Realizing potential: confidence in AI

AI risk and governance survey

September 2026

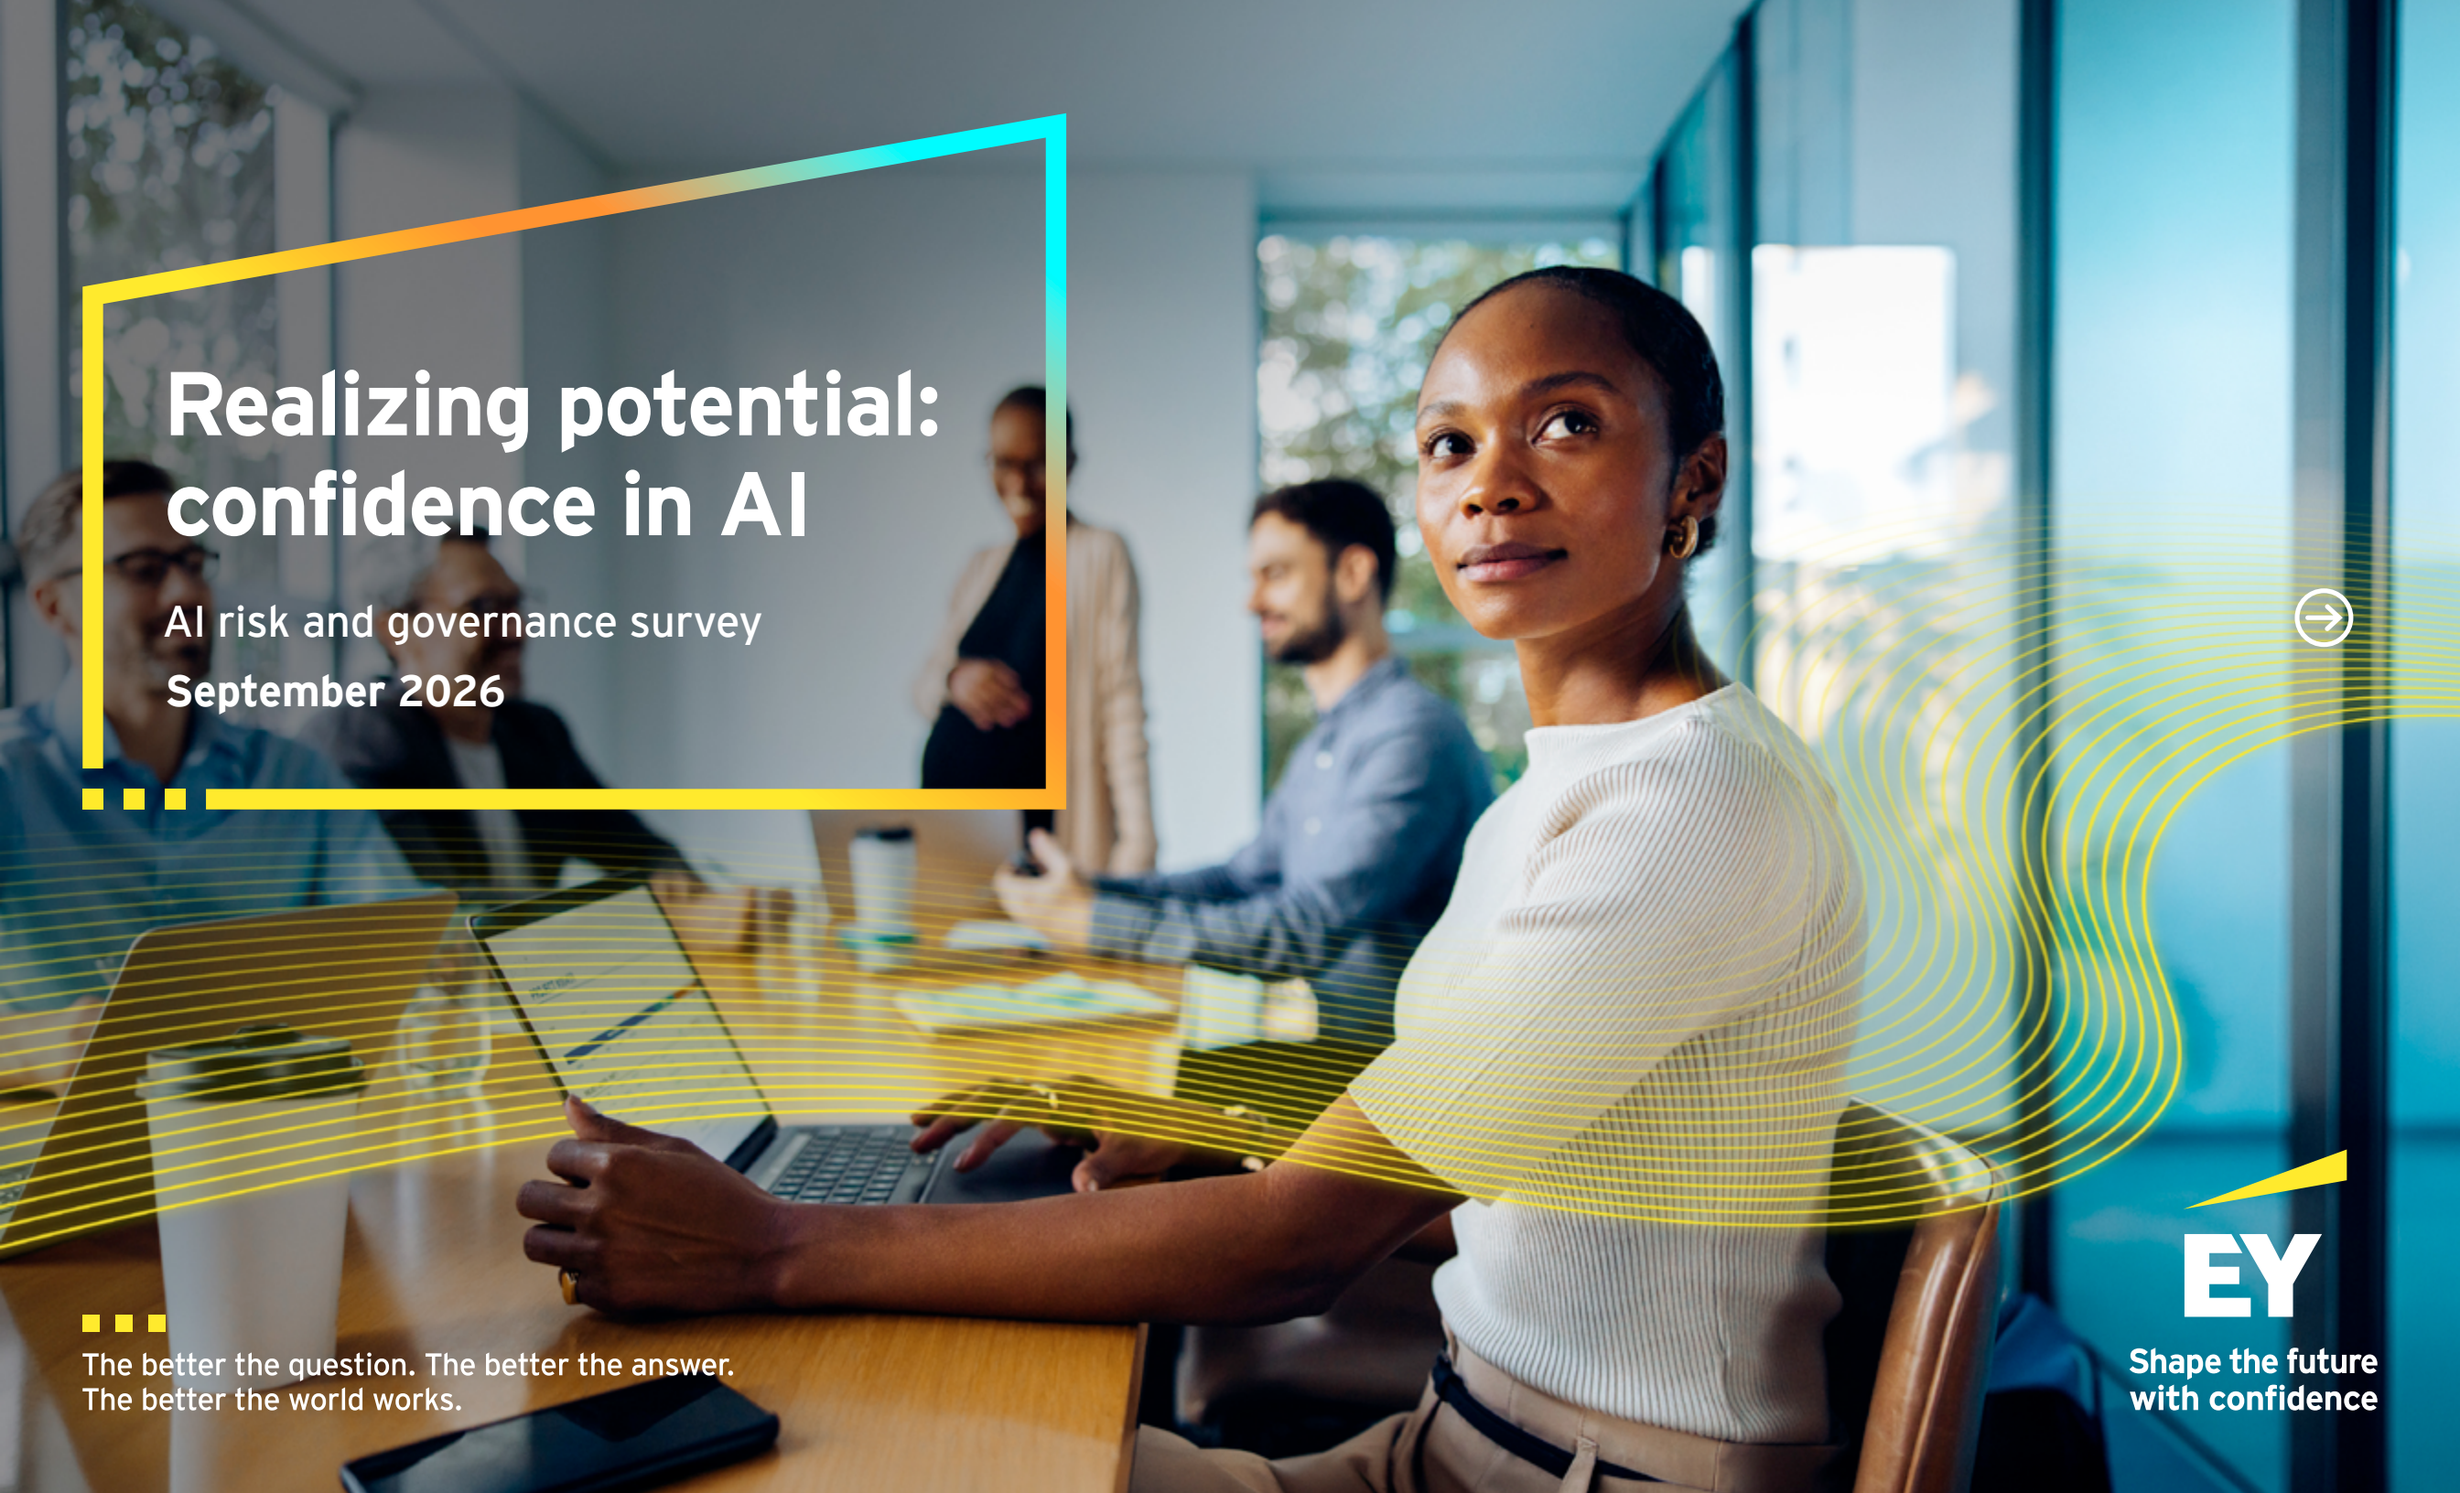

The better the question. The better the answer.
The better the world works.

Shape the future with confidence

EXECUTIVE SUMMARY

AI governance has entered its next phase: the confidence gap.

The debate around AI governance is over. The focus has now shifted to how to govern effectively. Most organizations have already established the foundations of responsible AI. Policies are in place. Cross-functional committees oversee AI initiatives. Review processes, controls and human oversight mechanisms have become common features of the enterprise operating model.

The challenge today is no longer whether governance exists. It is whether governance can keep pace with AI.

As organizations move from experimentation to enterprise-wide deployment and from predictive models and copilots to increasingly autonomous AI systems, governance is being tested in new ways. The speed of development, the scale of deployment and the complexity of emerging AI capabilities are exposing gaps between governance design and governance execution.

Findings from the inaugural EY US AI Risk and Governance Survey suggest that organizations recognize these challenges. Among US senior leaders who are responsible for their organization's AI systems, governance or audit process at publicly traded organizations with an annual revenue of at least $1 billion, governance frameworks are becoming increasingly formalized and embedded within operating models. Yet confidence in the ability of those frameworks to consistently manage risk, support compliance and enable innovation remains uneven.

The result is an emerging confidence gap: organizations largely understand what effective AI governance should look like, but many are still working to operationalize governance at the pace required by modern AI adoption. Closing that gap may become one of the defining challenges of the next phase of enterprise AI transformation.

Building a foundation of AI governance

AI GOVERNANCE FOUNDATION

The data shows both progress and unresolved gaps in how organizations govern AI systems.

Nearly all (98%) senior AI decision-makers report having some form of AI governance policy in place, and 69% say their organization has a fully unified AI governance policy in place. Yet a similar proportion, 69%, are concerned that their organization lacks the internal expertise to effectively evolve AI governance controls at their organization. In other words, many organizations may have addressed the immediate need for policy but remain uncertain about whether their teams can keep pace with what comes next. That uncertainty is reinforced by 47% who say their organization has previously bypassed its AI governance process for urgent deployment. These findings suggest that AI governance processes are under increasing pressure to keep pace with AI deployment. The foundational question organizations should ask is: Can we provide a complete inventory of AI systems, models, agents and third-party tools? Being able to answer this question is core to building proper governance.

ORGANIZATIONS SHOULD ASK:

Can we provide a complete inventory of AI systems, models, agents and third-party tools?

