How do you build a security roadmap for a shifting AI terrain?

Agentic AI, autonomous threats and the future of defense

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

EY

Shape the future with confidence

About the research

In December 2025, Ernst & Young LLP commissioned a study of 500 senior information security leaders across a broad range of industries including, banking and capital markets, wealth and asset management, insurance, oil and gas, technology, industrial products, private equity, media and entertainment, consumer products, retail, life sciences, and healthcare. Respondents were leaders at the director level or above, including a mix of C-suite leaders, who manage information security for companies with annual revenues of at least $500m. The goal of the study was to better understand how leaders are quantifying the value of artificial intelligence (AI) in security operations and defining their two-year budget and technology plans. Anchored in research, this study presents cyber leaders with insights on how to get the most out of their AI resources and investments.

1 The AI landscape in cybersecurity

The promise of AI in cybersecurity brings with it both the opportunities of automation and productivity and the threat of supercharged capabilities for those who seek to do harm. In both cases, the technology is turning out to be transformative.

As urgent as the threat is, both the speed and the stakes with which AI is transforming the enterprise demand a holistic and responsible approach. One factor — trust — is often a barrier to adoption, as features and function drive urgency for adoption, but confidence in outputs inhibits users from embracing new AI workflows. Trust in AI cannot be treated as an abstract sentiment or a single control; it must be engineered as a system, spanning cybersecurity, governance, compliance, transparency, explainability and ethics. Without this trust layer, even the most advanced AI capabilities struggle to scale beyond pilots into mission-critical operations.

The dichotomy of AI as protector of the enterprise and a looming threat manifests in what might seem to be contradictory sentiments from security professionals. Nearly all security leaders believe AI is a core defensive solution for cybersecurity (96%) and are already deploying AI in cybersecurity operations (95%). At the same time, 96% say AI-enabled cybersecurity attacks are a significant threat to their organization and less than half (46%) are strongly confident in their organization's ability to defend against a major security breach enabled by AI. Notably, 67% of respondents report still being in pilot mode — either experimenting with AI technology or exploring a strategy.

Preparedness becomes a matter of urgency since about half (48%) of senior security leaders stated that at least one quarter were AI-enabled in the past year.

It's important to note that security professionals recognize that the threat must be met by more than just technology. Greater spending, a commitment to human-in-the-loop frameworks, and cybersecurity governance maturity must coalesce to create both a trusted operational framework and deliver return on investment. Eighty-five percent of senior security leaders who are using AI in cybersecurity believe their organization's current cybersecurity budget is insufficient to meet AI-enabled threats. Only 20% of senior security leaders state their AI cybersecurity governance framework is fully optimized and embedded in organizational culture, leaving a wide gap for improvement.

The study showed a mix of results and progress, bright spots and room for improvement, and a universal sense that AI is foundational to building a modern cybersecurity program.

of security leaders say AI-enabled cybersecurity attacks are a significant threat to their organization.

Transformative AI

2 Financial realities: budgets, costs and investment trends

Security leaders are moving toward autonomous defenses. Virtually all respondents are confident that the strategic use of AI will transform their organization's proactive (99%) and defensive (99%) cybersecurity strategies.

However, the seriousness with which management is funding defensive efforts has not kept pace. Eighty-five percent of senior security leaders who are using AI in cybersecurity state their organization's current cybersecurity budgets are insufficient for AI-enabled threats, in part due to macroeconomic instability. Take healthcare for example, AI's limited uptake can be partly explained by the high cost of the technology, staff training and workflow adaptation.[1]

of security leaders say strategic use of AI will transform their proactive cybersecurity strategies.

The good news is security professionals expect their budget to increase, despite economic pressure. The number of senior security leaders dedicating at least 25% of their cybersecurity budget to AI solutions for cybersecurity specifically is expected to rise over the next two years from 9% today to 48%. From an overall spend standpoint, two-thirds (67%) of senior security leaders who use AI in cybersecurity expect to spend at least $5m two years out, and a third (34%) expect to spend at least $10m on the same time horizon.

