Ponemon INSTITUTE

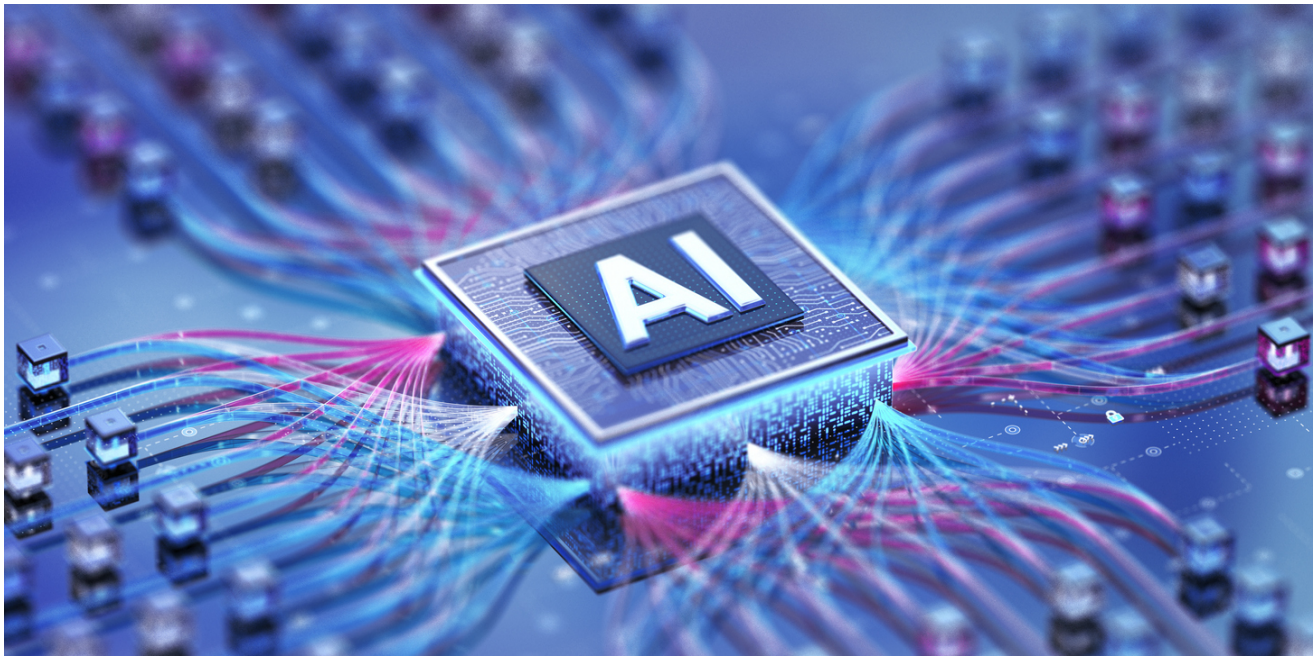

Managing Risks and Optimizing the Value of AI, GenAI & Agentic AI

Sponsored by OpenText

Independently conducted by Ponemon Institute LLC

March 2026

Managing Risks and Optimizing the Value of AI, GenAI & Agentic AI

Sponsored by OpenText
March 2026

The purpose of this research is to gain insight into how organizations are safeguarding their organizations from risks created by AI, GenAI and Agentic AI while still being able to benefit from their use. As shown in this research, organizations believe AI governance is important to realizing AI's value for security and business purposes. AI governance is the process of creating policies, assigning decision rights and ensuring organizational accountability for risks and investment decisions in the application and use of AI technologies.

AI's value is dependent upon good governance, non-human identity management and explainability. Respondents were asked to rate their organizations' focus on governance of AI systems from 1 = low focus to 10 = very high focus. As shown in Figure 1, 53 percent of respondents say their organizations are highly or very highly focused on governance (7+ responses on the 10-point scale).

Non-human identity management (NHIM) is the process of discovering, securing and managing digital credentials for applications, machines and automated processes that lack a human operator. This practice is considered crucial for cybersecurity because the number of these identities is growing rapidly and they are often overlooked, creating significant security vulnerabilities.

Respondents were asked to rate the priority of managing non-human identities on a scale from 1 = low priority to very high priority. Fifty-one percent of respondents say it is a high or very high priority (7+ responses on a 10-point scale).

Explainability in AI is the ability to understand how an AI system makes decisions or predictions, and why it made them. It is also known as explainable machine learning or explainable AI (XAI). Respondents were asked to rate the importance of explainability in AI-powered solutions from 1 = low importance to 10 = highly important. Forty-seven percent of respondents say explainability is very or highly important (7+ on the 10-point scale).

Figure 1. How focused are organizations on AI governance, the priority of managing non-human identities and the importance of explainability?
On a scale from 1 = low focus/importance/priority to 10 = very high focus/importance/priority, 7+ responses presented

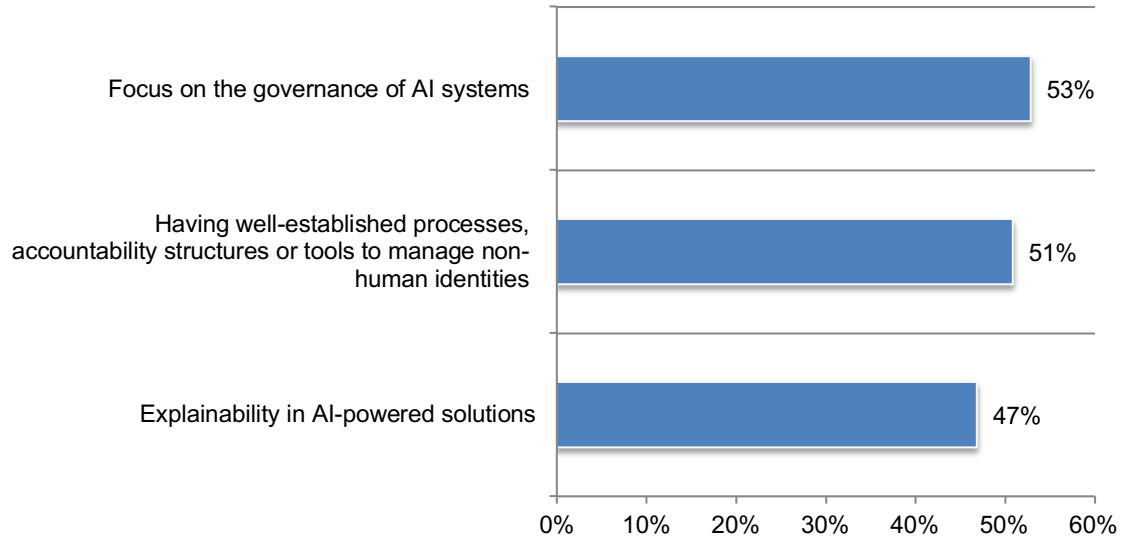

Summary of research findings

Most organizations have not achieved AI maturity, which impacts its effectiveness and value. Only 21 percent of respondents say they have achieved AI maturity. In the mature stage, AI in cybersecurity activities is fully deployed and security risks are assessed. Effectiveness of AI is measured with KPIs and C-level executives are regularly informed about AI's ability to prevent and reduce cyberattacks.

The top governance challenges are the lack of staff and budget needed to implement and maintain AI-based technologies. Fifty percent of respondents say AI deployments require too much staff to implement and maintain AI-based technologies and 44 percent of respondents say the staff does not have enough time to integrate AI-based technologies. Forty-six percent of respondents say there is insufficient budget for AI-based technologies.

As part of their AI governance strategy, 43 percent of respondents say their organization has adopted a Risk-based AI governance approach that focuses on identifying, assessing, and mitigating AI-related risks (like bias, security threats, or ethical issues). Of these respondents, 72 percent say they use the NIST AI risk management framework (27 percent), use the ISO framework for AI systems (30 percent) or use both (22 percent of respondents).

Procedures and policies for the ethical use of AI and compliance with regulations is the most important step in a risk-based approach. Sixty-one percent of respondents say their organizations create policies and procedures for AI usage, ethical considerations and compliance with regulations. Fifty-nine percent of respondents say they identify potential risks such as bias in training data or model drift and assess their likelihood and impact. Fifty-three percent of respondents say their organizations continuously monitor the speed of AI performance, evaluate decisions and update models to prevent performance degradation or the emergence of harmful behaviors.

Organizations are skeptical that AI models can learn robust norms and make safe decisions autonomously. Only 47 percent of respondents say their AI models can learn robust norms and make safe decisions autonomously. Less than half (48 percent) say it will be possible in the future to have AI systems that reason and make autonomous decisions based on ethics, regulations and laws to avoid misuse. More than half of respondents (51 percent) say human oversight is needed in AI governance because of the speed in which attackers can adapt.

Errors and inaccuracies are the biggest barrier to AI effectiveness. The two biggest obstacles are errors and inaccuracies in AI decision rules (45 percent of respondents) and errors and inaccuracies in data inputs ingested by AI technology (engine) (40 percent of respondents). Other barriers are the inability to create a unified view of AI users across the enterprise (39 percent of respondents) and the heavy reliance on legacy IT environments (38 percent of respondents).

Organizations face great difficulty in minimizing business and security risks created by AI. Sixty-two percent of respondents say it is very or extremely difficult to minimize model and bias risks such as the breach of ethical and responsible AI principles in the language model development leading to discriminatory or unfair outputs.

Sixty-one percent of respondents say it is very or extremely difficult to minimize such data risks as error propagation and misleading and harmful content caused by low-quality data used to train generative AI models. It is very or extremely difficult in minimizing such prompt or input risks such as misleading, inaccurate or harmful responses due to unsophisticated prompts or questions being provided to the AI model, according to 58 percent of respondents. Fifty-six percent of respondents say it is very or extremely difficult to minimize user risks such as the unintended consequences due to users becoming unwitting parties to the creation of misinformation and other harmful content.

While 59 percent of respondents say AI makes it highly difficult to comply with privacy and security regulations and mandates, only 41 percent of respondents say their organizations have data privacy policies specifically for the use of AI. Fifty-nine percent of respondents with a privacy policy say it is important to conduct assessments of how AI impacts their commitments to privacy. Another 59 percent of respondents say policies should require data stewardship requirements to mitigate uses of personal information that are adverse or unfair to individuals the data relates to.

Risks created by GenAI are data misinformation and concerns about security and privacy. Only half (50 percent) of respondents say their organizations take steps to reduce security risks pertaining to the usage of GenAI. Of these respondents, 55 percent of respondents say their organizations conduct training programs to familiarize everyone with the risks and rewards of GenAI. Fifty percent of respondents say their organization uses experts to validate rough draft generative AI outputs.

The lack of proper risk and security controls and complex system integration are reasons not to adopt Agentic AI. Thirty-eight percent of respondents say their organizations have fully adopted (15 percent) or partially adopted (23 percent) Agentic AI. Respondents report an average of 23 percent of employees use AI agentic agents that autonomously perform sequences of tasks such as coding, email response or data queries. Nineteen percent of organizations have no plans to adopt Agentic AI. The primary reasons are the lack of proper risk and security controls (43 percent of respondents) and complex system integration (38 percent).