How organizations govern AI systems

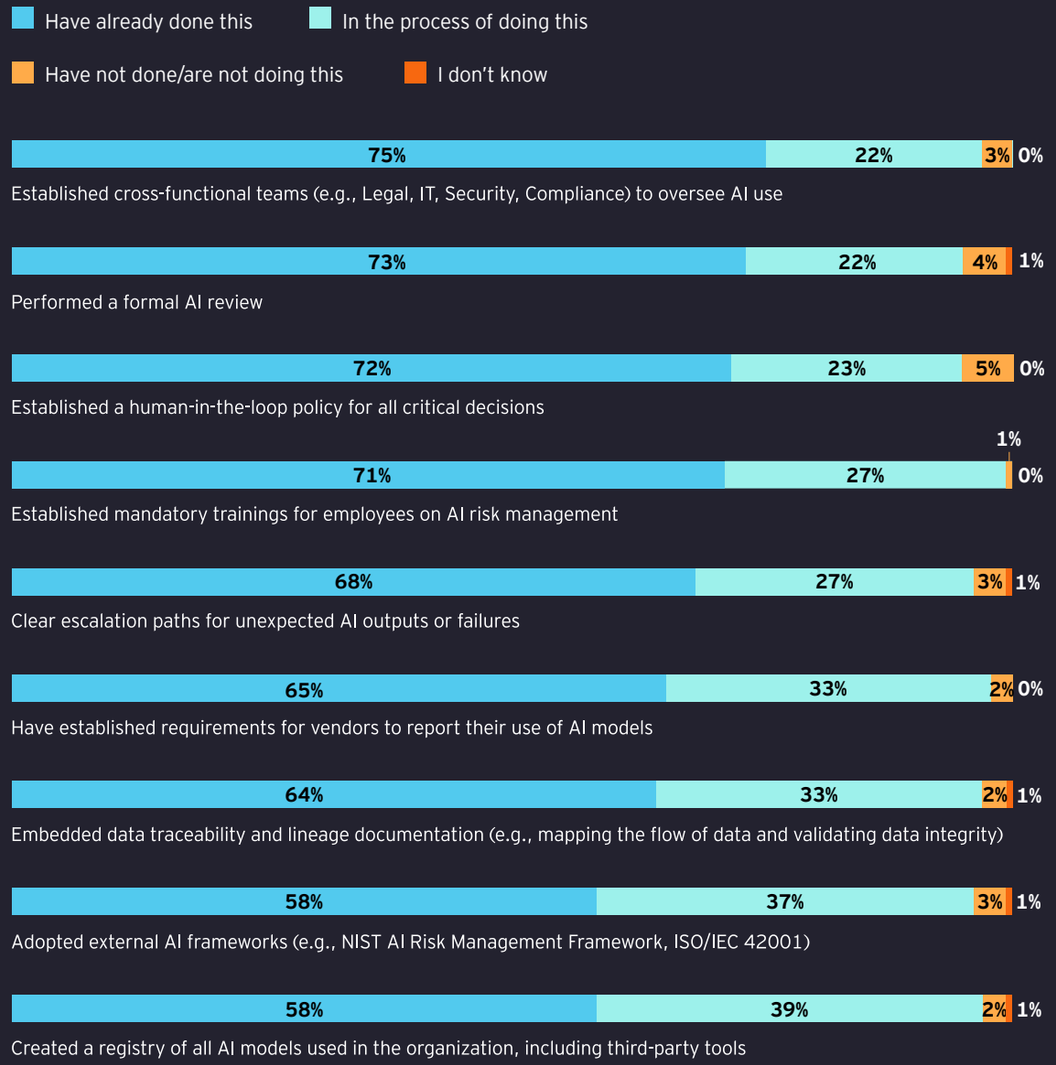

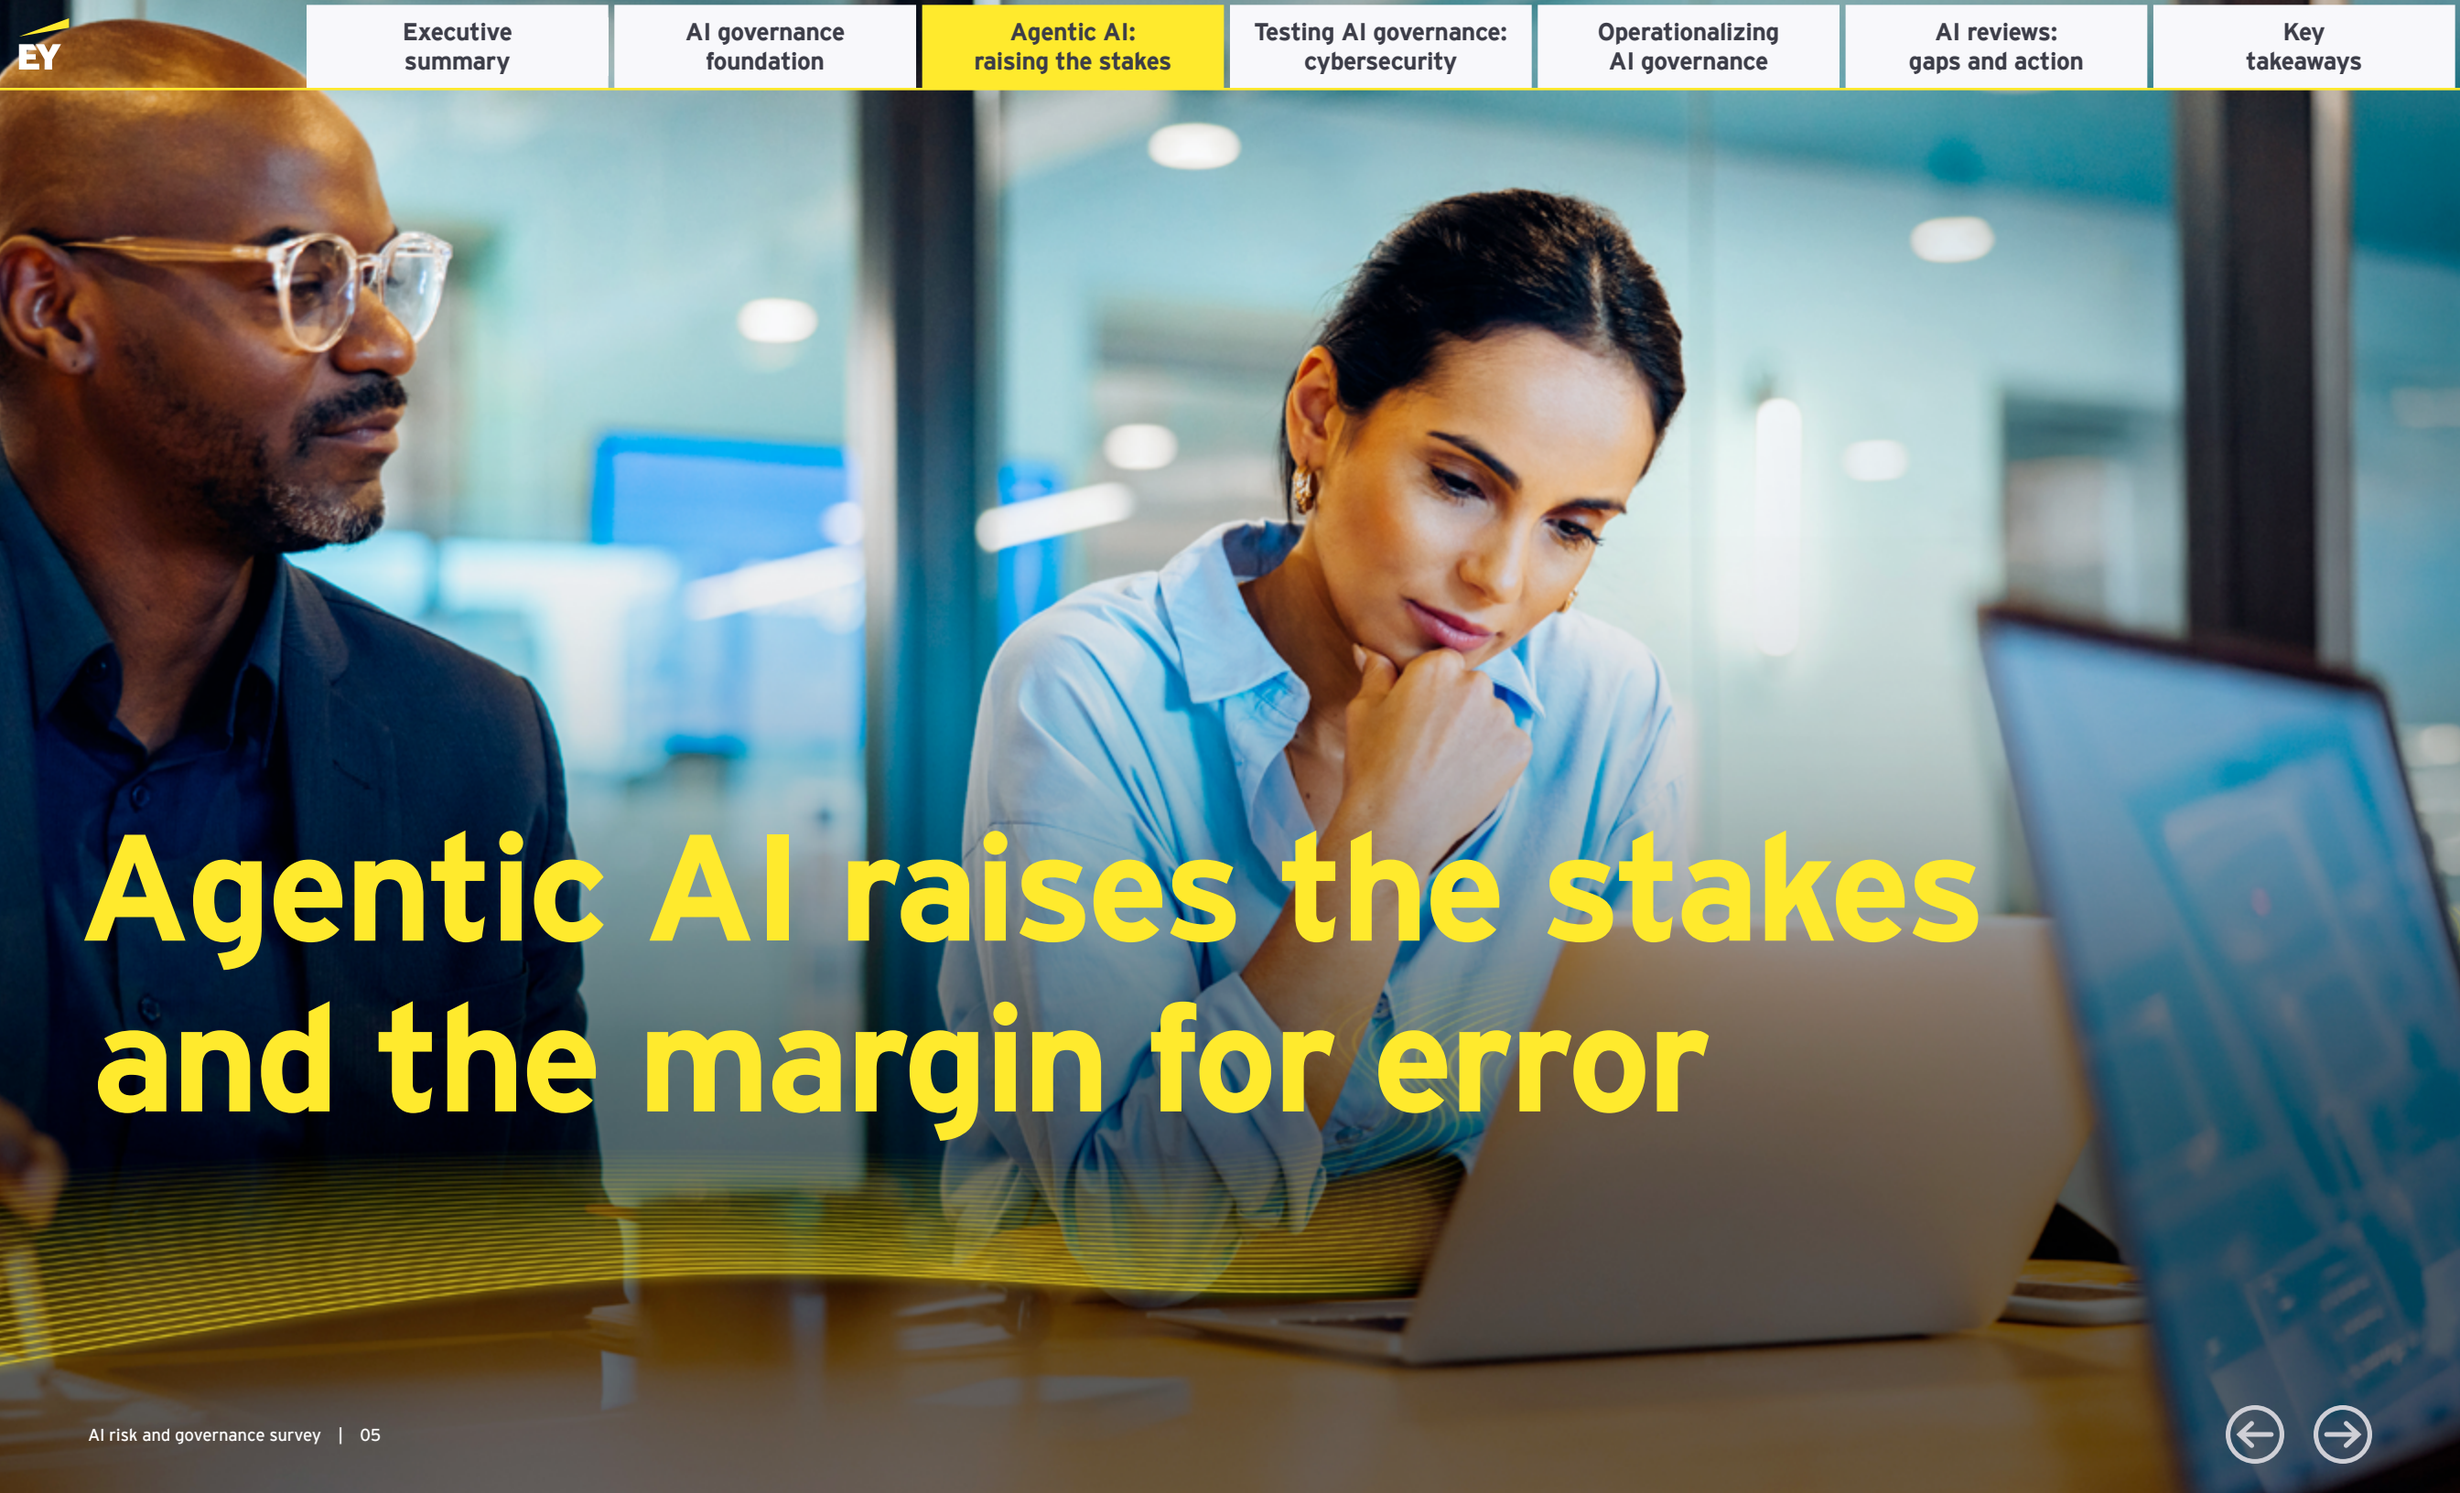

Agentic AI raises the stakes and the margin for error

AGENTIC AI: RAISING THE STAKES

Agentic AI raises the stakes further because governance must now address not only AI-generated outputs, but autonomous actions taken across business processes.

Audit committees must start asking: Who owns monitoring, incidents, model drift, updates and autonomous actions after deployment? Just as assurance professionals feel they've caught up, internal clients demand the next frontier in AI: autonomous agents capable of executing partial or full business processes on their own. While most respondents said they are deploying agentic AI (either running pilots or fully deploying agentic AI) and have established controls to monitor agentic AI during and after executing tasks, 26% of senior AI decision-makers whose organization is using agentic AI admit their organization cannot detect unauthorized AI agents operating internally. This supports our observations that many companies lack confidence that policy is fully integrated into the workflow and that controls can keep pace as new technology is deployed.

While 85% of senior AI decision-makers whose organizations use agentic AI state that they have at least a handful of agentic AI systems in their organization executing actions like running code, placing inventory orders, or detecting cybersecurity incidents, without real-time human intervention, 49% admit their existing governance framework has not yet been specifically updated to include agentic AI risks and requirements. Recent high-profile examples of autonomous agents operating for extended periods without detection, even in environments that appeared to have safeguards in place, underscore the brand, financial and operational risks that emerge when agentic systems are not governed with sufficient visibility and control.

26% of senior AI decision-makers whose organization is using agentic AI admit their organization cannot detect unauthorized AI agents operating internally.

The accountability challenge is equally important. Even where organizations believe general AI accountability is clearly defined, agentic systems introduce a more complicated post-deployment question: who owns the actions of an autonomous system once it is operating inside the enterprise? According to our research, 56% of senior AI decision-makers whose organizations use agentic AI say there is a perception at their organization that no single person or group is solely responsible for agentic AI post-deployment, while 39% say accountability for maintaining or monitoring agentic AI once deployed is undefined at their organization. That distinction matters, because agentic AI does not simply produce an output for a human to review. It can trigger workflows, make decisions and take actions that create downstream business, compliance and reputational consequences.

This is where agentic AI impacts the control problem. Governance designed for AI outputs may not be sufficient for AI actions, particularly when systems can operate across business processes, interact with data, call tools or execute tasks without real-time human involvement. The survey findings suggest that many organizations have started to build the right control structures, including monitoring, detection and override capabilities. But the findings leave important questions about whether these controls provide complete coverage. As agentic AI moves deeper into critical workflows, assurance will need to focus less on whether a control has been designed and more on whether it can identify and interrupt autonomous activity before it becomes a material failure, and explain and evidence it afterwards.