For those who do make the investment, rewards are materializing. Outside of cybersecurity specifically, another recent EY US AI Pulse Survey: Wave 4 polled senior leaders on AI investments and ROI.[2] Of leaders of companies that were currently investing $10m or more in AI across all business units or teams, 71% saw significant AI-related productivity gains over the previous year compared to only 52% where investments were less than $10m.

99%

of security leaders say strategic use of AI will transform their defensive cybersecurity strategies.

1 Andrei Kasyanau, "Balancing The Cost of AI in Healthcare: Future Savings Vs. Current Spending," *Forbes*, https://www.forbes.com/councils/forbestechcouncil/2024/04/17/balancing-the-cost-of-ai-in-healthcare-future-savings-vs-current-spending, April 2024

2 The dividend age: "How AI is turning promise into payoff," *EY*, December, 2025

3 Return on investment: progress is slow, but optimism is high

We are beginning to see some early returns and efficiency gains as organizations adopt agentic AI in cybersecurity, but the numbers paint a picture of gains not yet realized.

About half (46%) of senior security leaders who use AI in cybersecurity report a return of less than $1m when using agentic AI solutions for cybersecurity, with an additional 12% reporting either not tracking returns or saw no savings at all. While these results vary by company, with many firms only automating mundane tasks which limit cost reductions, there's still plenty of room for improvement.

The next frontier will be to expand agentic AI to more core functions over time and shift employees' efforts to more strategic tasks. Cybersecurity professionals see targeting new areas where they expect to be largely run by agentic AI within two years (**FIGURE 1**).

FIGURE 1: Cybersecurity functions predicted to be largely agentic-run in two years

Firms aren't just tracking top-line or bottom-line results when it comes to AI, especially in cybersecurity. Risk reduction and competitiveness come into play as well. Senior security leaders recognize the benefits of using AI in cybersecurity in areas like identifying anomalies (67%), enhancing threat detection (60%), and in a predictive capacity, identifying vulnerabilities that might be targeted (59%). These nonmonetary but tangible benefits are expected to lead to a positive impact on metrics such as mean time to recovery (MTTR), mean time to detect (MTTD) and reduce false positive rate (FPR). Overall, security professionals overwhelmingly believe (97%) that their organization's competitive advantage in the marketplace hinges on the maturity of their agentic AI cybersecurity defense in the next two years.

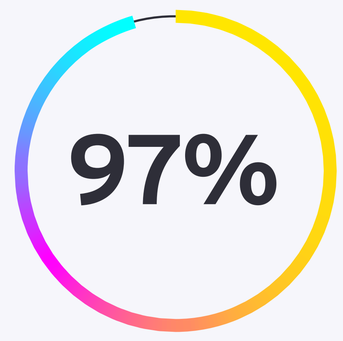

of security professionals overwhelmingly believe their organization's competitive advantage in the marketplace hinges on the maturity of their agentic AI cybersecurity defense in the next two years.