Organizations rate the effectiveness of Agentic AI in automating manual tasks and removing human error as low. Only 44 percent of respondents rate the effectiveness of automating manual tasks for security teams with limited human intervention as high. Only 39 percent of respondents say Agentic AI is highly effective in removing human error by systematically retrieving comprehensive threat intelligence.

The risk of data theft caused by AI agents is high. Agentic AI can both be a powerful tool for preventing data theft by automating security tasks and a source of new vulnerabilities that require updated data protection strategies. However, the malicious use of Agentic AI significantly increases the risk of data theft. Fifty-five percent of respondents believe AI agents will significantly increase data theft (29 percent) or moderately increase the risk of data theft (26 percent).

AI agents make detection of an intrusion more difficult. The malicious use of AI agents significantly complicates intrusion detection by enabling faster, stealthier attacks that mimic legitimate behavior, creating convincing deepfakes, automating complex reconnaissance, and poisoning training data. Sixty-six percent of respondents say AI agents make detection of an intrusion more difficult (41 percent) or moderately more difficult (25 percent). Only 12 percent of respondents say AI agents **do not** make intrusion detection more difficult.

Of the organizations that have deployed GenAI, only 45 percent of respondents rate its effectiveness in generating actionable insights to support better security decisions as highly effective. Fifty-two percent of respondents say their organization has fully deployed (18 percent) or partially deployed (34 percent) GenAI. Twenty-three percent plan to deploy in the future. Respondents report that an average of 25 percent of work tasks has been automated by GenAI. Twenty-five percent of respondents have no plans to adopt GenAI. The reasons are the lack of in-house expertise and concerns about accuracy, bias and misinformation.

AI is slightly more effective in reducing the time to detect anomalies, new patterns and emerging threats. Fifty-one percent of respondents say AI is very or highly effective in reducing the time to detect anomalies, new patterns and emerging threats. Less than half of respondents (48 percent) say effectiveness is high in threat detection and hunting to provide deeper contextual insights and reducing the manual workload.

Part 2. Key findings

Ponemon Institute surveyed 1,878 IT and IT security practitioners knowledgeable about their organizations' use of AI for cybersecurity and business purposes in North America (611 respondents), EMEA (496 respondents), AsiaPac (428 respondents) and LATAM (343 respondents).
In this section, a deeper dive into the global research is presented. The complete findings are shown in the Appendix. The report is organized according to the following topics.

- The importance of AI governance to achieving the balance between value and risk
- The difficulty of AI deployment and risk management
- Privacy, security and ethical considerations in AI
- The benefits and risks of GenAI and Agentic AI
- AI security and business risk perception gaps between the C-suite and those in the trenches
- Regional differences

The importance of AI governance to achieving the balance between value and risk

Most organizations have not achieved AI maturity which impacts its effectiveness and value. Figure 2 presents the stages of AI deployment maturity. Forty-two percent of respondents say their organizations are in the early adoption stage, which is characterized by a defined and planned cybersecurity AI strategy, but it is only partially deployed. In the full adoption stage, 37 percent of respondents say AI in cybersecurity technologies are mostly deployed with C-level support and adequate budget.

Only 21 percent of respondents say they have achieved AI maturity. In the mature stage, AI in cybersecurity activities is fully deployed and security risks are assessed. Effectiveness of AI is measured with KPIs and C-level executives are regularly informed about AI's ability to prevent and reduce cyberattacks.

Figure 2. What best describes the maturity of your organization's use of AI
Only one choice permitted

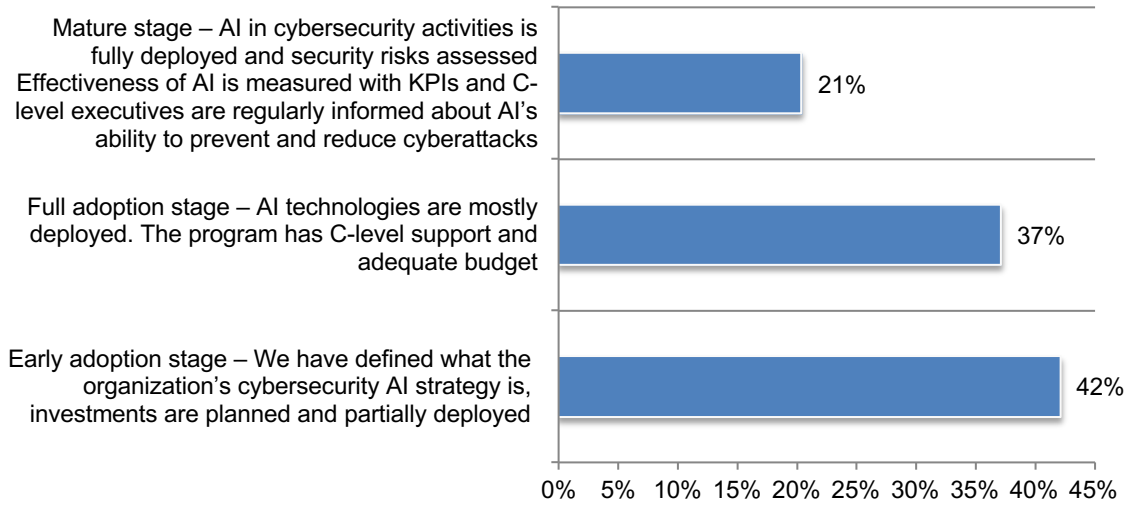

The top governance challenges are the lack of staff, budget and time needed to implement and maintain AI-based technologies. According to Figure 3, 50 percent of respondents say AI deployments require too much staff to implement and maintain AI-based technologies and 44 percent of respondents say the staff does not have enough time to integrate AI-based technologies. Forty-six percent of respondents say there is insufficient budget for AI-based technologies.

Figure 3. What are the organizational and governance challenges experienced when deploying AI?
Three responses permitted

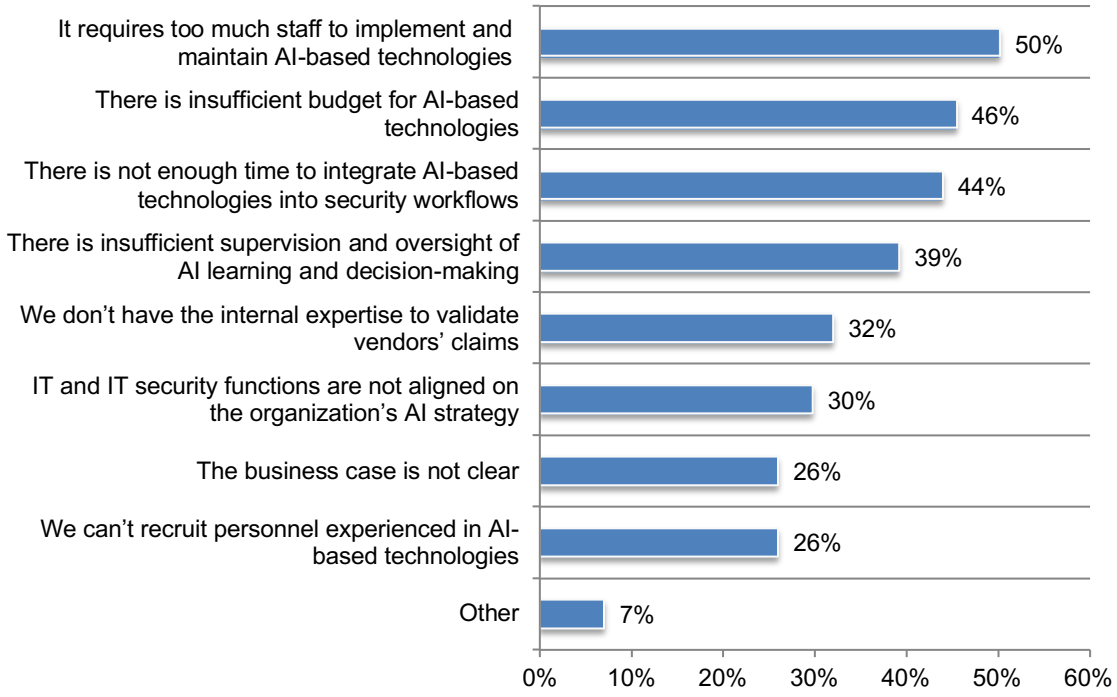

As part of their AI governance strategy, 43 percent of respondents say their organization has adopted a Risk-based AI governance approach that focuses on identifying, assessing, and mitigating AI-related risks (like bias, security threats, or ethical issues). Resources are prioritized to focus on high-risk applications across the AI lifecycle (inception to retirement) to ensure safety, fairness, and compliance. Of these respondents, 72 percent say they use the NIST AI risk management framework (27 percent), the ISO framework for AI systems (30 percent) or use both (22 percent of respondents).

Procedures and policies for the ethical use of AI and compliance with regulations is the most important step in a risk-based approach. The steps in a risk-based approach are shown in Figure 4. Sixty-one percent of respondents say their organizations create policies and procedures for AI usage, ethical considerations and compliance with regulations.

Fifty-nine percent of respondents say they identify potential risks such as bias in training data or model drift and assess their likelihood and impact. Fifty-three percent of respondents say their organizations continuously monitor the speed of AI performance, evaluate decisions and update models to prevent performance degradation or the emergence of harmful behaviors.

Figure 4. Steps taken in a risk-based approach to minimize AI risks
More than one response permitted