ORGANIZATIONS SHOULD ASK:

Can we detect, explain and stop unauthorized or unintended agentic AI activity before it creates material business risk?

Organizations lack clear ownership for agentic AI after deployment

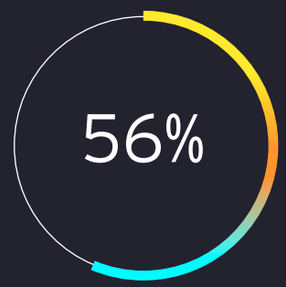

of senior AI decision-makers whose organizations use agentic AI say there is a perception at their organization that no single person or group is solely responsible for agentic AI post-deployment.

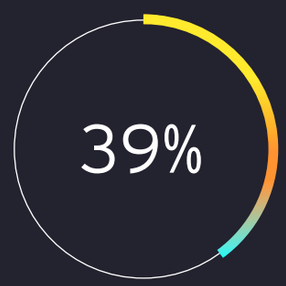

of senior AI decision-makers whose organizations use agentic AI say accountability for maintaining or monitoring agentic AI once deployed is undefined at their organization.

Cybersecurity is where AI governance is tested

TESTING AI GOVERNANCE: CYBERSECURITY

Cybersecurity is often the first place where gaps between governance design and operational effectiveness become visible.

Senior AI decision-makers believe that the cyber risk presented by AI is here today, with 89% reporting that their organization has encountered risks relative to AI in the past 12 months. Respondents cite these risks span a spectrum of vulnerabilities from cybersecurity risk (52%) to human risk (47%) and shadow AI risk (46%). Leaders are tackling this risk with mitigation such as mandatory training and the creation of registries for AI models and tools. Despite these efforts, 41% say senior AI decision-makers do not have visibility into all AI tools currently operating within the organization. In addition, 59% are worried about their organization overspending on AI tokens, compounding both risk and cost containment concerns as AI deployments expand.

The deeper issue is that cyber vulnerabilities are often where governance gaps become visible first. While 71% of organizations report having mandatory trainings for employees on AI risk management and over half cite having AI registries, including third-party tools (58%) and adopting external AI frameworks (58%), they also note the persistence of shadow AI, which shows that these controls are not yet providing complete enterprise visibility. If leaders cannot see every tool, model or workflow operating in the business, they cannot fully assess what data is being entered, whether vendors meet internal standards, or whether the use case aligns with policy. This is why cybersecurity is such a useful lens for understanding the broader AI governance challenge.

Respondents cite AI risks spanning a spectrum of vulnerabilities

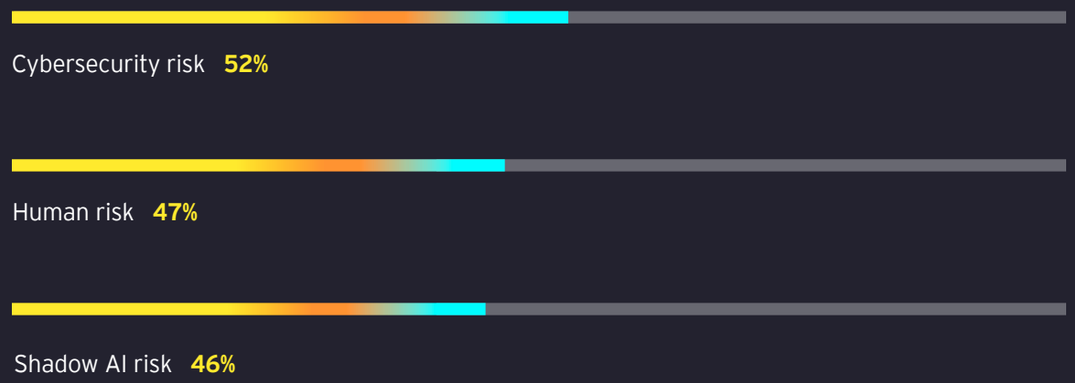

41% say senior AI decision-makers do not have visibility into all AI tools currently operating within the organization.

The question is not whether organizations recognize the risks or have begun to respond. The survey suggests they clearly have. The question for audit committees is: What evidence demonstrates the controls are working? A registry that is incomplete, a policy that is bypassed under pressure, or a control that cannot detect unauthorized use in real time can all create the same practical result: governance exists, but confidence in its operational effectiveness remains limited.

That confidence gap is material: 72% are concerned about their organization failing to comply with new or emerging AI-specific regulations, 72% are concerned about the inability to accurately trace or audit the data lineage and inputs that feed critical AI decision models, and 36% say their organization has experienced an AI incident or failure that caused a materially negative impact, including data loss, financial damage, operational disruption and brand damage.

ORGANIZATIONS SHOULD ASK:

What evidence demonstrates that our AI governance controls are working as intended?