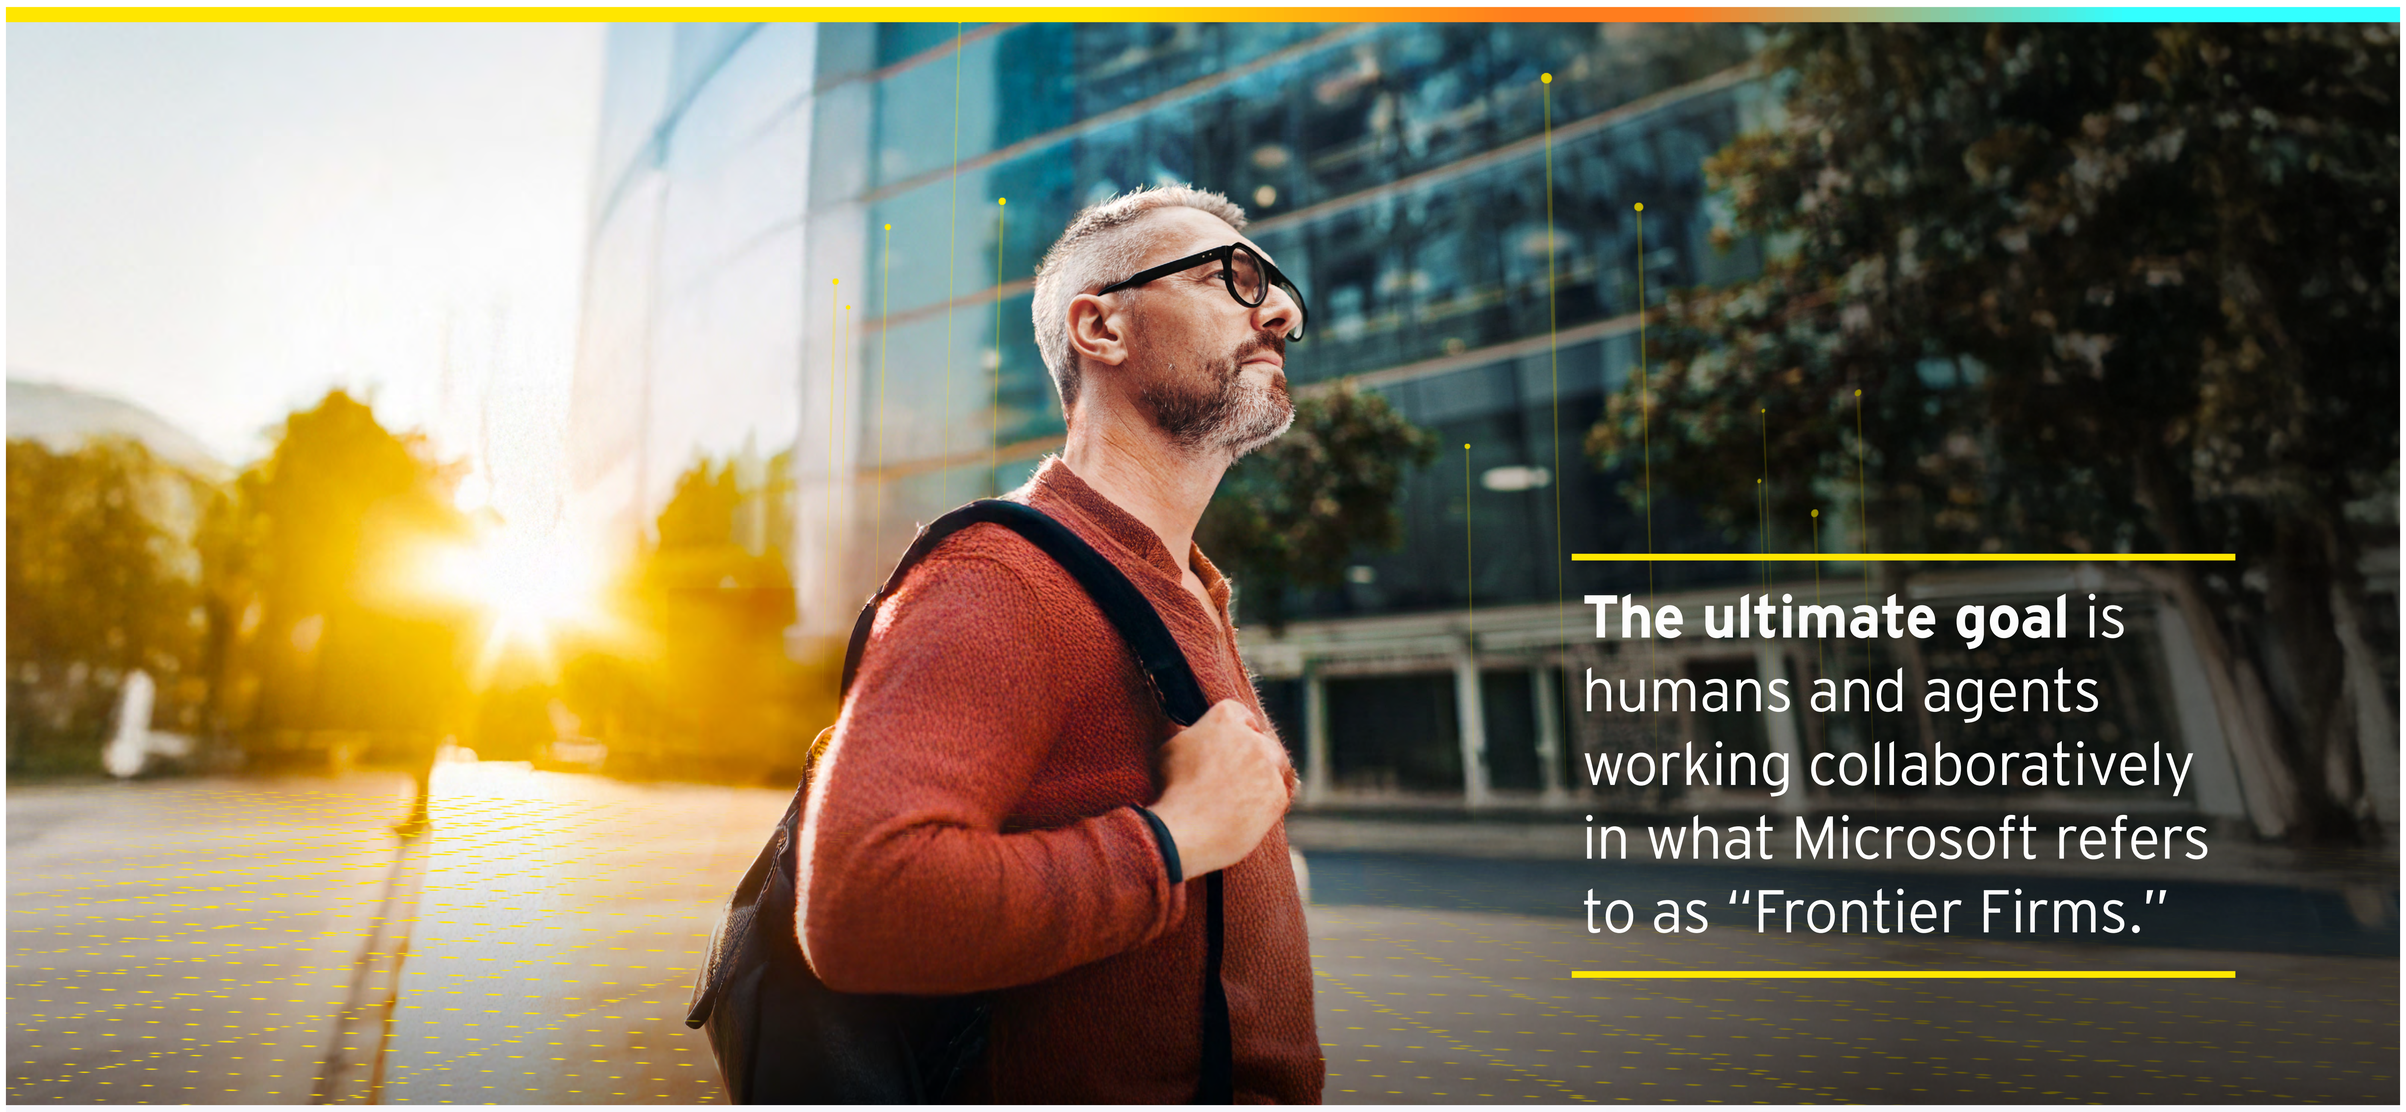

4 Human-in-the-loop cybersecurity: talent gaps and governance risks

As organizations accelerate their adoption of AI-driven cybersecurity tools, the human element has never been more critical.