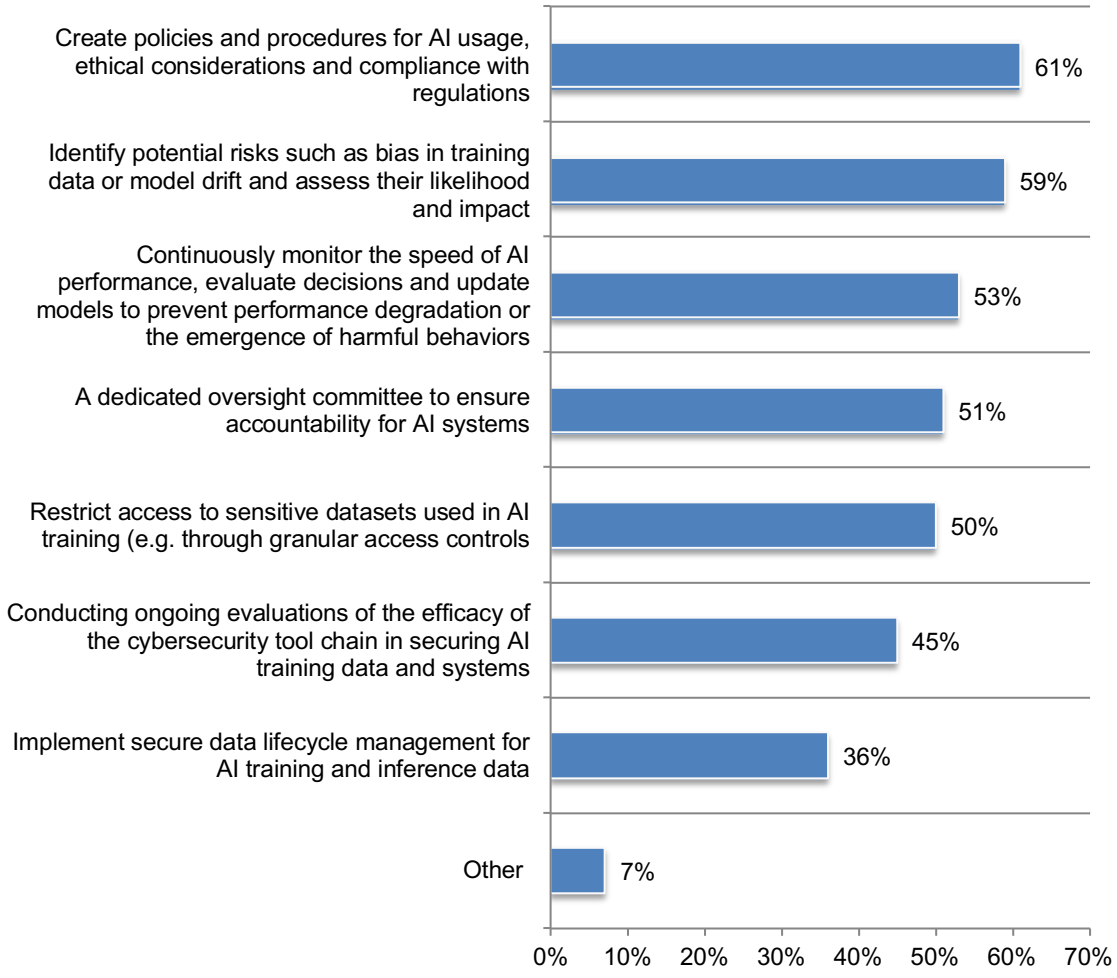

The functions most involved in AI governance are IT and IT security. As shown in Figure 5**,** 68 percent of respondents say the CIO (32 percent), the CISO (21 percent) and CTO (15 percent) are involved in AI governance and practices. Also mostly involved are the Chief Compliance and Privacy Officer (29 percent) and the Chief Legal Officer (23 percent).

Thirty-four percent of respondents say their organization has or will have a Chief AI Officer (CAIO). As the maturity of AI deployment evolves, more organizations may consider hiring a CAIO. The role of a CAIO is to lead the AI strategy, from vision and implementation to governance, ethics, and value creation. They help ensure AI is ethical, compliant, and effective in achieving strategic goals. They are considered different from CIOs or CTOs by focusing specifically on AI's transformative impact.

Figure 5. Which of the following functions are involved in AI governance and policies in your organization?
More than one response permitted

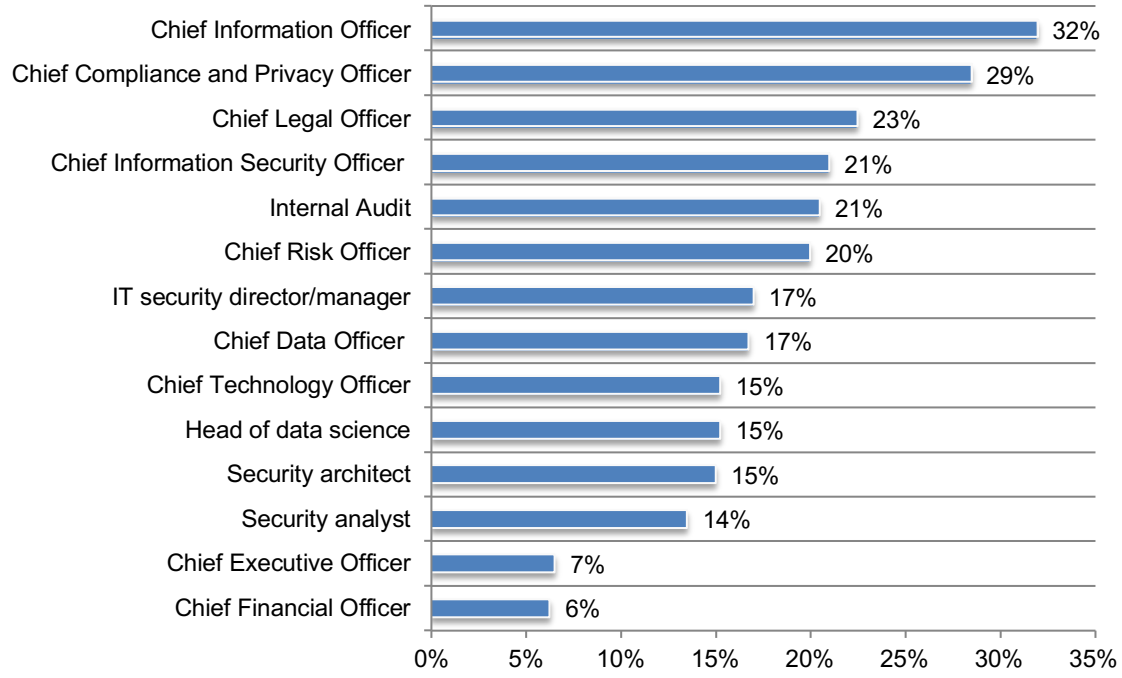

Organizations are skeptical that AI models can learn robust norms and make safe decisions autonomously. As shown in Figure 6, only 47 percent of respondents say their AI models can learn robust norms and make safe decisions autonomously. Less than half (48 percent) say it will be possible in the future to have AI systems that reason and make autonomous decisions based on ethics, regulations and laws to avoid misuse. More than half of respondents (51 percent) say human oversight is needed in AI governance because of the speed in which attackers can adapt.

Figure 6. Perceptions about AI decision making and the need for human oversight
Strongly agree and Agree responses combined

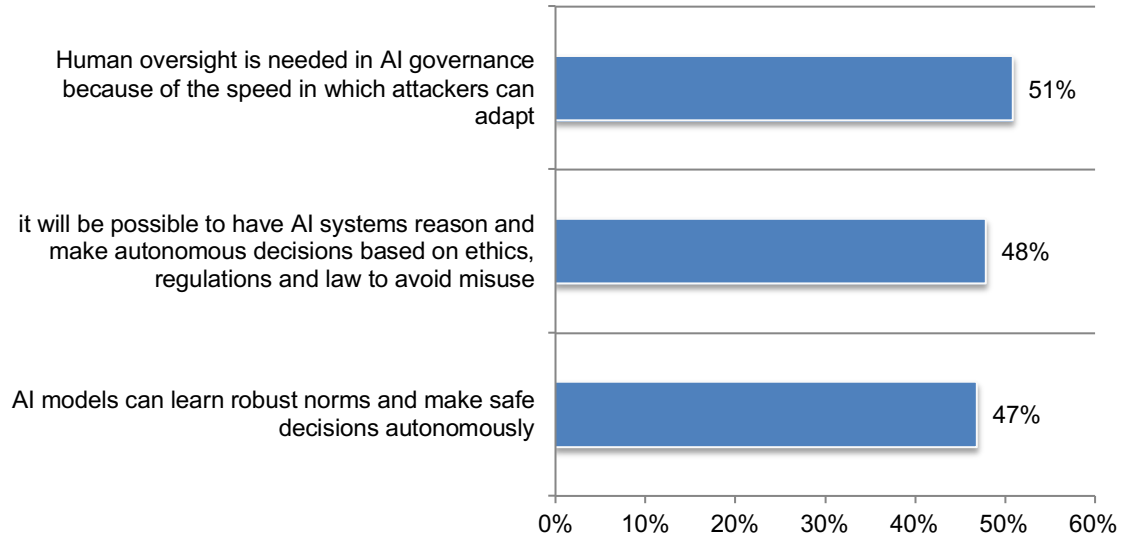

The difficulty of AI deployment and risk management

Errors and inaccuracies are the biggest barrier to AI effectiveness. As shown in Figure 7, the two biggest obstacles are errors and inaccuracies in AI decision rules (45 percent of respondents) and errors and inaccuracies in data inputs ingested by AI technology (engine) (40 percent of respondents). Other barriers are the inability to create a unified view of AI users across the enterprise (39 percent of respondents) and the heavy reliance on legacy IT environments (38 percent of respondents).

Figure 7. Barriers to the effectiveness of AI-based technologies used in organizations
Three responses permitted

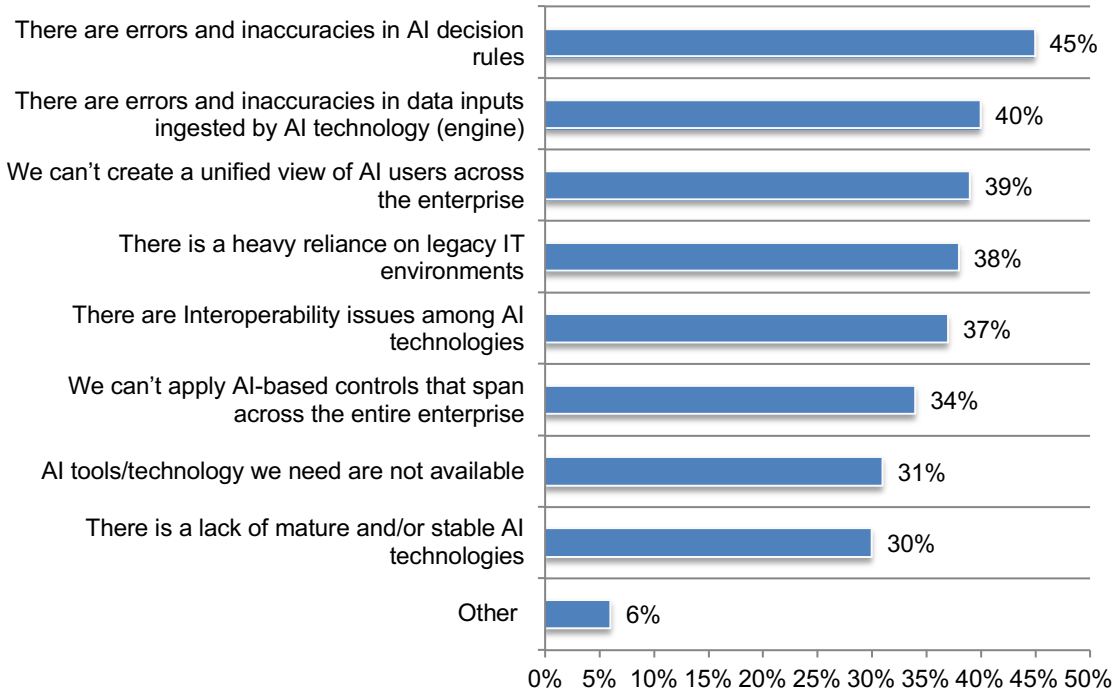

AI is slightly more effective in reducing the time to detect anomalies, new patterns and emerging threats. Respondents were asked to rate the effectiveness of the use of AI in threat detection and hunting to reduce the time to detect anomalies, new patterns and emerging threats that might otherwise go undetected on a scale from 1= low effectiveness to 10 = high effectiveness. As shown in Figure 8, 51 percent of respondents say AI is very or highly effective in achieving this goal (7+ responses on the 10-point scale).