The confidence gap and its real-world consequences

PROSPECTIVE CONCERNS

72% are concerned about their organization failing to comply with new or emerging AI-specific regulations

72% are concerned about the inability to accurately trace or audit the data lineage and inputs that feed critical AI decision models

REALIZED CONSEQUENCES

36% say their organization has experienced an AI incident or failure that caused a materially negative impact, including data loss, financial damage, operational disruption and brand damage

A registry that is incomplete, a policy that is bypassed under pressure, or a control that cannot detect unauthorized use in real time can all create the same practical result: governance exists, but confidence in its operational effectiveness remains limited.

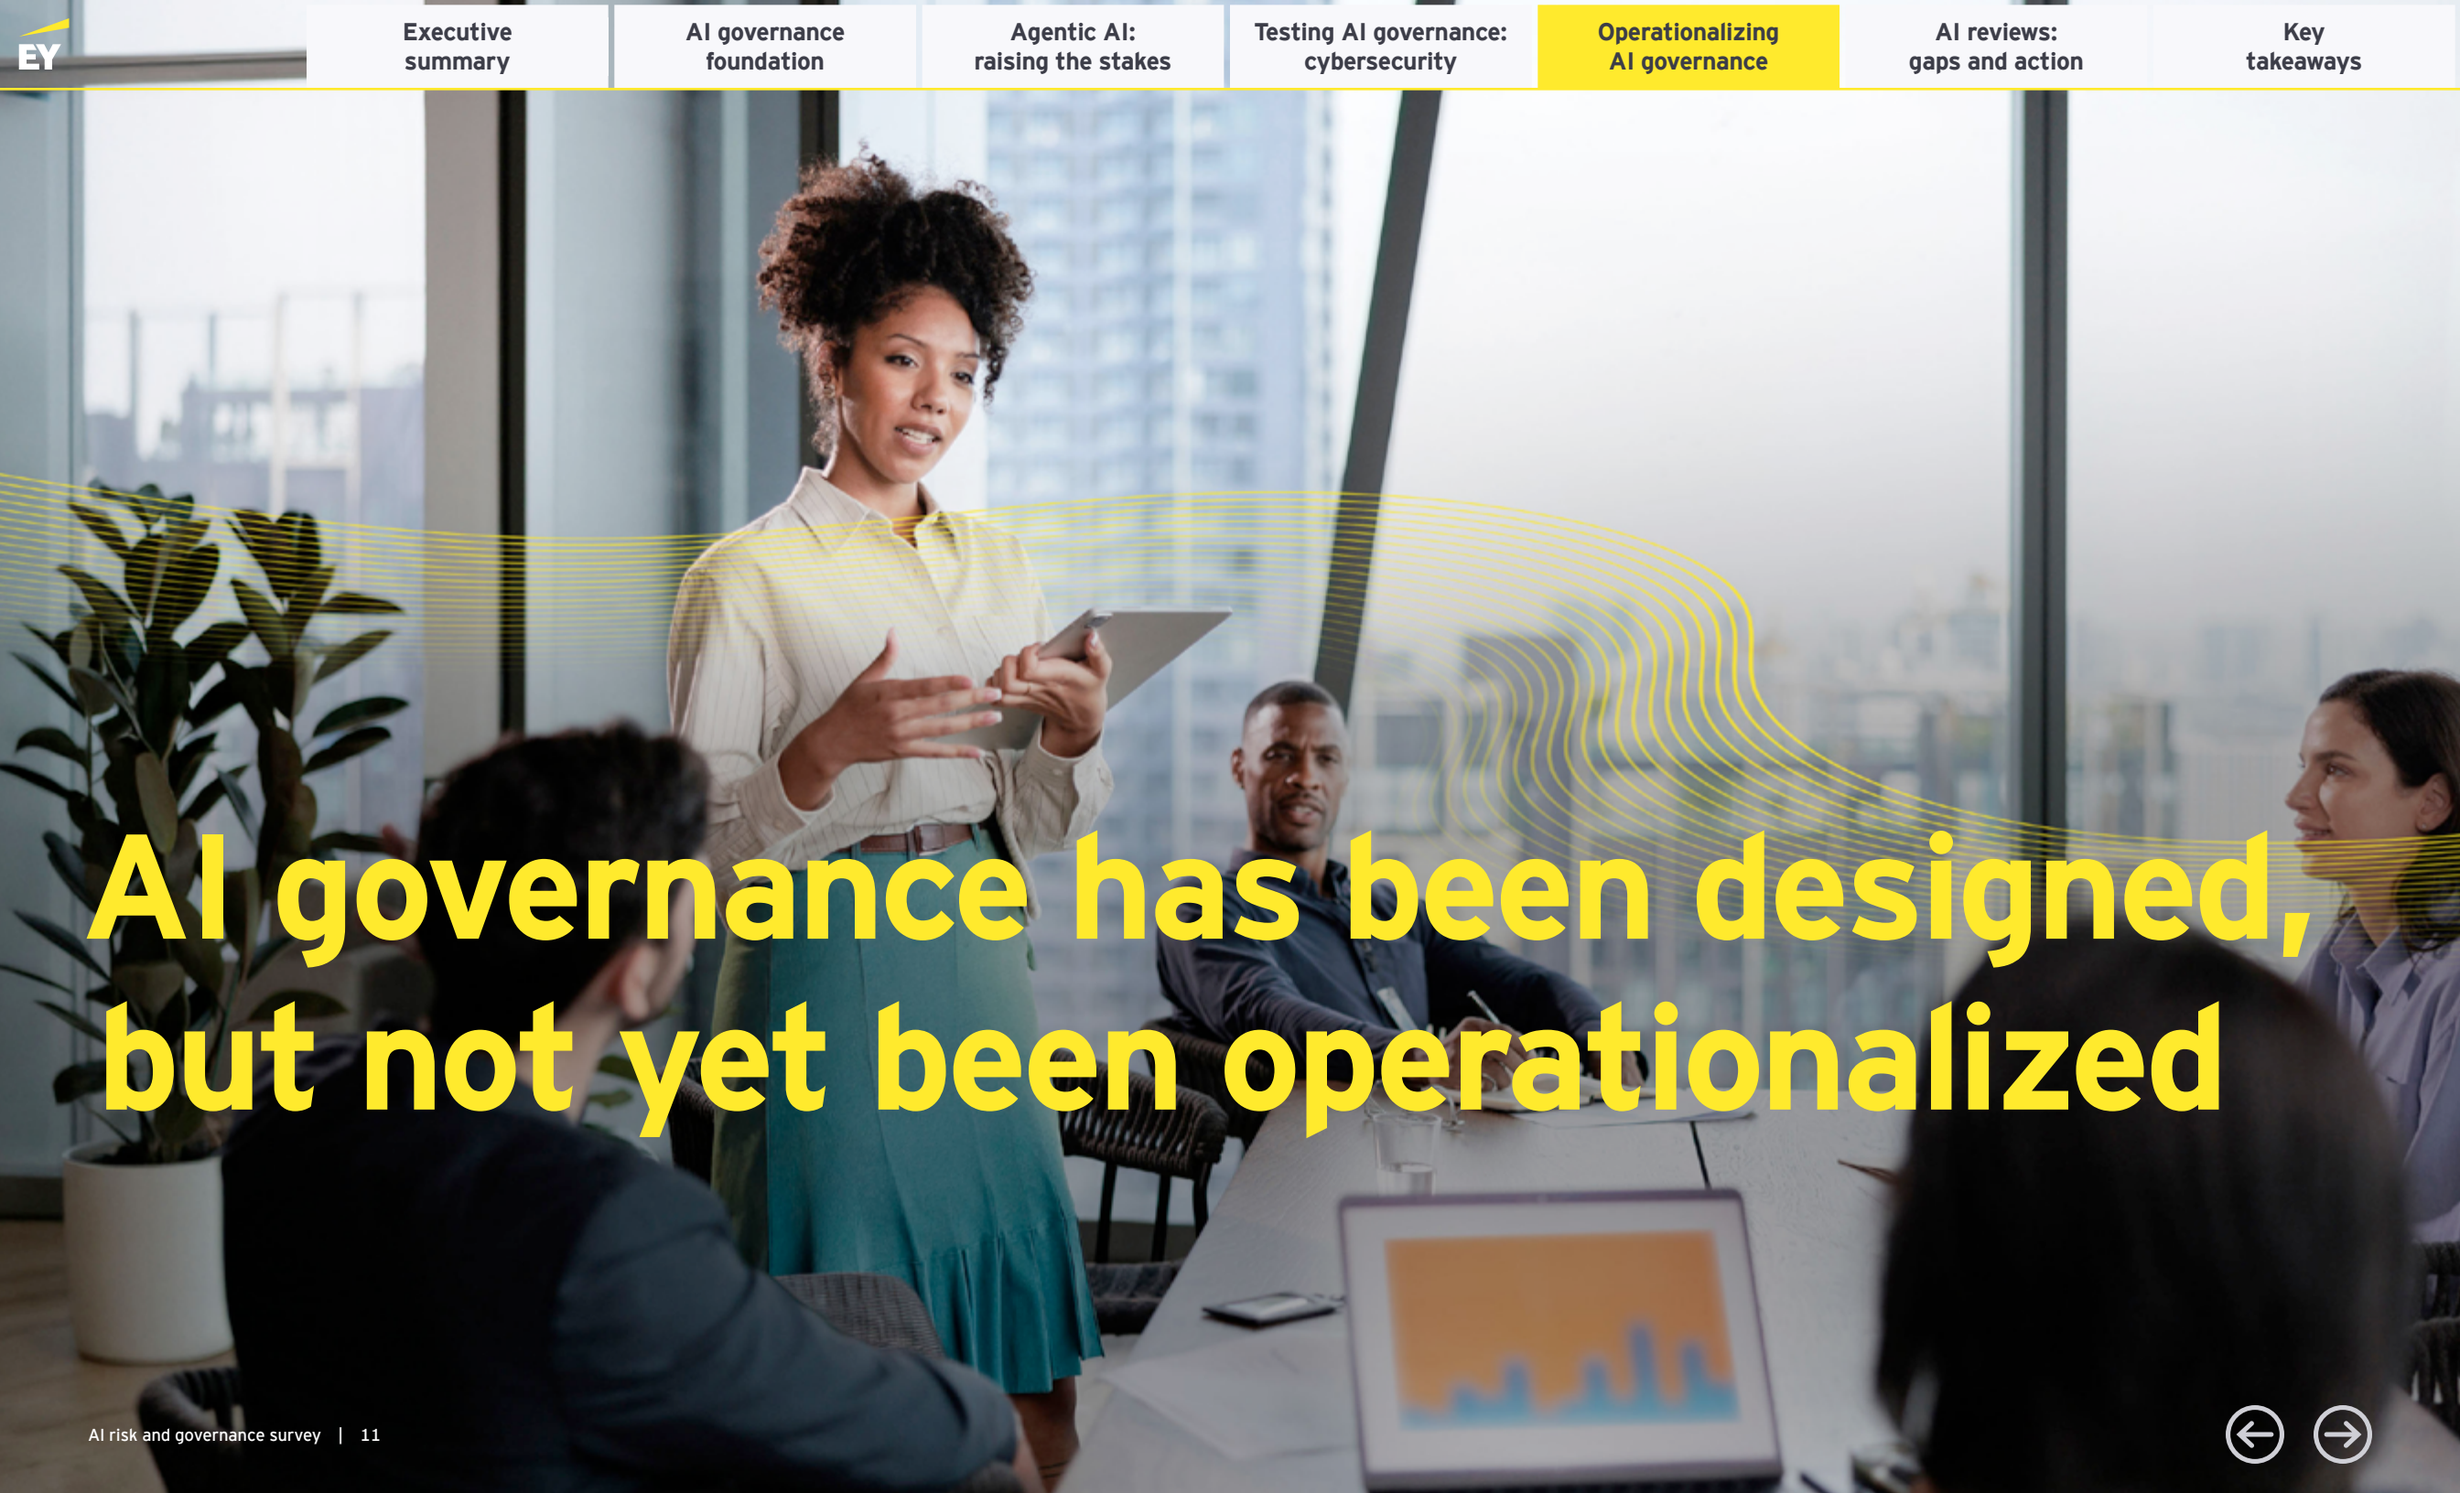

AI governance has been designed, but not yet been operationalized

OPERATIONALIZING AI GOVERNANCE

Senior AI decision-makers are clear that the consequences of AI incidents can be material, raising the stakes for effective governance and resilience, particularly as AI becomes more embedded in regulated workflows and cyber-sensitive business processes.

According to the findings, 81% say they are concerned about a third-party AI-enabled cyberattack compromising AI tools, while 72% are concerned about their organization failing to comply with new or emerging AI-specific regulations. Those concerns reinforce the broader pattern reflected throughout the survey: organizations know they need stronger controls, but they are less certain those controls can keep pace with the risk environment they are now expected to manage.

