Barndoor.ai

ORGANIZATIONAL BEHAVIOR & AI GOVERNANCE

2026 | Report

Half your workforce is ignoring internal AI policies. Why is governance breaking down?

Barndoor.ai

Executive Summary

This research, commissioned by Barndoor and conducted with 155 professionals between February-March 2026, examines why employees connect AI to their most sensitive systems even when they understand the risks. The answer points to a governance problem that policy alone cannot fix.

Key Findings

- **9 out of 10 employees** are now using AI on the job, with nearly 30% of them moving towards agentic AI.
- **50% of employees** are giving it persistent access to at least one internal workplace application such as email, internal databases, and API keys.
- When asked to rank the riskiest agentic behaviors, respondents correctly identified making changes to internal systems and handling financial workflows as the **top two dangers.**
- **Organizations are struggling with governance and control.** Nearly half **(48.4%)** of the employees are using non-approved AI tools, selected because of their ease of use.
- Employees are working under confusion: **30%** say they work under loose AI governance. Nearly **20%** don't know what their organization's policy even is.

There is a governance gap that existing frameworks aren’t ready for. Identity alone isn’t enough to govern what AI agents do once it gets access. With **28%** of employees using agentic AI, this is a growing blind spot for enterprises.

Employees will pick the easier AI tool

It’s widely understood that most of us will choose convenience over security - your employees, likewise, will knowingly accept security risks in exchange for productivity. The research found something even more precise: when employees were asked what drives their use of non-approved AI tools, nearly 50% cited “ease of use” as their #1 reason. This reason beats out capability **(42%)**, speed **(38%)**, data privacy concerns **(38%)**, and even direct pressure to get work done **(40%)**. While not a new phenomenon, the stakes are higher in a multi-AI ecosystem.

49% cite ease of use as the primary reason for selecting non-approved AI tools, higher than capability, speed, or job pressure.

The implication is that if your company-approved AI tool is harder to use than whatever free tool already open in your employee's browser, policy will not override convenience. Your employee will generally pick the easier tool, ignoring whatever security policy you have in place.

Throughout the research process, respondents were prompted for qualitative feedback. When asked why they use unsanctioned AI tools, one respondent said: ***"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."***

Employees want choices

Think your workforce is following that "Claude-only" policy? This research found that employees have self-established a multi-tool environment for themselves, outside corporate security best practices and guidance. Over **58%** of employees say they use at least 2-3 different AI tools for work; **7%** use 4-6 tools. They'll use ChatGPT for general tasks, Claude for accuracy sensitive work, Gemini because it's integrated into Google Workspace, GitHub Copilot inside their dev environment. Each connection carries its own exposure risk, operates under different terms of service, and most likely none of them are being tracked.

Here's how respondents explained it:

> "I switched tools depending on the complexity of the task at hand. For simple tasks, ChatGPT would suffice. For more involved tasks that I relied on for accuracy, I used Claude or Gemini."

> "I switched from one tool to another because the other tool had more specific knowledge about the inner workings of our company, and I needed a question answered about my company and the process."

> "I use Claude for coding and ChatGPT for other tasks."

Employees are giving AI persistent access to internal data

A lot of the risk exposure conversation centers around what employees are pasting into prompts: financial information, customer records, internal documents. The research goes further - 50% of employees admitted to giving AI persistent access to at least one workplace application. This is real risk, but these findings surface a new type of risk where employees are granting AI tools persistent, authenticated access to enterprise systems and applications:

- **1 out of 3 employees** connected AI to their work email
- **1 out of 5 employees** connected AI to internal databases
- **1 out of 10 employees** shared API keys and developer credentials to AI
- **Nearly 1 out of 10 employees** connected customer data to AI

The rate at which the connection surface is expanding is growing exponentially. The Model Context Protocol - the open standard that lets AI agents connect directly into the enterprise - has grown from roughly **100 servers** at its launch in November 2024 to over **16,000 today**. This means that employees - in just a few clicks without IT involvement - can connect an AI agent to any number of enterprise applications.

