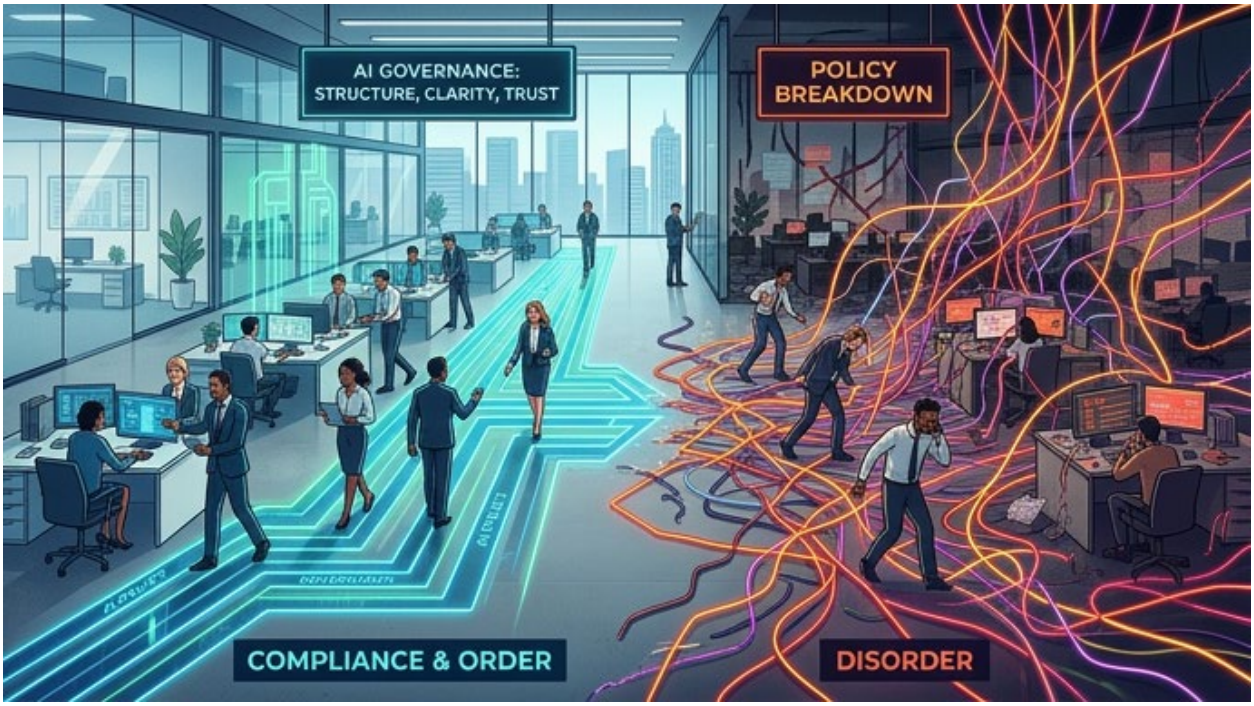

AI GOVERNANCE SERIES | ADVISORY PERSPECTIVE

Your Employees Aren't Ignoring AI Policy.

Your Operating Model Hasn't Caught Up with Them.

By Greg Aldrich | Global CIO & Strategic Advisor | April 2026

Barndoor AI published a report this month with a headline designed to alarm: "Half your workforce is ignoring internal AI policies."

The data behind it is real and worth taking seriously. In a survey of 155 professionals conducted in February and March 2026, they found:

- **1 in 3** employees connected AI to their work email
- **1 in 5** connected AI to internal databases
- **1 in 10** shared API keys or developer credentials with AI tools
- **Nearly 1 in 10** connected customer data to an AI system
- **28%** are now using AI that takes autonomous actions — not just generating suggestions, but executing tasks

And the Model Context Protocol — the open standard that lets AI agents connect directly into enterprise systems — has scaled from roughly 100 servers at its launch in late 2024 to over 16,000

today. That means employees can now wire AI into enterprise applications in a few clicks, without IT involvement, without approval, and without oversight.