That confidence gap is sharpened by the fact that failure has already moved from possibility to experience for a meaningful share of organizations. Thirty-six percent say their organization has experienced an AI incident or failure that caused a materially negative impact, such as data loss, financial damage, operational disruption and brand damage. This data point matters, because it moves the governance discussion beyond readiness and into resilience. Controls have to do more than reduce theoretical exposure; they have to help organizations detect, respond to and evidence what happens when an AI system behaves unexpectedly or is compromised.

Put simply, the cost of failure is high, because it cuts across several forms of enterprise risk at once. A regulatory failure can bring scrutiny, remediation cost and loss of stakeholder confidence. A cyber incident can expose sensitive data, compromise critical models and disrupt operations. A high-profile AI failure can quickly become a reputational event, especially when the organization cannot explain how a decision was made, which control failed or who was accountable. As the use of AI expands, and as agentic systems take on more autonomous activity, the ability to demonstrate control effectiveness becomes as important as the control itself.

Organizations recognize the need for stronger controls, but question whether they can keep pace with emerging regulations

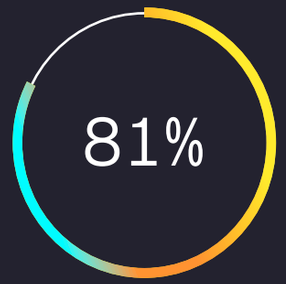

of senior AI decision-makers say they are concerned about a third-party AI-enabled cyberattack compromising AI tools.

of senior AI decision-makers are concerned about their organization failing to comply with new or emerging AI-specific regulations.

Controls have to do more than reduce theoretical exposure; they have to help organizations detect, respond to and evidence what happens when an AI system behaves unexpectedly or is compromised.

