Report by Lever Talent, Inc.	EY AI Risk and Governance Survey
Who was surveyed	484 Director+ leaders at \$10M to \$500M companies	202 senior AI executives at \$1B+ public companies
Agents without review	61% can't say a person reviews every autonomous workflow	85% have agents that act without real-time human involvement
Policy vs. practice	45% have no single owner for people and AI working together	98% have policies, but 47% have bypassed them for urgent deployments
Readiness for agents	Half haven't defined who does what or who owns errors	49% haven't updated governance for agentic AI; 26% can't detect unauthorized agents
What it costs	Technology-only adopters reach extreme value just 5.7% of the time	36% have had an AI incident with material negative impact

EY put it plainly:

"The biggest agentic AI risk is that human oversight hasn't evolved accordingly."

John McLain, EY Americas Assurance Technology Risk AI Leader

The good news is in the same data. Among EY respondents who ran formal AI assurance reviews, 64% significantly modified a quarter or more of their AI systems, and 29% paused a quarter or more. Oversight works when it actually happens. The problem is that most organizations only look once a year, while their agents work every minute.

Legal flag: *This section discusses legal liability and regulatory direction in general terms. It is not legal advice. Organizations should consult counsel about their specific obligations.*

Human in the loop, by design

At AI Rebooted, we build governance so a person is accountable for every AI-assisted outcome, and we put that governance where the work actually happens. As my colleague Meghan Gilson puts it, the human isn't just in the loop. The human IS the loop.

Most governance lives in a policy binder. Ours lives in the flow of work: the chats employees open every day, the agents running multi-step workflows, and the vibe-coded tools teams build on their own. Governance 1st, our AI Governance Operating System, is grounded in the NIST AI Risk Management Framework and works in three layers.


1. **Set it up.** We help the organization stand up a governance committee, define who owns which AI uses, and decide where a human must review before work moves on. This is the design work the research says half of companies have skipped.
2. **Operationalize it.** Governance 1st and its browser extensions apply guardrails in real time inside ChatGPT, Claude, Gemini, Microsoft Copilot, and hundreds of other AI tools, so policy shows up at the point of use.
3. **Monitor and escalate.** Our patent-pending technology continuously evaluates AI activity against more than 200 guardrails in ten behavioral categories. When something flags, the governance committee is notified for corrective action. A human decides what happens next.

The ten behaviors we monitor:

Category	What it catches
Hallucination	Confident answers that aren't true
Bias	Outputs that treat people or groups unfairly
Behavioral drift	A model that behaves differently over time
Prompt injection	Hidden instructions that hijack the AI
Jailbreaking	Attempts to get around safety controls
Refusal calibration	Refusing valid work, or accepting what it shouldn't
Bigotry	Discriminatory or hateful content
Logical fallacies	Flawed reasoning presented as sound
Partisanship	Political slant in business outputs
Misinformation	False or misleading claims

Notice that only two of the ten, prompt injection and jailbreaking, are classic security concerns. The other eight are about whether the work is right. That is the part most governance tools on the market miss, because they were built to solve a cybersecurity problem.

A person decides what happens when AI misbehaves



Governance 1st oversight loop · 5 steps

Every flag routes to people, and every fix feeds back into the policy and guardrails, so governance gets sharper as AI use grows.

Start with the foundation: Policy FastTrack

Every governance program starts with a policy that says who can use AI, for what, and who reviews the result. Without it, there is nothing to operationalize and nothing for the committee to enforce. Yet many mid-market organizations still don't have one, or have a generic template that doesn't reflect the laws they actually operate under.

Having a policy isn't the same as having one that works. EY found that about two-thirds of senior AI executives worry their organization lacks the internal expertise to design (63%), implement (63%), or evolve (69%) AI governance controls. A policy nobody can operationalize is a policy that gets skipped when deadlines hit.

Policy FastTrack closes that first gap fast. It gives organizations a starting AI policy built around the regulations that apply to them, so leaders can move from "we should have a policy" to a working foundation without months of drafting. Learn more at laws.ai-rebooted.ai.

A policy is step one, not the finish line. The research is clear that design and accountability drive value, not documents alone. Policy FastTrack gets the rules written. Governance 1st puts those rules to work in every chat, agent, and vibe-coded tool, and puts a human in charge of what happens when something goes wrong.

Draft note for Michelle: Please confirm Policy FastTrack's feature description, and pricing or offer details before publishing. Policy FastTrack launches October 28, about a week after the report publishes the week of October 19, which makes a strong combined announcement window.

Five questions for your leadership team

If you can't answer yes to all five, your AI is secured but not governed.

- ☐ Do we have a written AI policy that reflects the laws we operate under?
- ☐ Is one named person or committee accountable for how people and AI work together?
- ☐ For every agent running multi-step work, do we know who reviews it before the work moves on?
- ☐ Are we monitoring AI output for accuracy, bias, and drift, not just data leakage?
- ☐ When AI gets something wrong, does a human get notified and own the fix?

About AI Rebooted

AI Rebooted is an AI governance and workforce transformation company. We help organizations adopt AI responsibly by setting up governance, operationalizing it where the work happens, and keeping people accountable for every AI-assisted outcome. Our Governance 1st platform is grounded in the NIST AI Risk Management Framework and includes pre-built documentation aligned with NIST AI RMF, ISO/IEC 42001, GDPR, and the EU AI Act.

AI is the orchestra. Your people are the conductors. We make sure the music is right.

Schedule a demo: calendly.com/scott-berry-dsm/opportunities-discussion

Explore Policy FastTrack: laws.ai-rebooted.ai

Sources

- *2026 Co-Intelligent Workforce Report by Lever Talent, Inc.*, Expert Contributor Brief (September 2026). Full report publishing the week of October 19, 2026. Methodology and crosstabs available from Lever Talent on request.
- Ernst & Young LLP, *EY AI Risk and Governance Survey*, press release, September 15, 2026. Online survey of 202 US senior AI decision-makers at publicly traded companies with \$1B+ annual revenue, fielded May 28 to June 15, 2026; margin of error plus or minus 7 points at 95% confidence. ey.com/en_us/newsroom/2026/09/ey-survey-finds-that-autonomous-ai-implementation-outpaces-oversight-yielding-an-ai-governance-gap