That API key stat is one that also deserves attention. API keys and developer credentials are not data, but identity. An employee who has granted an unsanctioned AI tool access to production API credentials has given that tool the ability to authenticate as the enterprise in any system that credential is valid for.

One survey respondent stated that while they are aware of these risks, they continue to take these actions, “My company deals with a lot of personal banking information and I feel like using outside tools could compromise web safety.”

How employees rank AI risk (but are still doing these things)*

Sending messages or emails

Monitoring info or systems

Updating docs or workflows

Executing tasks

Making automatic decisions

Handling financial data

Making changes in internal systems

Anthropic moves in on enterprise

At the time of creating this report, Anthropic announced two initiatives: providing organizational admins with more Claude-usage controls; and a new plugin marketplace that lets companies build and distribute their own plugins.

The admin controls is a move in the right direction, but they only apply to Claude. If your employees are also using ChatGPT, Gemini, Copilot, or something else, you’re not covered. Even within Claude, these controls do not provide per-request or per-data element

permission, meaning access is granted to the entire plugin or MCP connector and not evaluated at the AI agent action level.

The plugin marketplace is a step toward enterprise customization, but breaks least privilege. These tools, once published, can be downloaded by anyone within the company. There's no way to provide a plugin to a specific team, role, or department.

Another risk is that even though a plug-in is not editable by a user, nothing prevents the user from copying the prompts in the plug-in skills and removing any guardrails they wish to avoid. This may give the admin a false sense of safety.

While these features may change, the fact remains that even when AI vendors release governance, they're only able to govern themselves.

Most organizations are running blind

When employees were asked how well their organization understands AI usage, **34.8%** said "very well" and **36.8%** said "somewhat well." In contrast, these same employees describe working in environments where AI is "allowed but loosely governed," where "there is no consistent policy," and an absence of formal rules rather than active oversight.