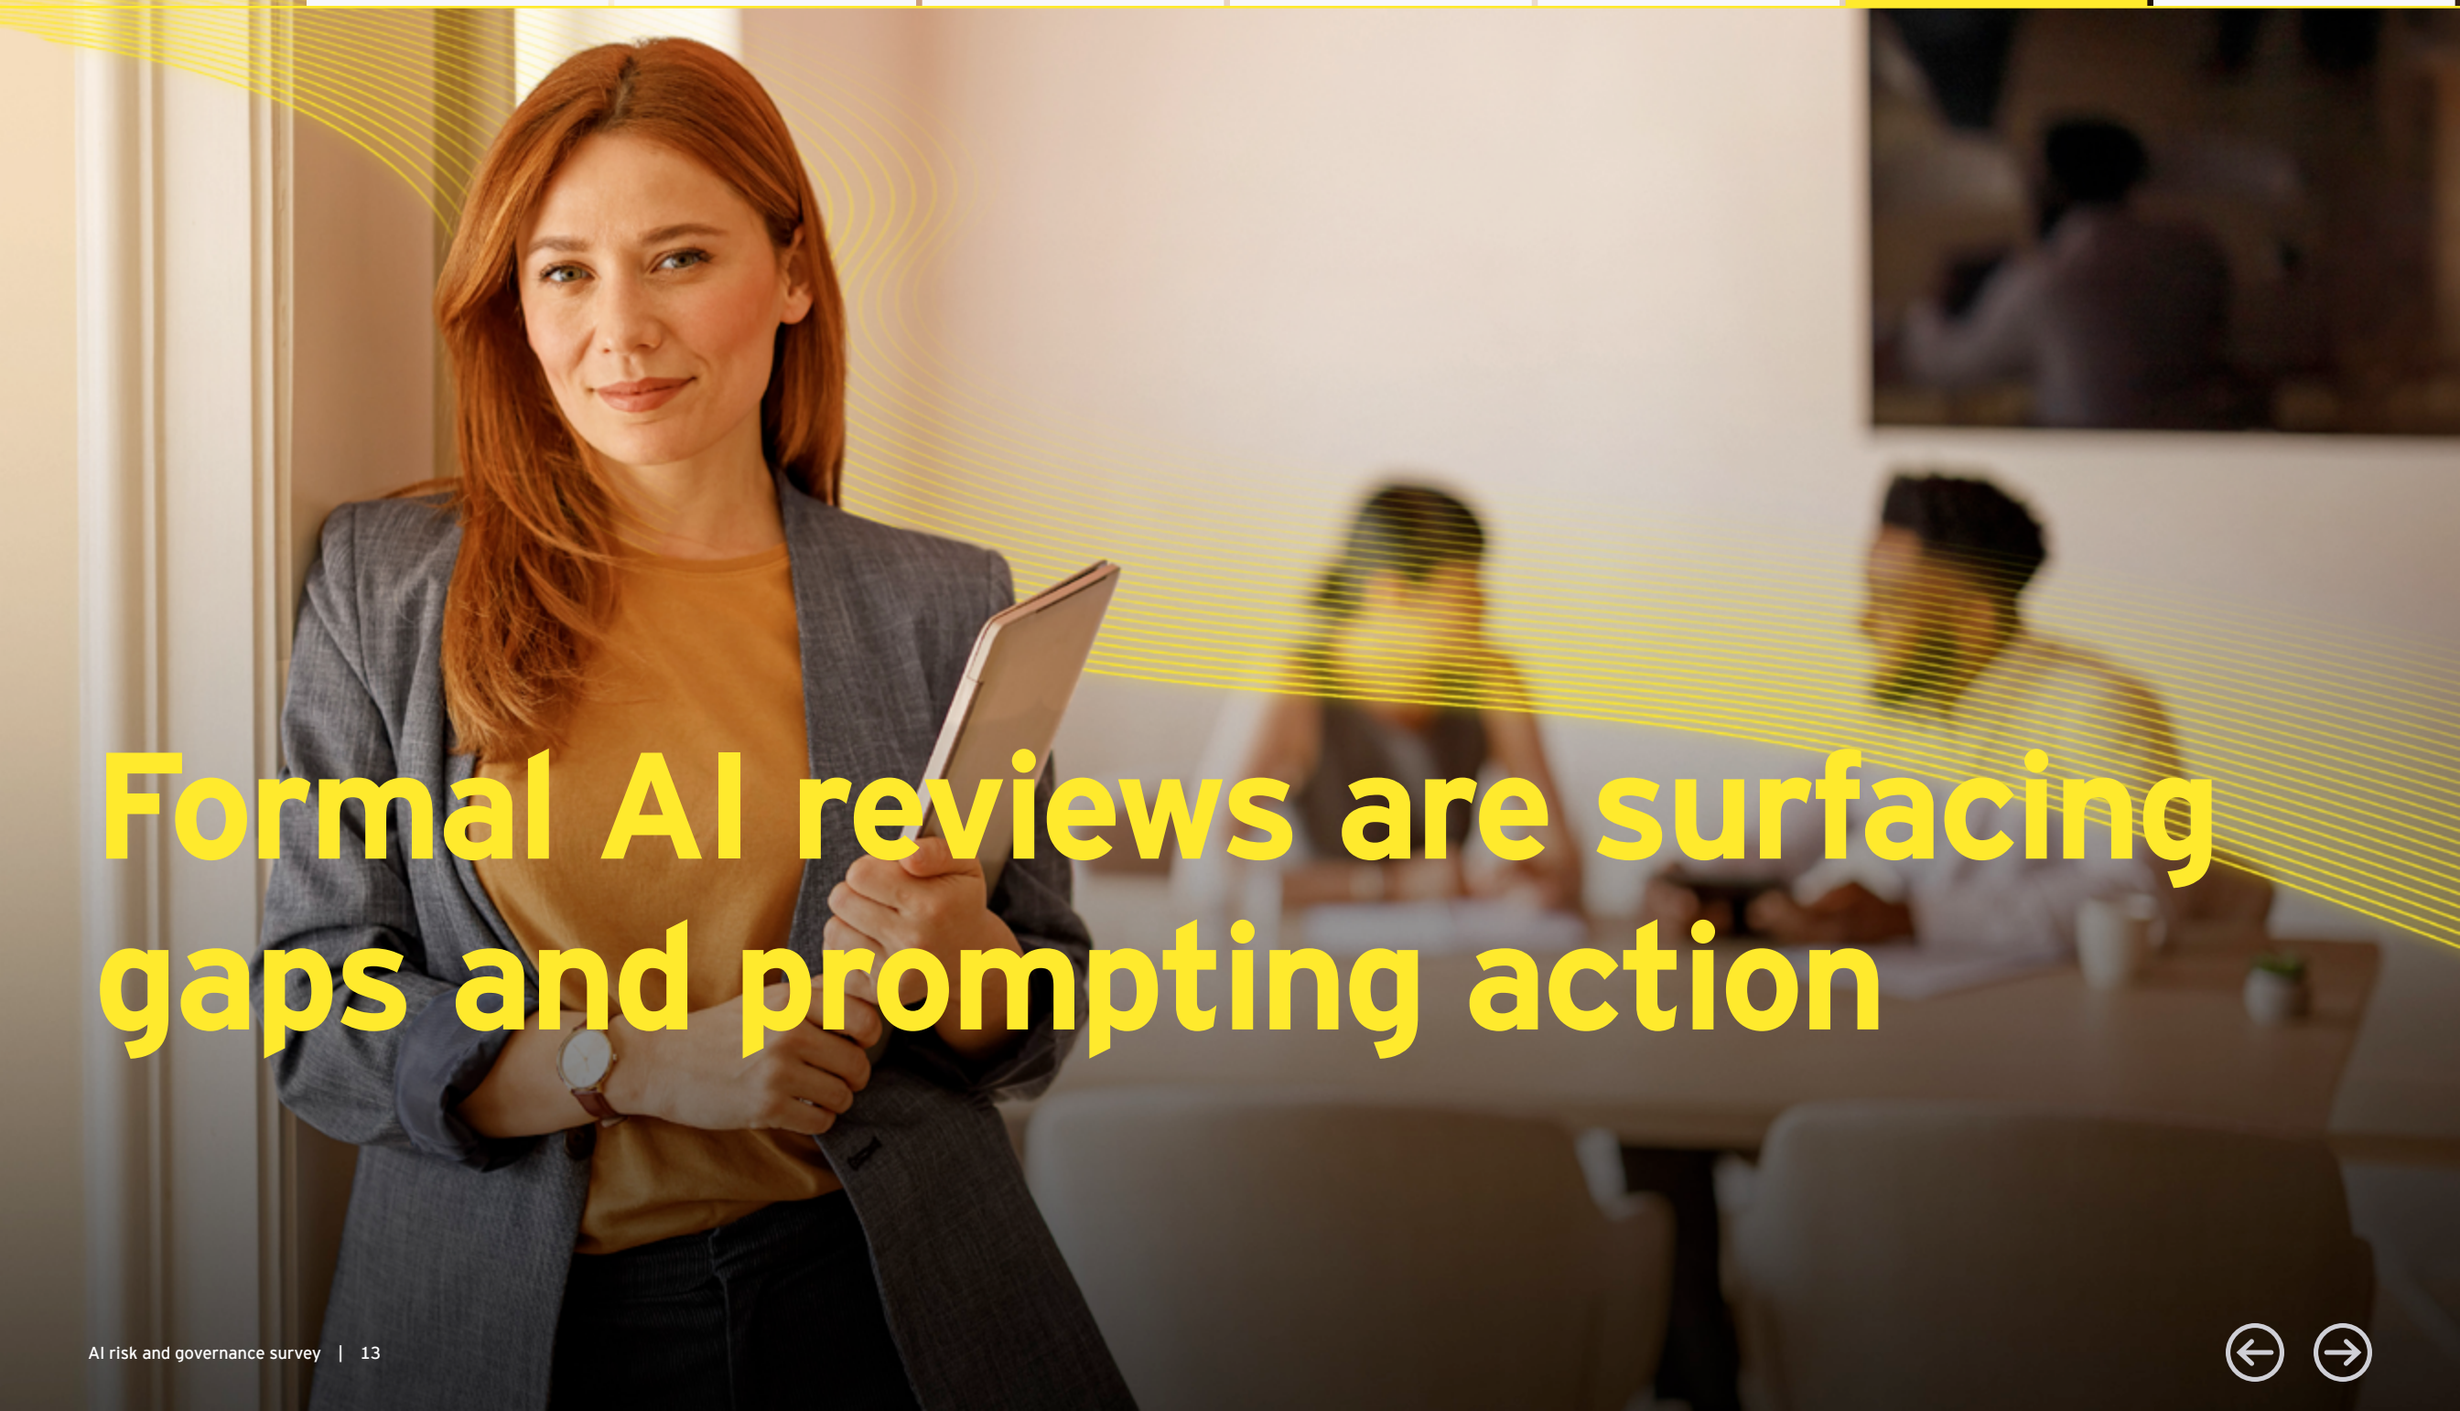

Formal AI reviews are surfacing gaps and prompting action

AI REVIEWS: GAPS AND ACTION

Formal risk and compliance reviews of AI systems are where the gap between governance design and governance effectiveness becomes more visible.

The research shows that formal review activity is becoming widespread, with 98% of senior AI decision-makers saying their organization has conducted a formal AI assurance review at least annually and 80% conducting them monthly or quarterly. That cadence suggests that many organizations recognize AI governance cannot be treated as a one-time approval exercise. As models, data, vendors and use cases change, assurance has to operate as a recurring discipline that evaluates whether controls are still aligned to the risks they were designed to manage.

The findings also show that these reviews are not purely procedural. Among organizations conducting formal AI assurance reviews, 92% found issues, and those findings led to meaningful action. Sixty-four percent of senior AI decision-makers said at least one-quarter of reviewed AI systems were significantly modified in the past year, while 29% said at least one-quarter were paused and later restarted and 25% said at least one-quarter were stopped entirely. Assurance reviews are surfacing issues that require remediation, escalation and, in some cases, stopping AI activity before risk becomes more consequential. The most common issues organizations find in their formal AI assurance reviews include data quality problems (57%), AI model drift (48%) and shadow AI (39%). Put together, the findings validate the urgent need for assurance reviews to determine whether agents are behaving as intended, that data lineage is reliable, and compliance obligations are met.

That is the role assurance can play in closing the confidence gap. Policies, registries and human-in-the-loop controls establish an important foundation, but formal AI reviews help assess whether those mechanisms are operating as intended and whether exceptions are being addressed. As AI becomes more autonomous and embedded in critical workflows, organizations will need assurance processes that do more than identify gaps after the fact. They will need reviews that are connected to ongoing monitoring, clear accountability, evidence-based remediation and governance decisions about when an AI system should be modified, paused or taken out of service.