Respondents were also asked to rate effectiveness in learning deeper contextual insights and reducing the manual workload from threat detection and hunting on a scale from 1 = low effectiveness to 10 = high effectiveness. Less than half of respondents (48 percent) say effectiveness is high (7+ on the 10-point scale).

Figure 8. Effectiveness in the use of AI in threat detection and hunting
On a scale from 1 = low effectiveness to 10 = high effectiveness, 7+ response presented

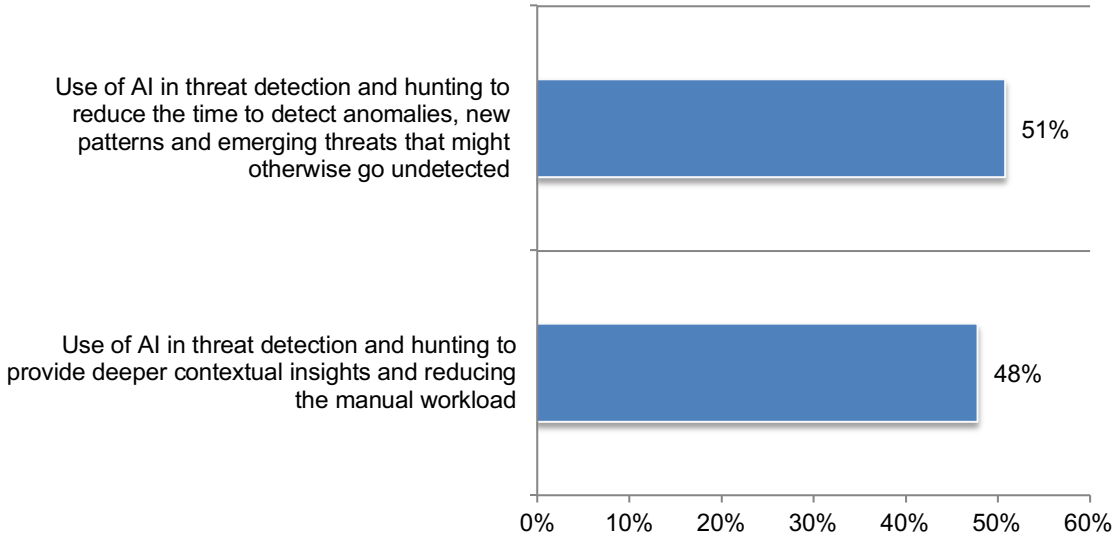

Organizations face great difficulty in minimizing risks created by AI. Respondents were asked to rate the difficulty in minimizing AI risks on a scale from 1 = low difficulty to 10 = extremely difficult. All responses shown in Figure 9 are 7+ on the 10-point scale. Sixty-two percent of respondents say it is very or extremely difficult to minimize model and bias risks such as the breach of ethical and responsible AI principles in the language model development leading to discriminatory or unfair outputs.

Sixty-one percent of respondents say it is very or extremely difficult to minimize such data risks as error propagation and misleading and harmful content caused by low-quality data used to train generative AI models.

It is very or extremely difficult in minimizing such prompt or input risks such as misleading, inaccurate or harmful responses due to unsophisticated prompts or questions being provided to the AI model, according to 58 percent of respondents.

Fifty-six percent of respondents say it is very or extremely difficult to minimize user risks such as the unintended consequences due to users becoming unwitting parties to the creation of misinformation and other harmful content.

Figure 9. The difficulty in minimizing AI risks
On a scale from 1 = low difficulty to 10 = extremely difficult, 7+ responses presented

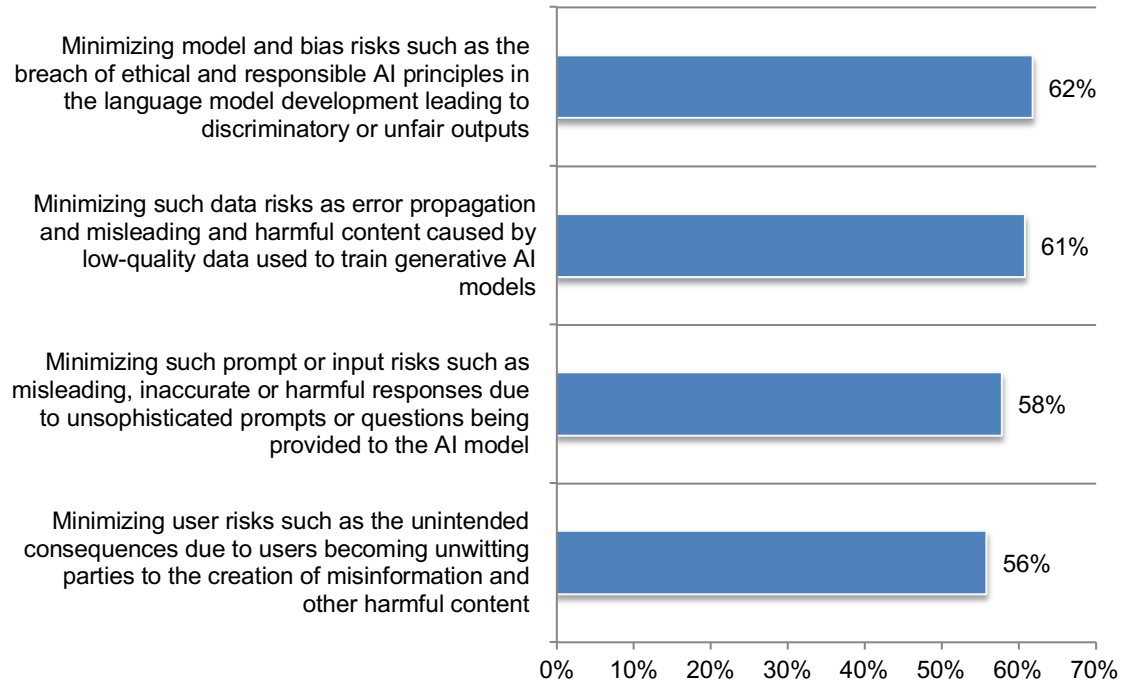

Privacy, security and ethical considerations in AI

While 59 percent of respondents say AI makes it is highly difficult to comply with privacy and security regulations and mandates, only 41 percent of respondents say their organizations have data privacy policies specifically for the use of AI.

As shown in Figure 10, 59 percent of respondents with a privacy policy say it is important to conduct assessments of how AI impacts their commitments to privacy. Another 59 percent of respondents say policies should require data stewardship requirements to mitigate uses of personal information that are adverse or unfair to individuals the data relates to.

Figure 10. What is included in privacy policies
More than one response permitted

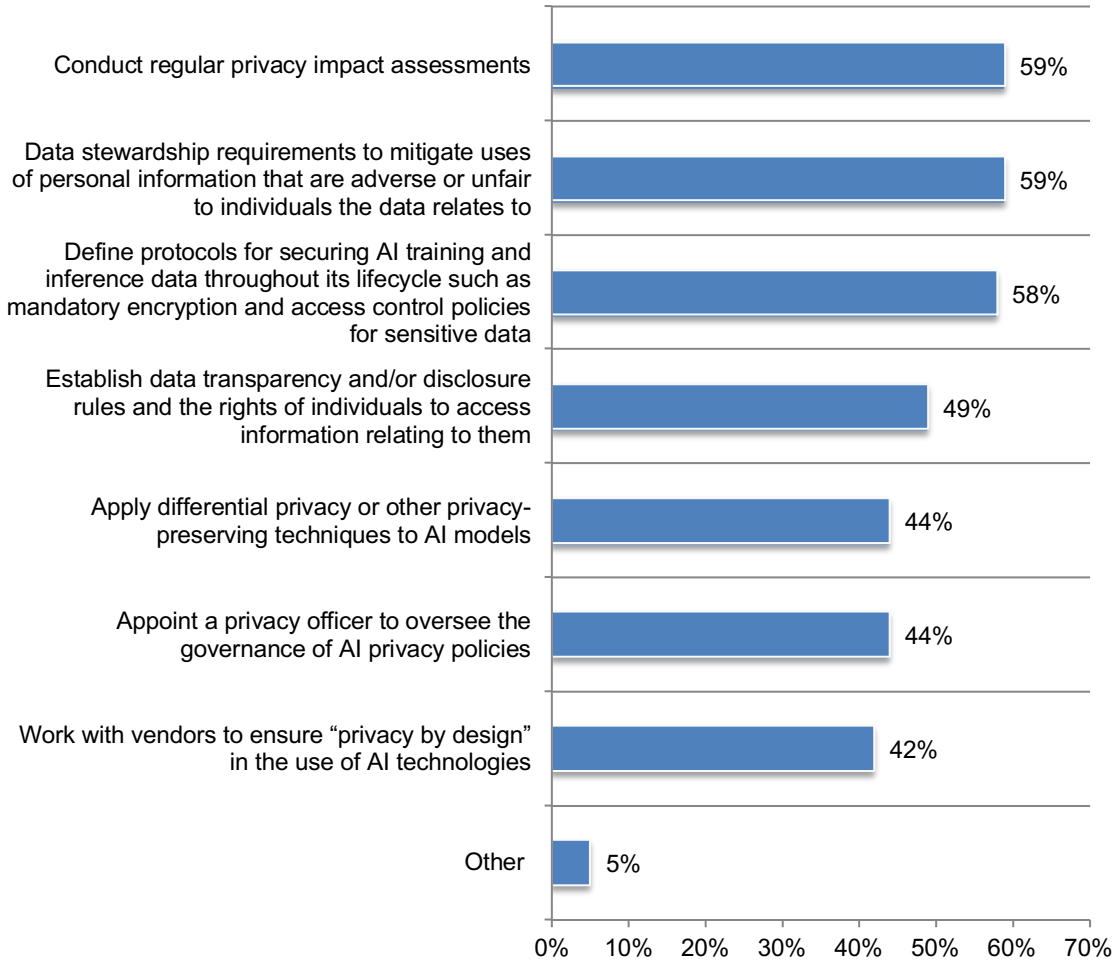

The benefits and risks of GenAI and Agentic AI

GenAI refers to a category of AI algorithms that generates new outputs based on the large language models they have been trained on. Dynamic-based generative AI models are better positioned to analyze complex systems, such as network infrastructures or software applications to identify vulnerabilities, detect novel threats and mitigate risks.

Fifty-two percent of respondents say their organization has fully deployed (18 percent) or partially deployed (34 percent) GenAI. Twenty-three percent plan to deploy in the future. Respondents report that an average of 25 percent of work tasks has been automated by GenAI.

Of the organizations that have deployed GenAI, only 45 percent of respondents rate its effectiveness in generating actionable insights to support better security decisions as highly effective.

The lack of in-house expertise and concerns about accuracy, bias and misinformation are the top two barriers to adopting GenAI. Twenty-five percent of respondents say their organizations have no plans to adopt GenAI and their reasons are shown in Figure 11.