Despite AI's ability to enhance analyst efficiency and shift focus toward higher-value tasks, most organizations continue to struggle with severe talent shortages and the need for skilled human-in-the-loop oversight. Human-in-the-loop cybersecurity is not merely a risk control, it is the primary mechanism through which organizations establish trust in AI-driven decisions. By embedding human judgment, context and accountability into AI workflows, enterprises create confidence that automation will augment, rather than undermine, responsible decision-making.

Eighty-five percent of senior security leaders say their organization maintains mandatory human-in-the-loop requirements for all critical security decisions and 98% believe that agentic AI in cybersecurity's successful return on investment depends on a clear human-in-the-loop oversight strategy. This makes it clear that human oversight remains a strategic safeguard, not a dispensable layer.

As organizations expand their use of AI-enabled security tools, analysts are tasked with validating and interpreting AI. Yet this need for deeper knowledge collides with a mismatched talent pool. Ninety percent of senior security leaders say their organization struggles with recruiting and retaining cybersecurity professionals with knowledge in AI-driven solutions or defenses, and 89% identify cybersecurity staff untrained on AI-enabled cyber attacks as their organization's greatest liability, a stark indication that AI does not reduce human risk when the workforce lacks the knowledge to govern it appropriately.

Training must be a bigger part of the answer, especially in the face of talent shortages. While there's a narrative that employees fear job displacement, there are many workers who see training and upskilling as a path to success. According to the EY Agentic AI Workplace Survey Results, 89% of desk workers (not just security employees), believe upskilling and re-skilling are crucial for staying relevant in an AI-augmented workplace.[3] Furthermore, 59% of desk workers cite lack of adequate training to develop agentic AI skills as an organizational barrier.

Underscoring this, AI is proving to be a significant force multiplier for the trained analysts organizations do have. Nearly all senior security leaders whose organization is using AI in cybersecurity say their organization's use of AI-driven cybersecurity solutions has demonstrably increased their human security analysts' operational efficiency (97%) and freed them up to focus on higher-value, strategic work (97%). Rather than replacing human knowledge, AI is pushing analysts toward roles that involve higher-order reasoning, governance and decision validation. The result is a cybersecurity landscape in which AI and human capability are deeply interdependent.

The ultimate goal is humans and agents working collaboratively in what Microsoft refers to as "Frontier Firms."[4] This concept manifests as organizations where operations are structured around on-demand intelligence and hybrid teams of humans plus agents. These firms can scale rapidly, operate with agility and generate value faster than their peers. This idea applied to cybersecurity means evolving security as a fully orchestrated automated function with humans at the center, collaborating with and directing AI agents.

3 "Unchanneled worker enthusiasm squanders agentic AI's promise," *EY*, October, 2025

4 "2025: The year the Frontier Firm is born," *Microsoft Work Trend Index Annual Report, Microsoft*, https://www.microsoft.com/en-us/worklab/work-trend-index/2025-the-year-the-frontier-firm-is-born?msockid=0a8508ac123f61690e411c5b13d96070, April, 2025

5 Governance as the foundation for responsible, scalable AI cybersecurity

As organizations race to adopt AI-driven cybersecurity capabilities, governance has emerged as the essential foundation for ensuring responsible, reliable and scalable deployment.

With most organizations already implementing or embedding governance frameworks into their core processes, governance is increasingly recognized as the key to converting AI's potential into real business value. Weak or immature governance, however, exposes organizations to critical risks from data breaches to compliance failures.