How often organizations have conducted a formal AI assurance review

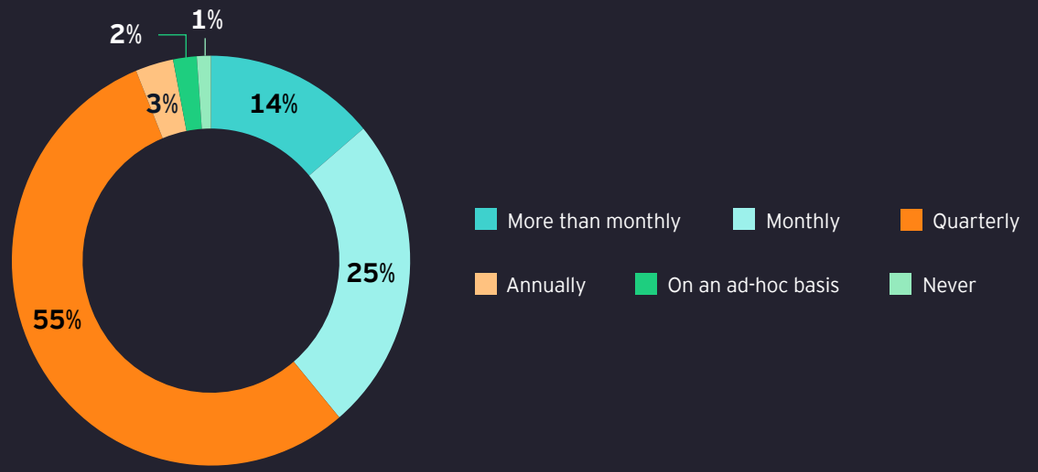

Due to rounding, totals may not add up to 100%.

98% of senior AI decision-makers said their organization has conducted a formal AI assurance review at least annually.

Closing the confidence gap

KEY TAKEAWAYS

The data points to a clear conclusion: AI governance has moved into the operating model, but confidence is lacking.

Organizations have built policies, established review processes, assigned oversight and begun to create the structures needed to manage AI responsibly. Yet the findings also show that formal governance and operational effectiveness are not the same thing. Controls can exist and still be bypassed. Registries can be created and still leave blind spots. Accountability can be defined broadly and still become unclear when autonomous agents begin acting inside the enterprise. And speed is the pressure point. Organizations will need help proving their AI governance is working and creating a continuing operating discipline.

This is the moment for decision-makers to move beyond governance as a framework and treat it as an operating discipline. That means:

- Creating visibility into the full AI estate
- Clarifying ownership across the lifecycle
- Embedding controls into the pace of deployment
- Making assurance a recurring source of evidence rather than a periodic checkpoint

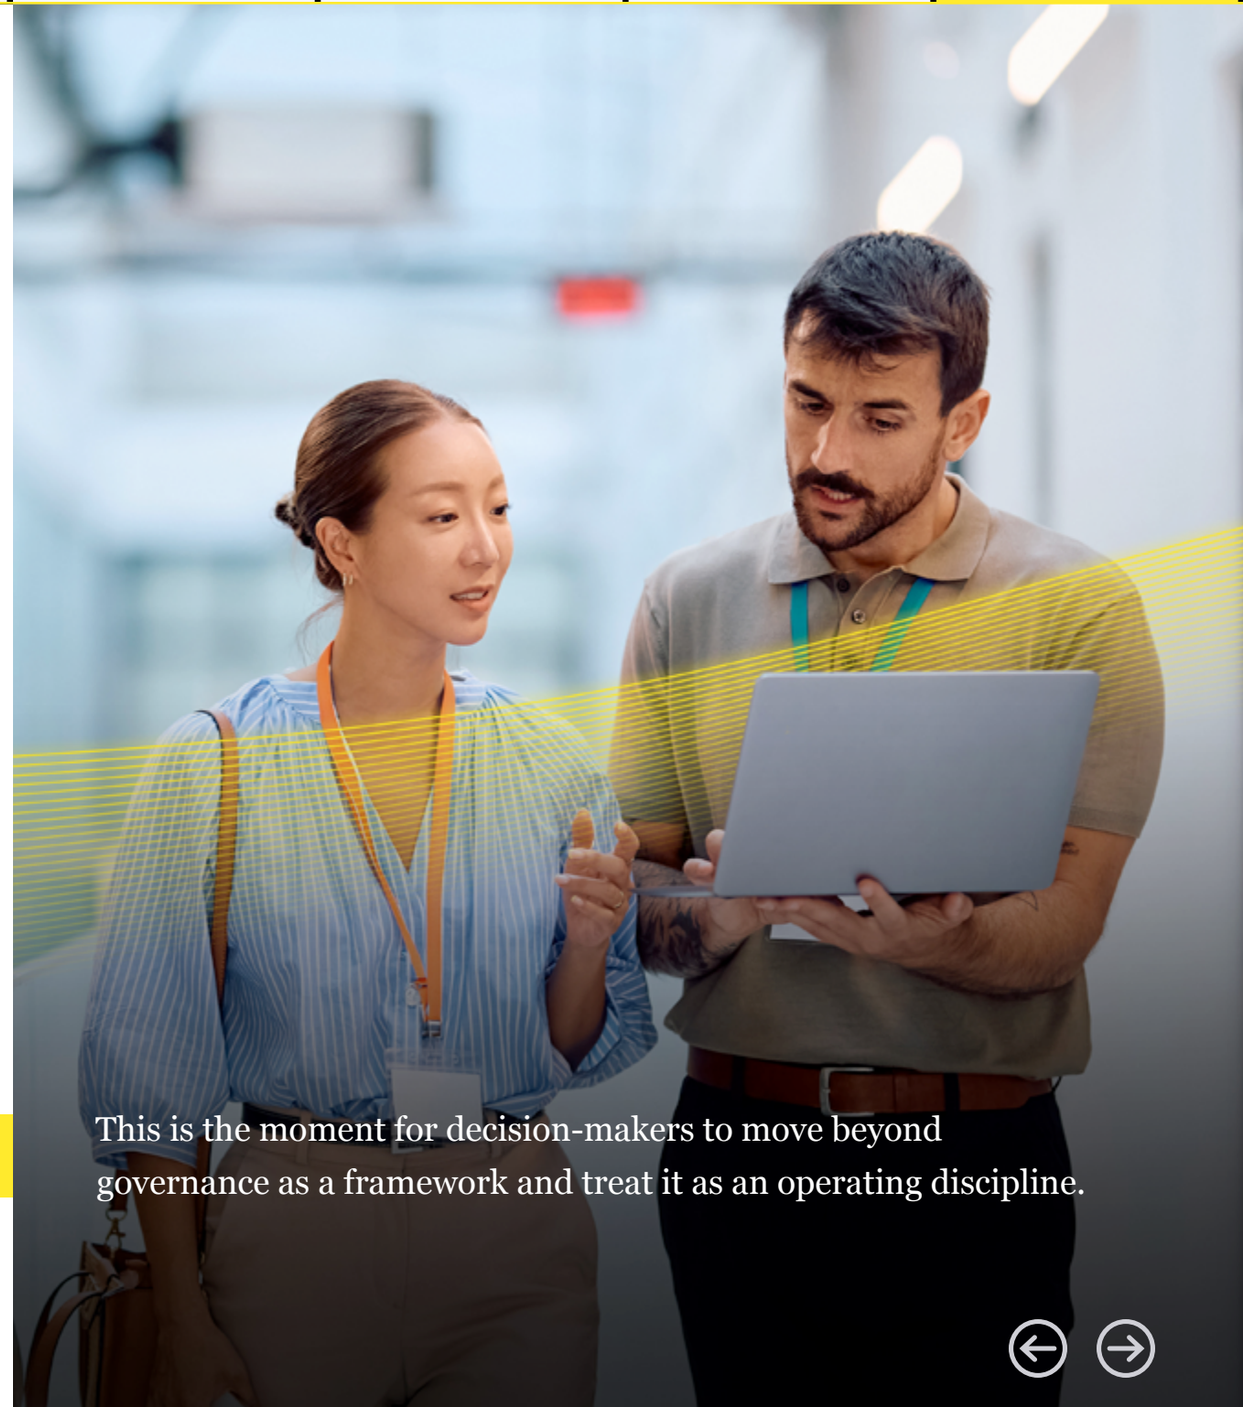

KEY TAKEAWAYS

The challenge is to make confidence scalable. To move faster while maintaining control.

Leaders should use this research as a baseline for asking harder questions:

- **Where is AI governance working?**
- **Where is it being stretched?**
- **Where can assurance help close the gap before risks become events?**

The organizations best positioned for the next phase of AI will be those that can pair ambition with evidence, speed with accountability and innovation with trust.

For boards and management teams, the findings point to the need to move beyond policy development and establish operational governance across the AI lifecycle. This includes maintaining visibility into authorized and unauthorized AI agents, defining accountability for autonomous decisions, and using independent reviews to determine whether controls are operating as intended.

AUTHORS

Richard J. Jackson
EY Global and Americas Assurance AI Leader

John McLain
EY Americas Deputy AI Leader

METHODOLOGY

EY Americas Assurance team commissioned an online survey among n=202 US senior AI decision-makers (board members, C-suite, VP+) at publicly traded companies with at least $1b+ in annual revenue. Respondents hold direct oversight over their organization's AI systems, governance or audit processes, with active deployment or pilot programs across traditional, generative or agentic AI. The survey was fielded between May 28 and June 15, 2026. The margin of error (MOE) for the total sample is plus or minus 7 percentage points at the 95% confidence interval.

This research was conducted by human researchers among human respondents. The fielding process was thoroughly supervised and screened for bots and AI agents to deliver the highest quality of data in accordance with industry standards.

EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multidisciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

US SCORE no. 32145-261US
2605-11926-CS
ED None

ey.com