Figure 11. Why would your organization not adopt GenAI?
Two responses permitted

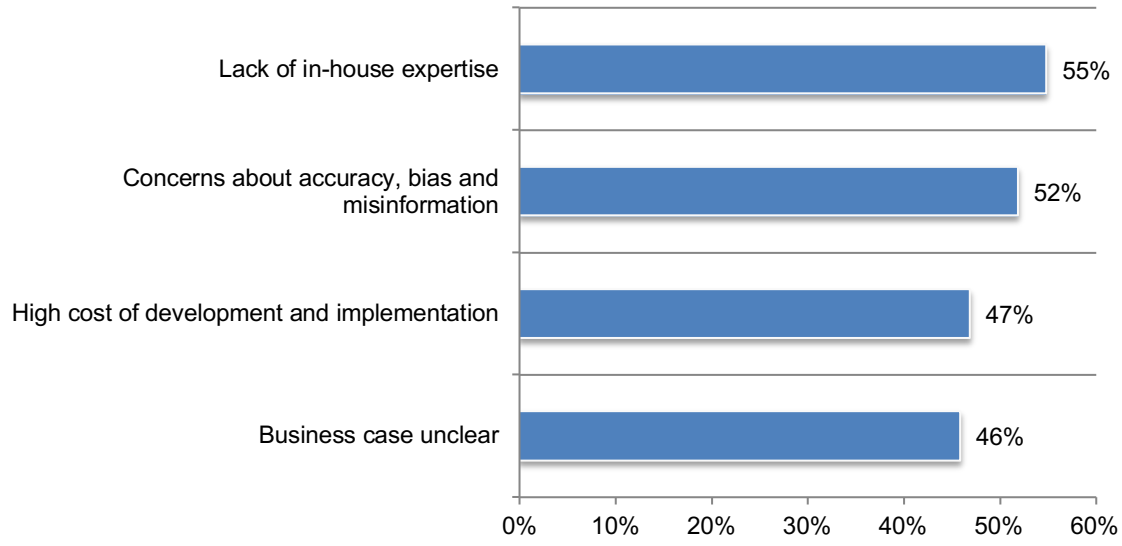

Risks created by GenAI are data misinformation and concerns about security and privacy. Only half (50 percent) of respondents say their organizations take steps to reduce security risks pertaining to the usage of GenAI. As shown in Figure 12, of these respondents, 55 percent of respondents say their organizations conduct training programs to familiarize everyone with the risks and rewards of GenAI.

Fifty percent of respondents say their organization uses experts to validate rough draft generative AI outputs. Validating generative AI output involves human-led critical review and checks for accuracy, bias, and logic, focusing on fact checking against reliable sources, assessing for bias, ensuring logical flow, and verifying source citations (or noting their absence). Key steps include cross-referencing facts, looking for diverse perspectives, asking if the information is current and relevant, and using critical thinking to spot hallucinations or inconsistencies.

Figure 12. Does your organization take any of the following steps?
More than one response permitted

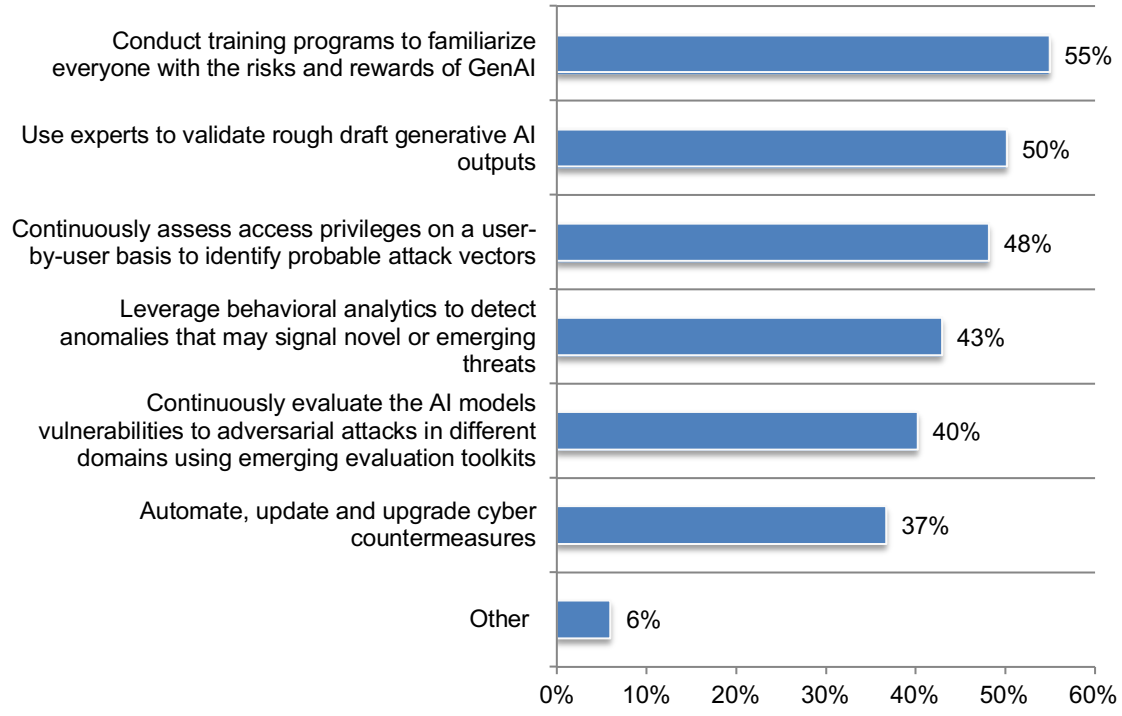

Agentic AI is a type of AI that can autonomously make decisions, take actions and learn on its own to achieve specific goals. It is characterized by autonomy, the ability to initiate and complete tasks without constant oversight and reasoning, where sophisticated decision-making is based on context.

Thirty-eight percent of respondents say their organizations have fully adopted (15 percent) or partially adopted (23 percent) Agentic AI. Respondents report an average of 23 percent of employees use AI agentic agents that autonomously perform sequences of tasks such as coding, email response or data queries.

The lack of proper risk and security controls and complex system integration are reasons not to adopt Agentic AI. Nineteen percent of organizations have no plans to adopt Agentic AI. The primary reasons are the lack of proper risk and security controls (43 percent of respondents) and complex system integration (38 percent), as shown in Figure 13.

Agentic AI risks involve new vulnerabilities from autonomous decision-making, like prompt injection, privilege escalation and unauthorized actions requiring evolved controls beyond traditional security, focusing on identity management (agent-specific), fine-grained access, continuous monitoring, human oversight, sandboxing and strict policy enforcement to secure their reasoning, memory, tools, and actions from misuse and compromise.

Figure 13. Why would your organization not adopt Agentic AI?
Two responses permitted

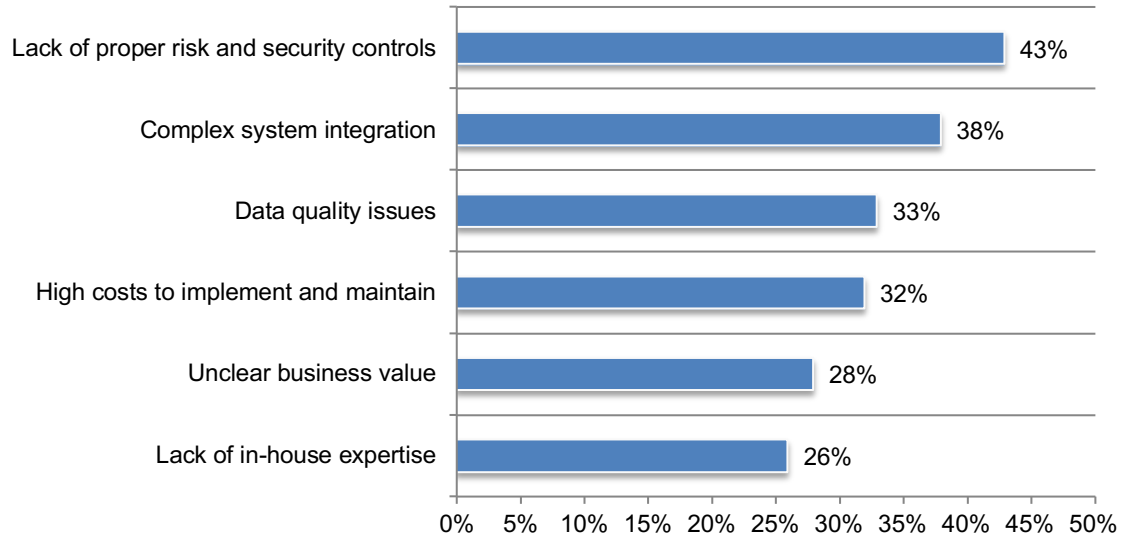

Organizations rate the effectiveness of Agentic AI in automating manual tasks and removing human error as low. Agentic AI significantly can reduce human error by automating tasks, ensuring consistency, and learning from patterns, leading to fewer mistakes in data entry, security, and other areas. However, it doesn't eliminate error entirely. Challenges to reducing human error are ensuring AI logic is sound, handling unforeseen issues, and managing the risks of autonomous decisions, often requiring human oversight.

Respondents were asked to rate the effectiveness of Agentic AI in automating manual tasks for security teams with limited human intervention from 1 = low effectiveness to 10 = high effectiveness. As shown in Figure 14, only 44 percent of respondents rate the effectiveness of automating manual tasks for security teams with limited human intervention as high (7+ on the 10-point scale). Only 39 percent of respondents say Agentic AI is highly effective in removing human error by systematically retrieving comprehensive threat intelligence,

Figure 14. Agentic AI's effectiveness in automating manual tasks and removing human error
On a scale from 1 = low effectiveness to 10 = highly effective, 7+ responses presented, 7+ responses presented

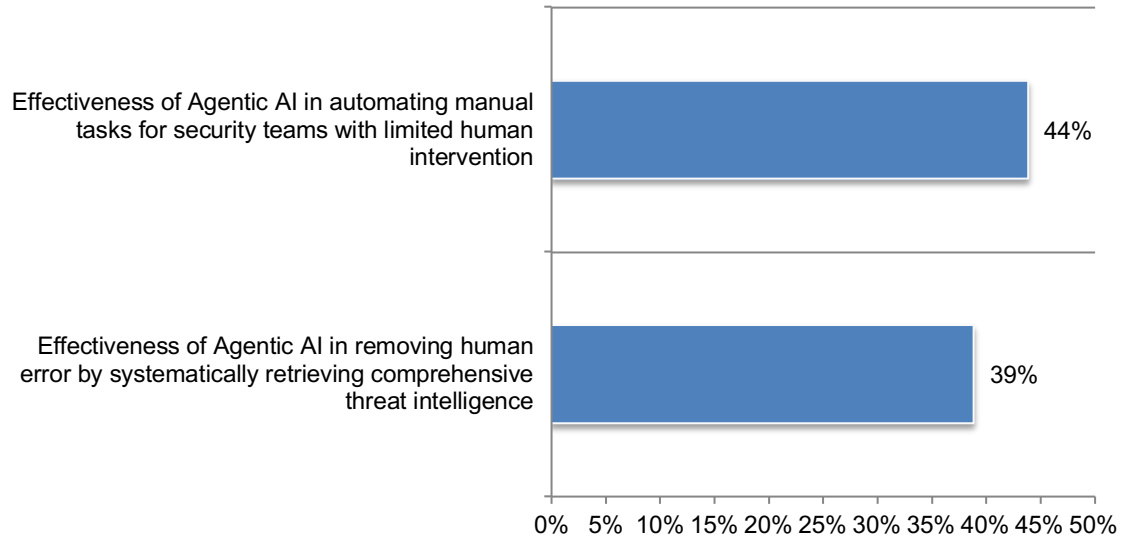

The risk of data theft caused by AI agents is high. Agentic AI can both be a powerful tool for preventing data theft by automating security tasks and a source of new vulnerabilities that require updated data protection strategies.