Most employees work without AI rules

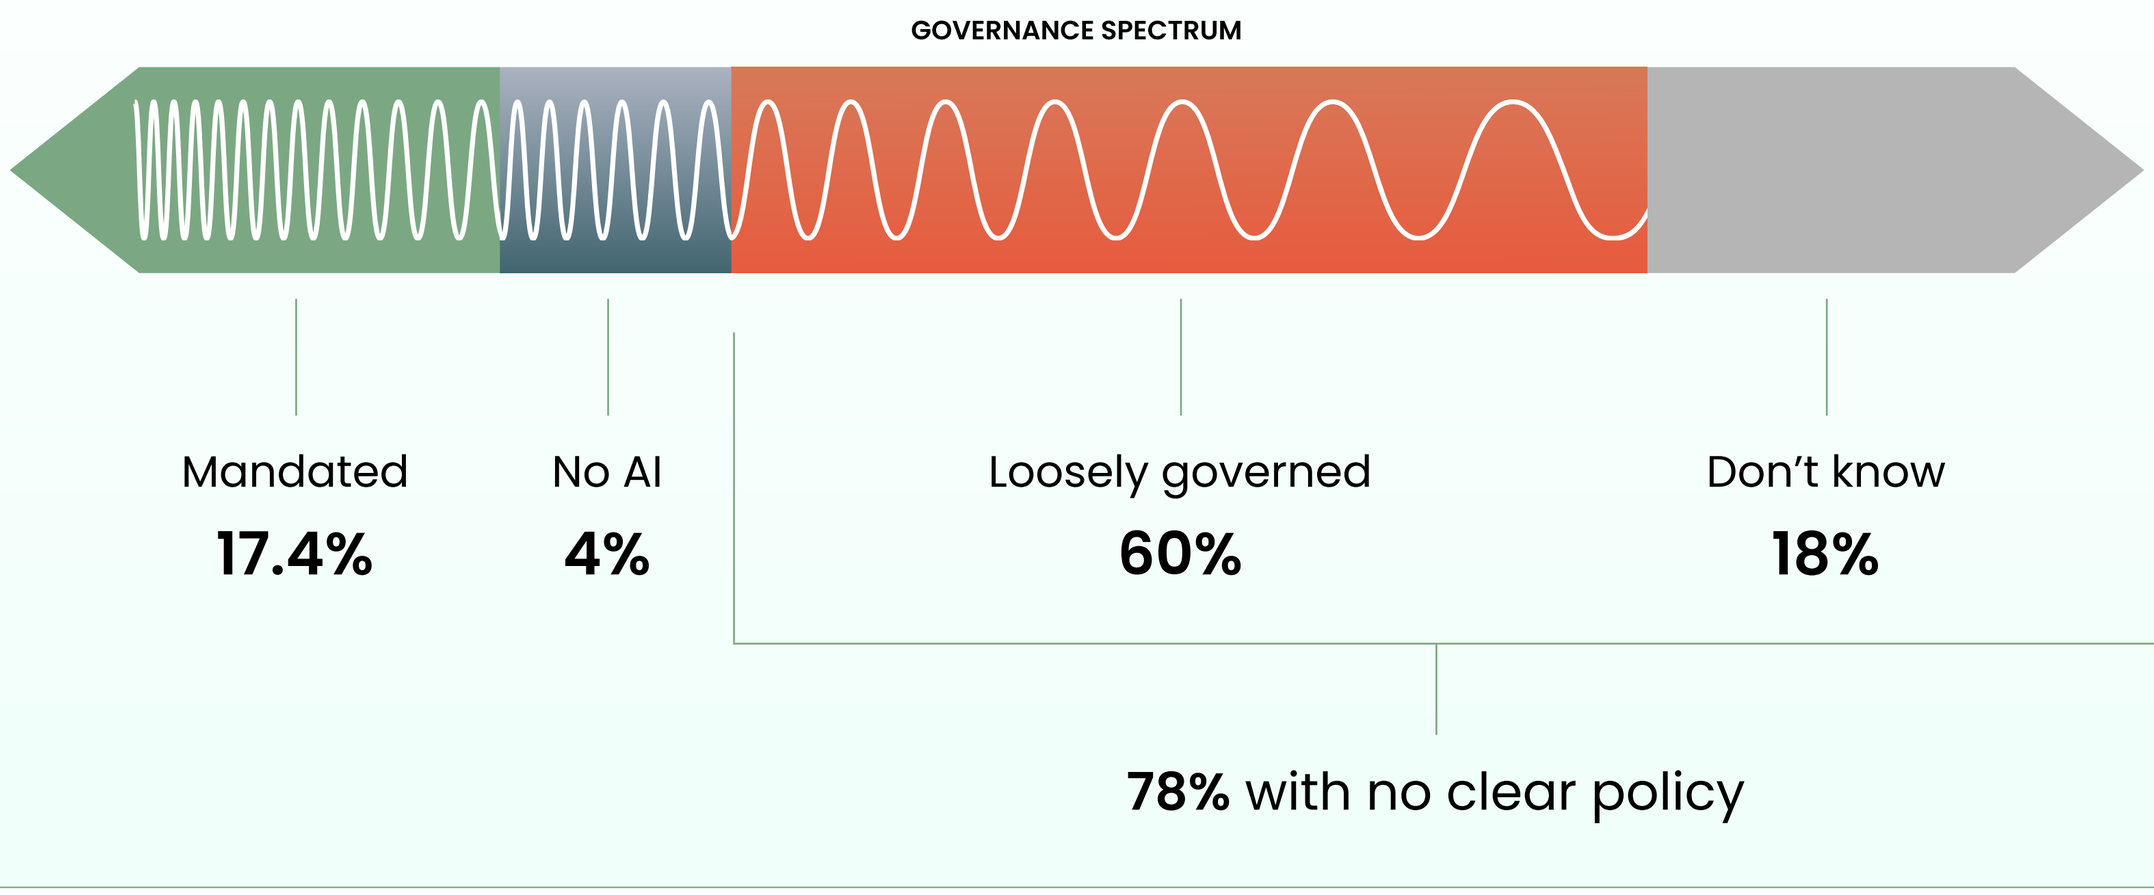

These results show that many employees are unclear or confused by their organization's AI policies. More employees work under loose governance than any formal structure, mandated or not. Many aren't even sure if their companies have AI governance.