The report frames this as a governance failure. Employees are bypassing policy. Organizations are running blind.

I'd frame it differently.

> ***"This isn't primarily a governance failure. It's an operating model failure under new conditions — and the distinction matters for how you respond."***

What the Data Actually Shows

The most important finding in this report isn't the behavior. It's the reason for the behavior.

When employees were asked why they use non-approved AI tools, nearly 50% cited ease of use as the main reason — higher than capability, speed, or even job pressure. One respondent put it plainly: "The company-provided AI cannot think correctly and give out proper responses for the question and issue. I want to use tools that are easy and effective."

Another: "AI isn't restricted; there are no approved AI tools, it's every man for himself."

And another: "My job expects me to get tasks done very quickly and accurately, but doesn't want us to use AI tools for 'plausible deniability' reasons. There are no published rules, so one doesn't know they've broken a rule until they are punished after the fact."

These aren't employees ignoring policy. These are employees optimizing for productivity in the absence of usable governance. That is a meaningfully different problem — and it requires a meaningfully different response.

> ***"People bypass controls when controls are not aligned with outcomes. That is not a compliance problem. It is a design problem."***

The Headline Misframes the Cause

The report's implicit conclusion — shared by most security and compliance narratives — is that the answer to shadow AI is greater visibility, stronger access controls, and better governance tooling. Barndoor's own product is positioned as a centralized control plane to govern AI agents across every tool and platform.

I don't disagree that visibility and control matter. The API key statistic alone is a serious signal: an employee who has given an unsanctioned AI tool access to a production API credential has granted that tool the ability to authenticate as the enterprise in any system that credential covers. That is not a minor exposure.

But more control layered on top of a misaligned system doesn't fix the misalignment. It adds friction to both the sanctioned and unsanctioned paths — and if the sanctioned path is harder to use, employees will continue to choose the easier one.

The report itself surfaces this tension without quite naming it: 78% of employees work with no clear AI policy, and nearly 20% don't know what their organization's policy even is. You cannot enforce

what hasn't been communicated. And you cannot communicate policy that hasn't been designed around how work is actually done.

Three Reframes Worth Considering

If your organization is looking at this data and asking, "How do we stop employees from doing this?" I'd suggest three reframes before you reach for the control lever.

1. Redesign workflows, not just controls

The employees in this report aren't choosing shadow AI because they want to create risk. They're choosing it because it helps them do their jobs. The right response to that signal is to make the approved path competitive — not to add friction to the unapproved one.

That means embedding AI into approved processes in ways that are genuinely useful. It means making sure that sanctioned tools are as capable and easy to use as the alternatives employees are already using. And it means designing AI governance into the workflow at the point where work happens, not as a gate that sits outside it.

Governance that exists only as policy will not survive operational reality. Governance embedded in workflow will.

2. Move control upstream

The shift from generative to agentic AI fundamentally changes the risk profile. When AI only answers questions, the human decides what to do with the output. When AI takes actions — sending emails, modifying database records, executing code, changing CRM data — the human is no longer the last line of defense.

The 28% of employees already using agentic AI represent the leading edge of a shift that will accelerate significantly. Governance frameworks designed for prompt-and-response AI are not adequate for action-taking AI. The control model has to move upstream — to the point at which the boundaries, permissions, and acceptable actions are defined before the agent is deployed, not discovered after something goes wrong.

This is precisely the argument I made in Part 2 of this series: specifications are the governance layer. Prompts guide behavior. Specs enforce it. The same logic applies here at the organizational level — you govern agentic AI by defining what it is and isn't permitted to do before it acts, not by monitoring what it did after the fact.

> ***"Governance frameworks designed for AI that 'answers' are not adequate for AI that 'acts.' The control model must move upstream."***

3. Accept and manage edge behavior

You will not eliminate shadow AI. The connection friction has collapsed too far. The tools are too accessible and too capable. The productivity incentive is too strong.