However, the malicious use of Agentic AI significantly increases the risk of data theft. Agentic AI systems can autonomously plan, reason, and execute complex tasks with minimal human intervention, allowing attackers to conduct large-scale, sophisticated, and rapid data exfiltration that can easily bypass traditional security measures. Fifty-five percent of respondents believe AI agents will significantly increase data theft (29 percent) or moderately increase the risk of data theft (26 percent), as shown in Figure 15.

Figure 15. Will the malicious use of AI agents increase the risk of data theft?

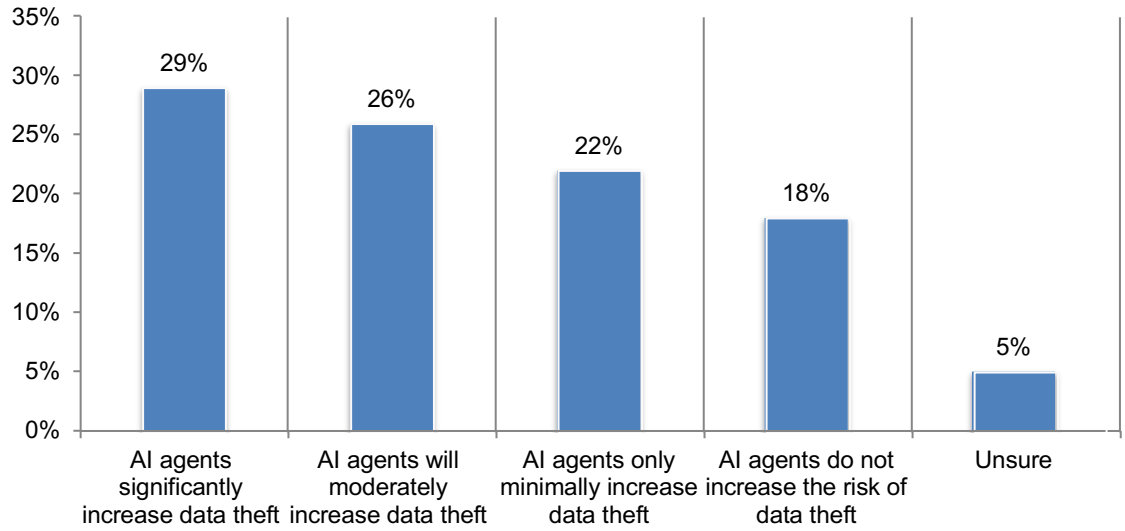

AI agents make detection of an intrusion more difficult. The malicious use of AI agents significantly complicates intrusion detection by enabling faster, stealthier attacks that mimic legitimate behavior, creating convincing deepfakes, automating complex reconnaissance, and poisoning training data. According to Figure 16, 66 percent of respondents say AI agents make detection of an intrusion more difficult (41 percent) or moderately more difficult (25 percent). Only 12 percent of respondents say AI agents **do not** make intrusion detection more difficult.

Figure 16. Will the malicious use of AI agents make detection of an intrusion more difficult?

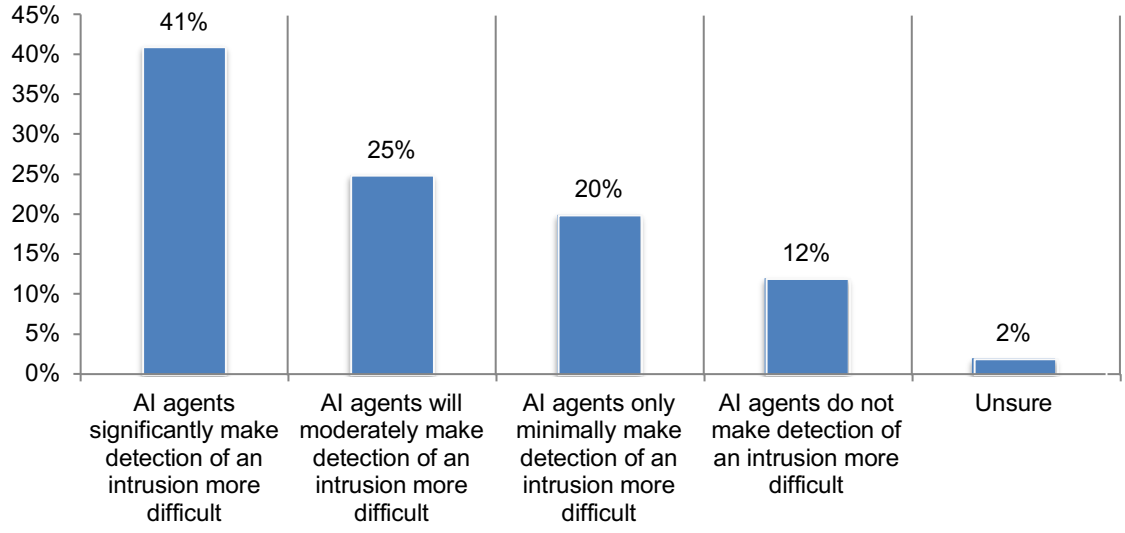

AI security and business risk perception gaps between the C-suite and those in the trenches

To understand the alignment between the C-suite and those in the trenches in minimizing AI risks, a special analysis was conducted based on the 56 percent of respondents who are at the C-suite level (supervisors and above) vs. 44 percent of respondents who are staff/technicians.

Following are some of the interesting differences between the C-suite and technician level positions.

C-level respondents have a more positive view of the effectiveness of their organization's security posture. Fifty-five percent of C-level respondents vs. 49 percent of staff/technician respondents say their organization's IT security posture is high or highly effective in mitigating risks, vulnerabilities and attacks across the enterprise.

Fifty-one percent of C-level respondents say their organization had only one cybersecurity incident in the past 12 months. In contrast, 60 percent of staff/technician positions say their organizations had **more than one** cybersecurity incident in the past 12 months.

Figure 17. Has your organization experienced one or more security incidents, including cyberattacks?

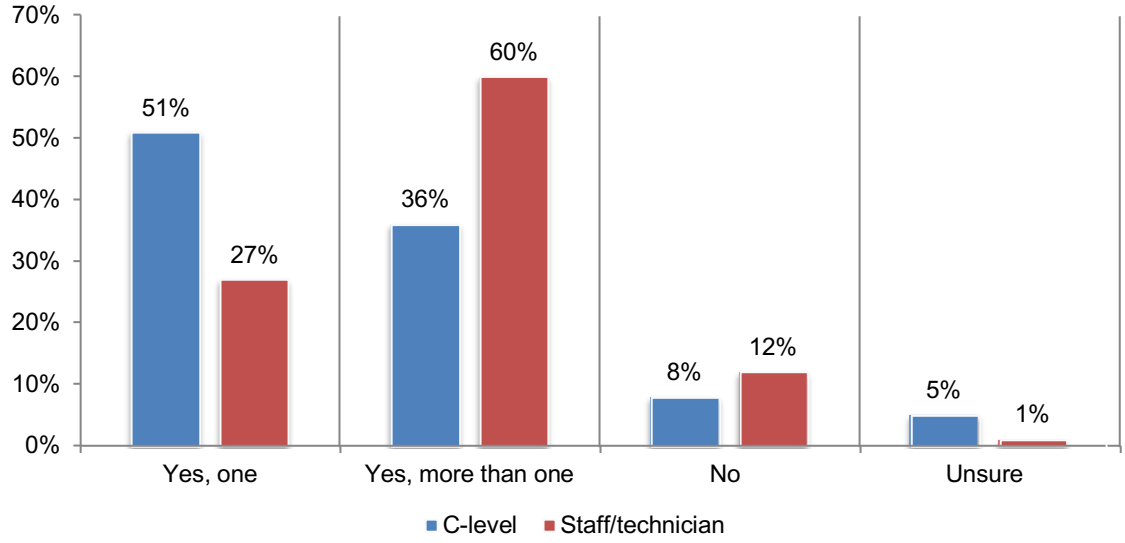

Forty-three percent of both C-level and technician respondents say their organizations take a risk-based approach to its AI strategies and initiatives. As shown in Figure 18, of these respondents, the biggest gap is that 81 percent of the C-level vs. 41 percent of technicians say as part of the risk-based approach their organizations create policies and procedures for AI usage, ethical considerations and compliance with regulations.

Sixty-five percent of C-level vs. 53 percent of technicians say their organizations identify potential risks such as bias in training data or model drift and assess their likelihood and impact. Another significant gap is that 61 percent of C-level respondents vs. only 39 percent of technicians restrict access to sensitive datasets used in AI training such as through granular access controls.

Figure 18. What steps does your organization take to minimize AI risks?