Direct respondent feedback:

> “AI isn’t restricted, there are no approved AI tools, it’s every man for himself.”

> “Because we are just out there blind, taking a chance and hoping we don’t get into trouble as there has been no real guidance from our department.”

> “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.”

Employees are on their own with AI

When asked how their organization's AI approval and restriction policies make them feel, 33% said they feel very confident and supported. Another quarter feel somewhat confident, but with reservations. The remaining 42% are either neutral, mixed, hesitant, or frustrated.

The qualitative responses offer some insights. Employees don't describe feeling actively blocked or punished, they describe working in a vacuum:

- *“AI isn't restricted, there are no approved AI tools, it's every man for himself"*
- *"Because we are just out there blind, taking a chance and hoping we don't get into trouble as there has been no real guidance from our department"*
- *"There is no consistent policy on AI usage and I'm not sure what to do"*
- *"We have access but no formal policy so usage is all over the place"*

Thinking about how AI is approved or restricted at your organization, how does that make you feel about using AI at work?

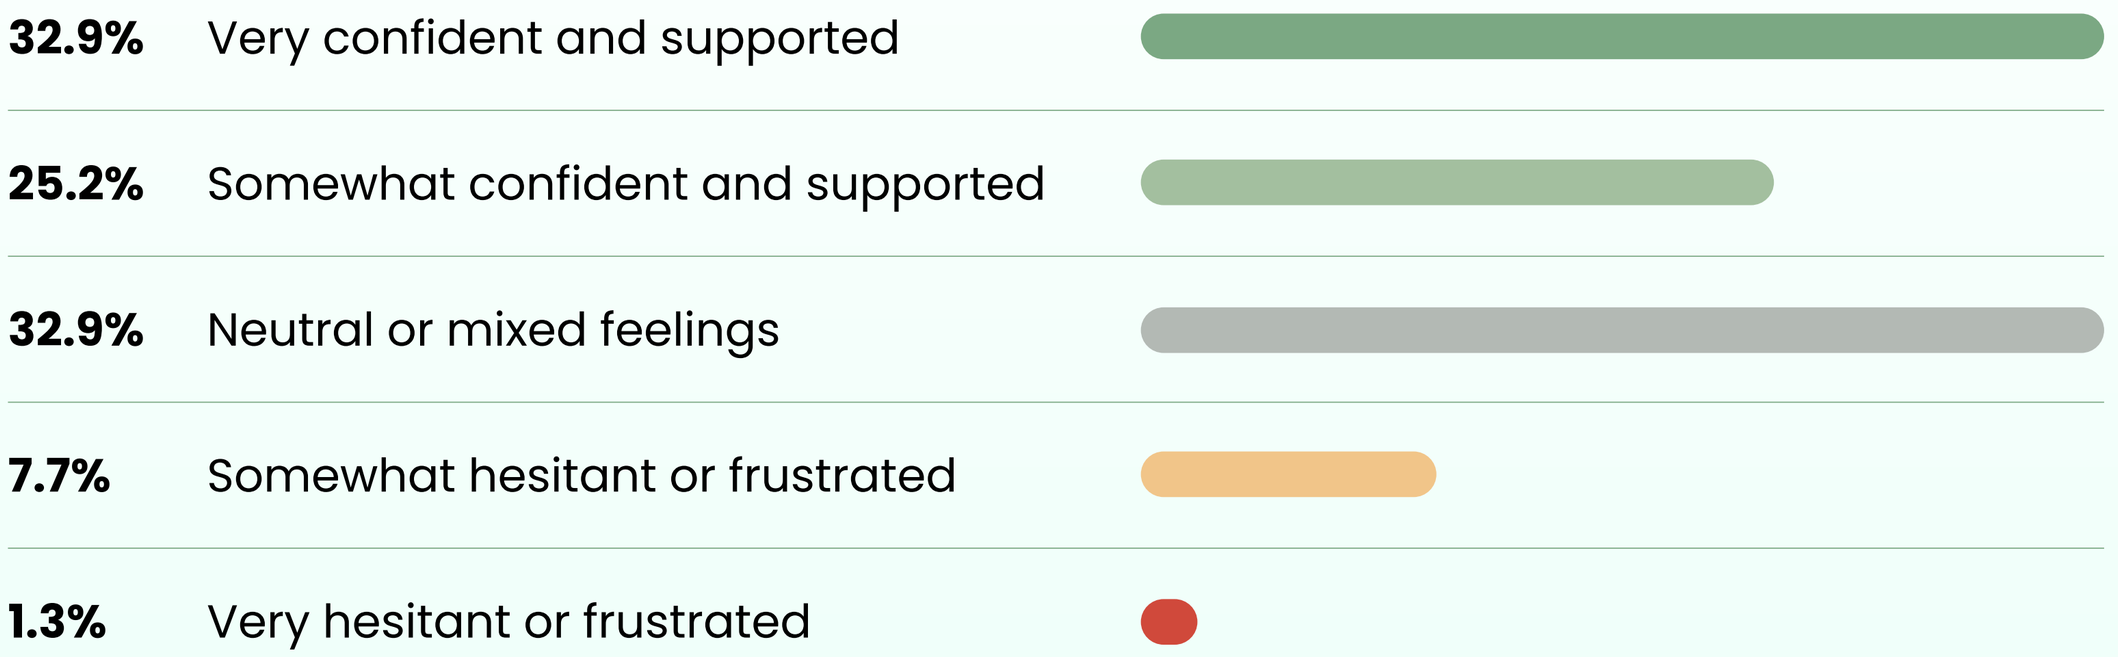

Barndoor.ai

Employees are moving from generative to agentic

Most enterprise AI governance frameworks were designed for generative AI: tools that respond to questions, requests, and prompts via an AI chat. The employee asks, the AI answers, the employee decides what to do with the output. The research supports this understanding, with nearly all respondents saying they use AI for generating or editing content and research.