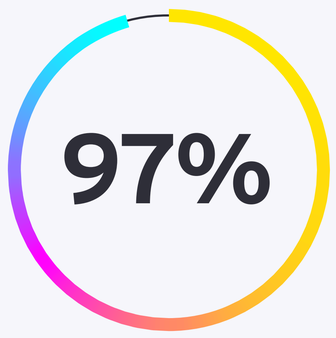

of security leaders report that a robust governance framework for AI in cybersecurity is essential to translating AI potential into profitable business value.

Organizations have already begun maturing their governance posture. While only 20% of senior security leaders state they've fully optimized and embedded an AI security governance framework into their organizational culture, some have made progress. More than half (51%) of senior security leaders report that governance frameworks are implemented and embedded into key processes and another 26% report their framework is fully rolled out and integrated across relevant business units (Figure 2). These gaps highlight a critical trust challenge: without governance embedded into day-to-day decision-making and culture, AI systems may function as designed yet still fail to earn the confidence required for broad adoption and long-term value creation.

These investments are not simply procedural. Organizations overwhelmingly believe governance is what unlocks the true value of AI in cybersecurity. Ninety-eight percent of senior security leaders whose organization has an AI governance framework for cybersecurity say governance frameworks have proven essential for the responsible use of AI, and 97% report that a robust governance framework for AI in cybersecurity is essential to translating AI potential into profitable business value.

These figures underscore a critical reality: without governance, AI becomes a capability full of promise but lacking guardrails; with governance, AI becomes a scalable, trusted, business-aligned system.

The risks of failing to invest in governance are equally clear. Weak or inconsistent governance leaves organizations exposed to some of the most damaging cybersecurity threats. Some of the top consequence are increased likelihood of data breaches (58%), compliance risks from agentic AI activity (58%), and data exposure (54%), which illustrate how governance gaps translate directly into real-world vulnerabilities. Instead of AI improving cyber resilience, poor oversight can inadvertently magnify organizational risk.

The pressures are further intensified by external uncertainty. With 91% of senior security leaders who have a cybersecurity framework expressing concern about the compliance risk associated with a lack of established government regulation for AI cybersecurity, governance is no longer merely an internal leading practice but a strategic imperative. In this unsettled regulatory environment, mature internal governance becomes the stabilizing force that confirms AI deployments remain safe, trusted, compliant and aligned with future requirements, regardless of how external policy evolves.

FIGURE 2: Maturity of organization's governance framework for AI in cybersecurity*