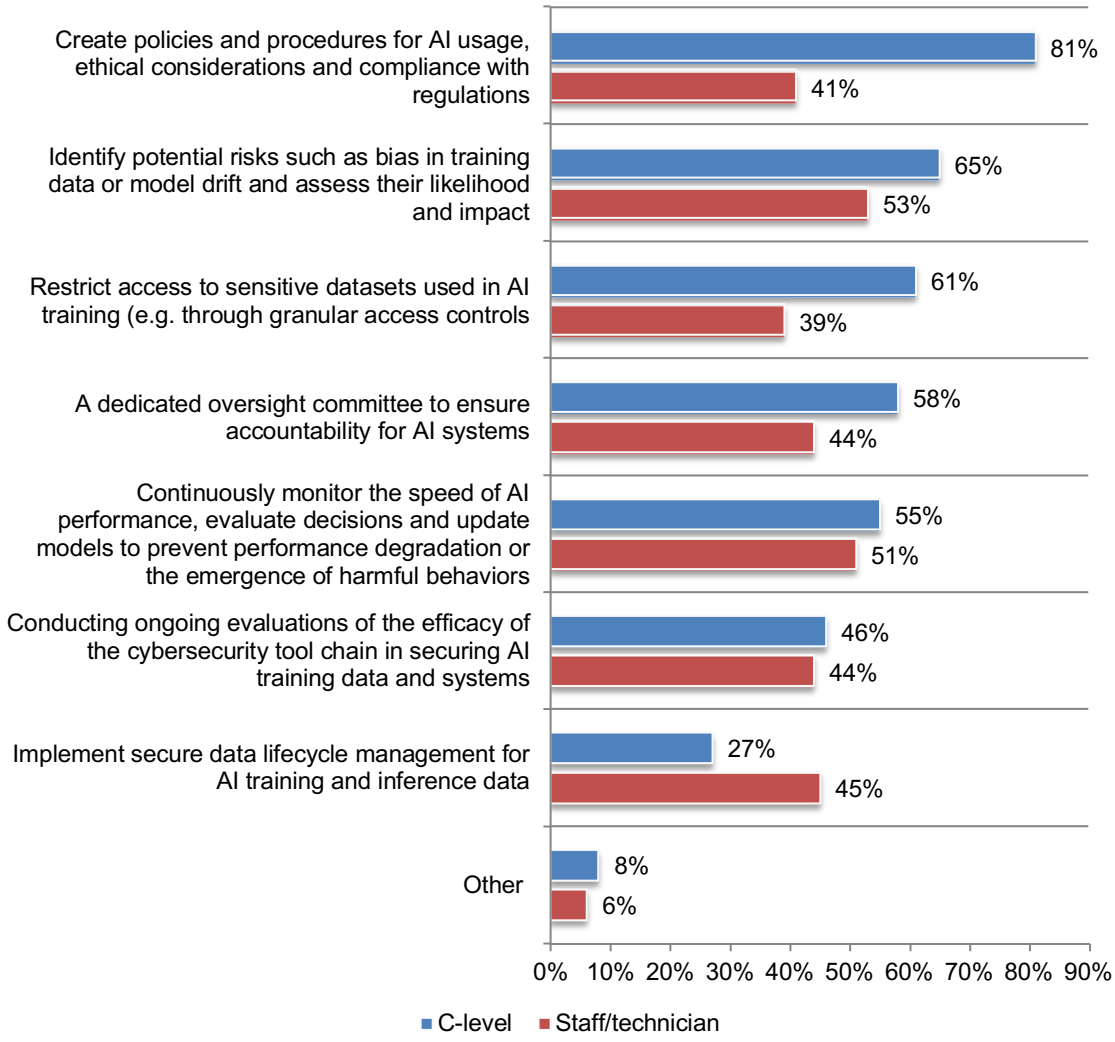

More C-level respondents say their organizations have tools or processes to investigate and respond to AI-generated or AI-assisted threats. Sixty-five percent of C-level respondents vs. 35 percent of staff/technician respondents say their organizations have these tools or processes in place to respond to the misuse of tools that result in data leakage and prompt injection model.

Prompt injection is a critical security vulnerability in Large Language Models (LLMs) where malicious, crafted inputs trick the model into ignoring its original developer instructions and executing unauthorized commands. By exploiting the inability of LLMs to distinguish between instructions and data, attackers can bypass safeguards, leak sensitive data, or manipulate outputs.

Figure 19. Does your organization have tools or processes in place to investigate and respond to AI-generated or AI-assisted threats such as the misuse of tools that result in data leakage and prompt injection model?

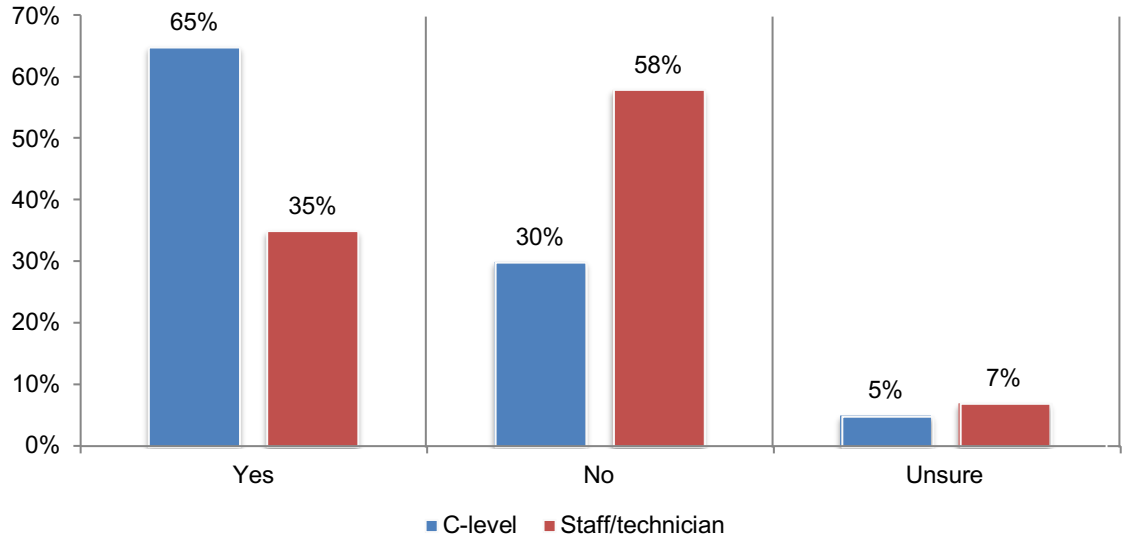

More C-level respondents are positive about the future ability of AI systems to reason and make autonomous decisions to avoid misuse. Fifty-six percent of C-level positions (24 percent + 32 percent) vs. 40 percent (16 percent + 24 percent of staff/ technicians say in the future it will be possible to have AI systems reason and make autonomous decisions based on ethics, regulations, and law to avoid misuse.

Figure 20. In the future, it will be possible to have AI systems reason and make autonomous decisions based on ethics, regulations, and law to avoid misuse.
Strongly agree and agree responses combined

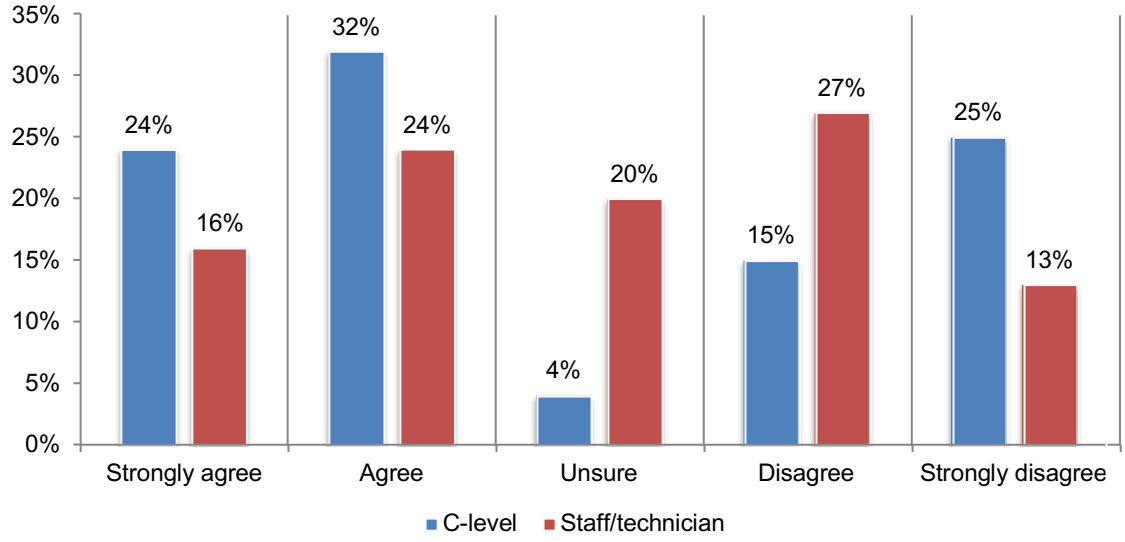

Staff/technicians who are in the trenches believe human oversight in AI governance is needed. Fifty-five percent of staff/technician respondents (25 percent + 30 percent) vs. 46 percent of C-level respondents (16 percent + 30 percent) say human oversight is needed in AI governance because of the speed in which attackers can adapt.

Figure 21. Human oversight is needed in AI governance because of the speed in which attackers can adapt
Strongly agree and agree responses combined

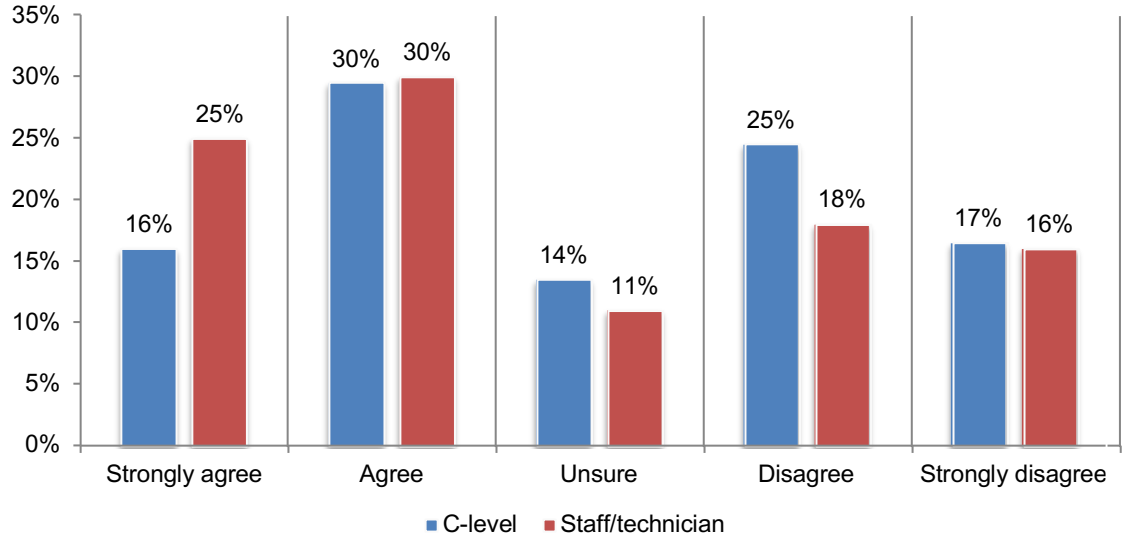

Regional differences

In this section, the findings from the following four regions are presented: North America (611 respondents), EMEA (496 respondents), Asia Pac (428 respondents) and LATAM (343 respondents).

Organizations in North America are most likely to take a risk-based approach to AI strategies. According to Figure 22, 51 percent of respondents in North America and 46 percent of respondents in EMEA have adopted a risk-based approach. Only 37 percent of organizations in Asia Pac have a risk-based approach.

Figure 22. Does your organization take a risk-based approach to its AI strategies and initiatives?