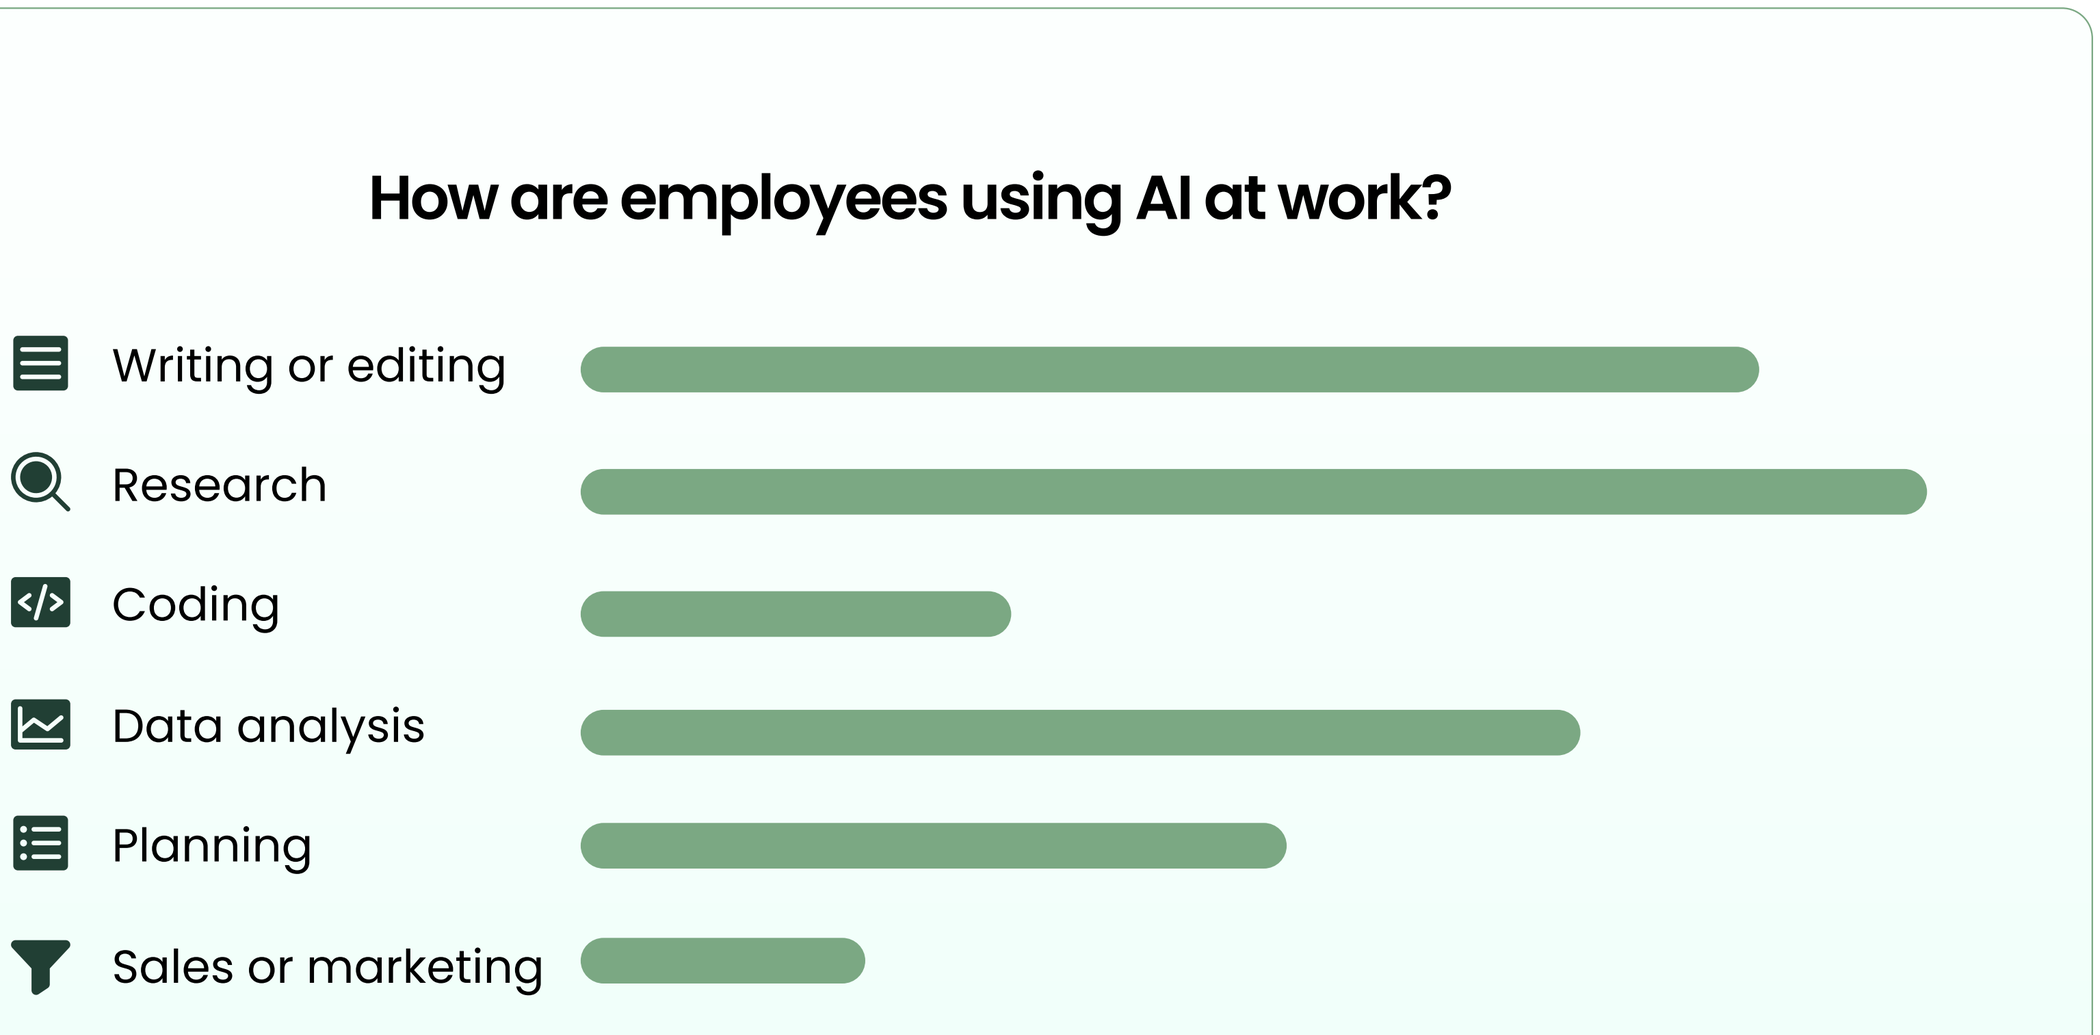

That said, a shift is occurring. The research found that a share of the workforce has moved beyond merely chatting with an AI. When asked whether their AI tools "go beyond giving suggestions and actually take actions," **28.3%** said yes.

Over **30%** report using AI tools that take actions on their behalf, not just generating suggestions but executing tasks autonomously.

Agentic AI means your AI has access to tools and applications that give them the ability to read and write emails, edit your calendar, change your CRM data, execute code, call APIs, and manage workflows. Their actions are concrete - an AI agent that sends an incorrect email, posts the wrong social media image, or modifies a database record has made a change that's harder to reverse. One respondent said, ***"AI that takes actions or runs workflows can be dangerous as it will become privy to private information. It may start thinking for itself and adjust to the workflow, so I don't trust it."***

Barndoor.ai

A third of employees are already using AI that acts

AI that answers

53% of employees

AI that acts

30% of employees

1 in 10 employees can’t identify the tools they’re using

Let’s reframe: Access isn’t the same as governance

Most organizations tackle the shadow AI issue through better policies, employee training, or stronger access controls.

This survey reveals why this isn’t enough. It’s not about the AI tool employees choose, it’s about what these tools are given access to once that choice was made. Remember our previous stats: one in three employees has granted AI persistent access to their work email. One in five has connected AI to internal databases and systems. One in ten has handed over API keys.