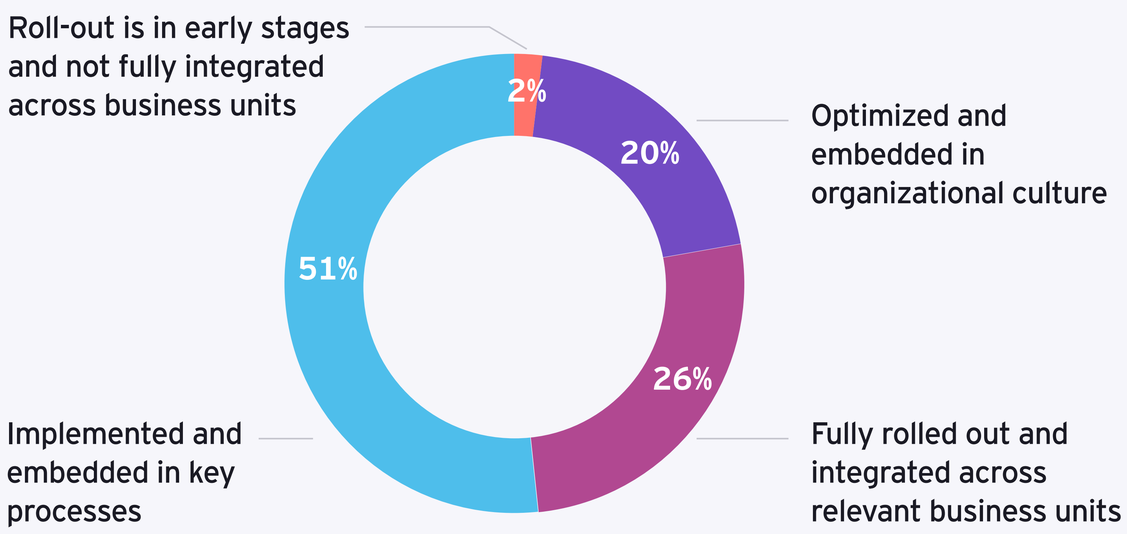

*Due to rounding, totals may not add up to 100%

Key takeaways and the way forward

The findings in this study illustrate a dual reality: AI has become indispensable for modern cyber defense, yet it simultaneously introduces new, complex risks that require urgent action. Security leaders must act in four key areas to drive value creation with AI:

1 Budget realities demand a reprioritization toward AI-driven cybersecurity

Cyber threats are advancing faster than most organizations' funding. Organizations must deliberately rebalance overall technology spend to prioritize AI-driven cybersecurity both to defend against AI-powered attacks and to avoid falling behind in the adoption of defensive capabilities. Underinvestment not only increases exposure to autonomous threats, but also limits organizations' ability to deploy AI responsibly, securely and at scale.

2 AI ROI requires deeper integration of agentic AI into core security functions

Meaningful return on investment emerges when agentic AI is embedded into cybersecurity. Organizations that move beyond pilots and task-level automation toward orchestrated, agent-driven security operations are better positioned to realize both financial returns and measurable improvements in resilience, speed and accuracy.

Failing to act in these critical areas can create risks for falling behind as AI reshapes both the threat landscape and the foundations of cybersecurity itself.

3 Human-in-the-loop oversight of AI and skills development are nonnegotiable

As AI systems take on greater autonomy, skilled human oversight becomes more critical, not less. Effective human-in-the-loop strategies depend on a workforce that understands how to validate AI outputs. Persistent talent shortages and skills gaps represent one of the greatest vulnerabilities in AI-enabled cybersecurity. Organizations must invest aggressively in re-skilling and upskilling their existing workforce so that humans can collaborate effectively with AI agents and maintain control over increasingly autonomous systems.

4 Governance is the foundation for trusted, scalable AI cybersecurity

Governance frameworks provide the guardrails so that AI systems are secure, compliant, transparent and aligned with organizational values and regulatory expectations. Without embedded governance spanning cybersecurity, compliance, ethics, transparency and explainability, AI initiatives risk amplifying exposure rather than reducing it. Organizations that treat governance as a living system — continuously improving and integrating into culture and operations — are best positioned to build trust, manage emerging risks and translate AI innovation into durable competitive advantage.

About the EY Cybersecurity Roadmap survey

The research was conducted via an online survey of the n=500 senior security leaders defined as full-time US employees at the director level and above (including 216 C-suite and 284 leaders below the C-suite level) who manage their organization's information security, including data and systems, at organizations with at least $500m in annual revenue across 12 industries. The survey was fielded between December 9, 2025 and January 8, 2026. The margin of error (MOE) for the total sample is +/-4 percentage points at the 95% confidence interval. Industries surveyed included banking and capital markets (n=50), wealth and asset management (n=50), oil and gas (n=30), consumer products (n=50), technology (n=30), industrial products (n=30), life sciences (n=50), private equity (n=30), retail (n=50), media and entertainment (n=30), healthcare (n=50) and insurance (n=50). Weighting was applied to distribute the sample evenly across all industries.

Authors

Ayan Roy
EY Americas Cybersecurity Competency Leader

Dan Mellen
EY Global Cyber Chief Technology Officer

Esther Lee
EY Americas Consulting, Cybersecurity

Ganesh Devarajan
EY Americas Consulting Cyber Risk Practice Leader

A special thanks to Martin Glowik and David Cooper for their contributions to this report.

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. 30287-261US
CS no. 2511-11675-CS

ey.com