The more realistic and more productive framing is: how do we bound it, monitor it, and design around it? That means formulating clear categories of acceptable and unacceptable use — not as a policy

document that employees never read, but as embedded guidance at the point of decision. It means monitoring AI behavior at the action level, not just the access level. And it means accepting that the edge of your organization will always move faster than your governance framework, and designing for that gap rather than pretending it doesn't exist.

The Missing Piece: Operating Model Design

What this report correctly identifies — and doesn't fully name — is that AI governance failures occur at the decision layer, not the policy layer.

The questions it raises but doesn't answer:

- Who owns AI usage at the edge of the organization?
- When an employee exposes customer data through an unsanctioned tool, where does accountability sit?
- What is the escalation model when something goes wrong?
- What are the acceptable risk boundaries — and who defines them?

These are operating model questions, not technology questions. And they don't get answered by a control plane. They get answered by organizational design: clear ownership, explicit accountability, defined decision rights, and governance embedded into how work is structured rather than layered on top of it.

A note on the report's methodology: 155 respondents across a broad range of functions is a directional signal, not a definitive study. The findings are credible, and the qualitative responses are distinctively valuable — but the specific percentages should be treated as indicative rather than precise. The pattern they describe, however, is consistent with what I observe in client environments across industries. The data exhibit a real dynamic, even if the exact numbers warrant appropriate skepticism.

Connecting the Thread

Read this report alongside the UiPath agentic orchestration document I analyzed recently, and a clear pattern emerges.

UiPath says: organizations need a coordination layer for agents, robots, and humans. Barndoor says: organizations are losing control of AI at the edge. Both are describing symptoms of the same underlying gap: the absence of a coherent operating and control model for AI-enabled work.

More tooling is not the answer to that gap — at least not as the first move. The first move is organizational: define who owns AI governance, design the workflow so approved AI is the easiest path, move control upstream before agents take action, and establish accountability models that function at the speed AI actually operates.

The platform follows the operating model. It should never precede it.

What Leadership Should Be Asking

If your organization recognizes itself in this data, these are the right questions to bring to your next leadership conversation:

- Are our approved AI tools actually competitive with the alternatives employees are already using? If not, we're designing for non-compliance.
- Do our employees know the policy? If 20% don't, we have a communication failure before we have a compliance failure.
- Have we defined what agentic AI is and isn't permitted to do — at the action level, not just the tool level?
- Who owns AI governance at the operational layer? Not the committee — the day-to-day accountability for how AI is used in workflow?
- Are we monitoring AI behavior at the action level? Access controls tell you who got in. Action monitoring tells you what happened inside.

The Barndoor report is a useful mirror. The employees it describes are not irresponsible. They are rational actors responding to an organizational environment that hasn't adapted to the tools they're using.

The answer to that is not more policy. It's a better-designed system.

> ***"Governance isn't breaking because employees are irresponsible. It's breaking because the system they're operating in hasn't adapted. That is a leadership problem, not a compliance problem."***

Part of an ongoing series on AI governance and organizational readiness. Part 1: The AI Governance Tightrope | Part 2: The New Control Layer | Part 3: The Operating Model for AI Governance | Advisory: Agentic Orchestration — Questions Worth Asking

---

About the Author

Greg Aldrich is a Global CIO and Strategic Advisor with 30+ years of experience helping boards, executive teams, and C-suite stakeholders navigate IT strategy, AI governance, and digital transformation. He serves as Senior Strategy Advisor at Blue Tree Technology Group and Senior Transition Architect at SDS Consulting, and currently serves as CIO at Andrew Wommack Ministries & Charis Bible College, where he chairs both the AI Committee and IT Steering Committee. He has advised organizations across financial services, gaming, healthcare, higher education, logistics, and nonprofit sectors.

Connect: linkedin.com/in/galdrich