Traditional identity access management prevents unauthorized access, but in this case, you’ve already granted access to AI that’s taking action across your systems. Your focus should now be what is AI doing right now and managing those actions. It’s not just about preventing who gets in, but what happens inside.

This means visibility into runtime AI actions, granular controls across permissions and context, and the ability to enforce policy at the individual agent level.

Survey methodology

The first annual Organizational Behavior & AI Governance report was commissioned by Barndoor in February-March 2026 across 155 professionals with titles in executive leadership, operations, finance, HR, legal, IT/security, engineering, data analytics, marketing, customer support, and sales.

The survey was fielded by Remesh, a hybrid qualitative-quantitative research platform that combines polling and open-ended responses to gather both statistical findings and verbal employee responses.

About Barndoor

Founded in 2024, Barndoor AI is the AI governance platform built for the agentic enterprise. Barndoor's centralized control plane gives IT and security teams the visibility and fine-grained access control they need to govern AI agents across every tool, team, and platform.

Venn.ai provides knowledge workers secure, governed access to agentic AI across the workplace apps they already use. Together, Barndoor and Venn address the AI governance gap from both the enterprise end, as well as individual empowerment. Learn more at Barndoor.ai and venn.ai.

Barndoor.ai

Venn.ai
POWERED BY Barndoor.ai