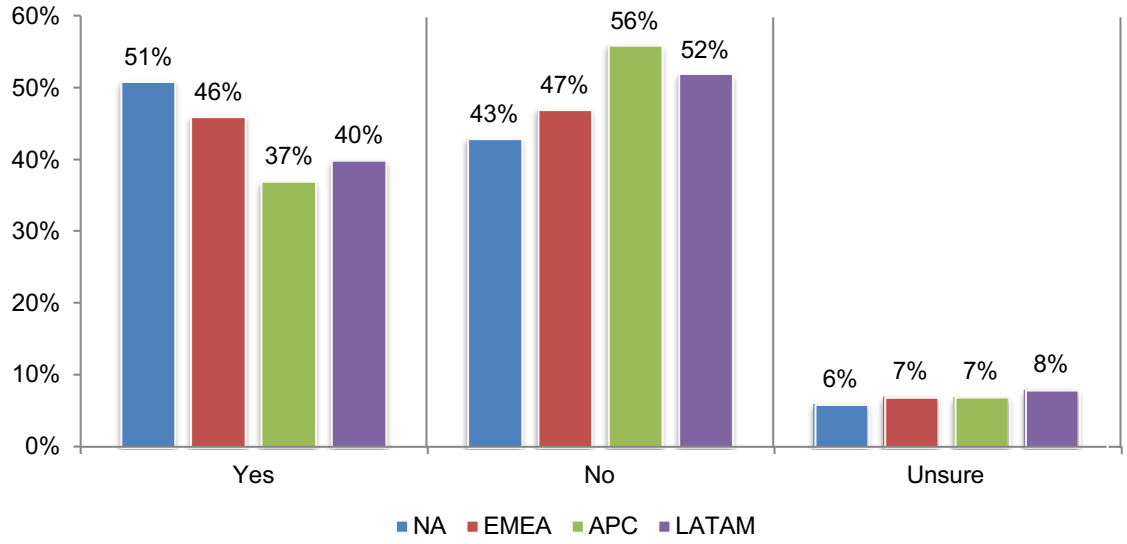

Most regions have adopted the NIST AI risk management framework, the ISO framework for AI systems or both. According to Figure 23, 31 percent of North American organizations have adopted the NIST AI risk management framework and 29 percent say they use the ISO framework. Asia Pac organizations are most likely not to adopt either framework (28 percent of respondents).

Figure 23. Does your organization use the NIST AI risk management framework and/or ISO framework for AI systems using machine learning to design and deploy trusted AI applications?

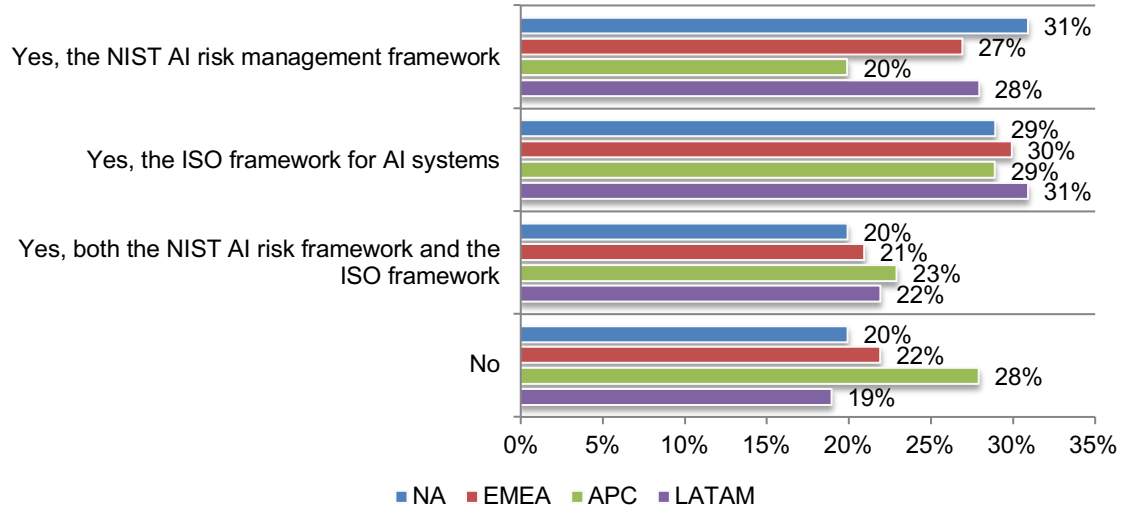

To respond to AI-assisted threats, organizations use a combination of AI-powered cybersecurity platforms that provide autonomous detection and response, behavioral analytics, and real-time threat intelligence. These tools process massive amounts of data faster than humans to identify and neutralize sophisticated attacks.

According to Figure 24, North American and Asia Pac organizations are most likely to have tools or processes in place to investigate and respond to AI-generated or AI-assisted threats, 57 percent and 51 percent, respectively.

Figure 24. Does your organization have tools or processes in place to investigate and respond to AI generated or AI-assisted threats?

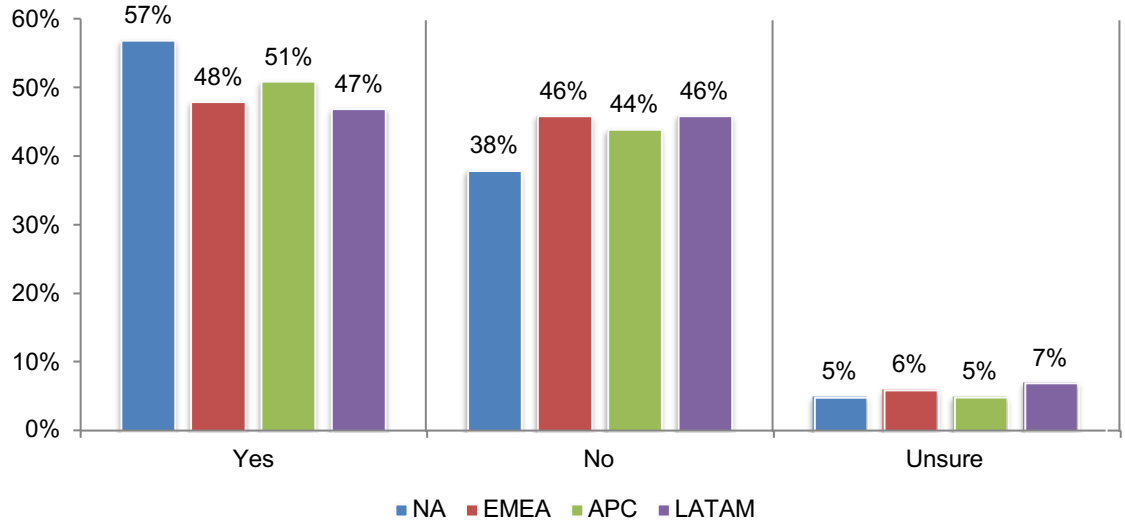

North American and Asia Pac organizations are most likely to rate the difficulty in safeguarding confidential data used in AI as very or extremely difficult, as shown in Figure 25.

Figure 25. The percentage of organizations rating the difficulty of safeguarding confidential data used in their organizations as very or extremely difficult
On a scale from 1 = low difficulty to 10 = extremely difficult, 7+ responses presented

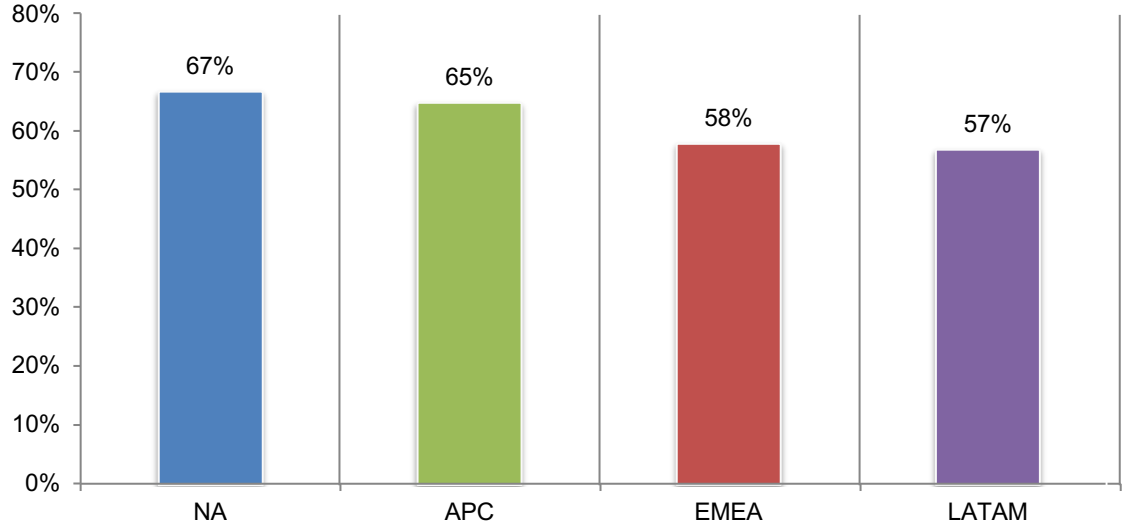

Lack of in-house expertise, concerns about errors and high cost are barriers to adopting GenAI across the regions. The percentage of organizations adopting GenAI in each region are: North America (56 percent), EMEA (53 percent) Asia Pac (53 percent) and LATAM (46 percent).

As shown in Figure 26, in North America and EMEA the biggest concerns are lack of in-house expertise and concerns about accuracy, bias and misinformation (54 percent, 56 percent and 56 percent, respectively). In Asia Pac reasons for not adopting are the lack of in-house expertise and an unclear business case (60 percent and 52 percent, respectively). LATAM will not adopt because of concerns about accuracy, bias and misinformation, and the high cost of implementation (51 percent and 54 percent, respectively).

Figure 26. Why would your organization not adopt GenAI?
Two responses permitted

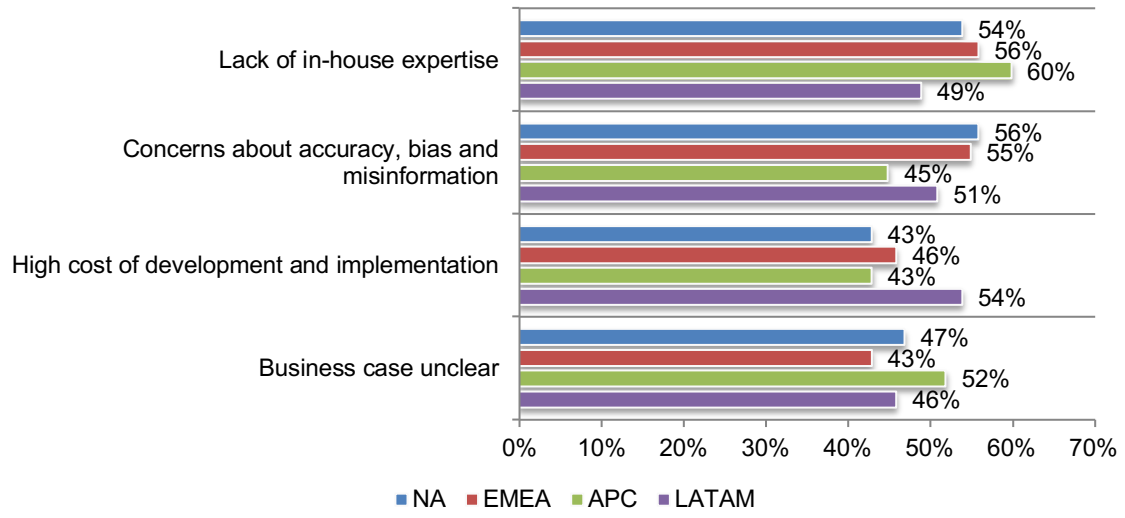

The risk of not having security control, data quality issues, and complexity of system integration are barriers to adopting Agentic AI. The percentage of organizations adopting Agentic AI in each region are: North America (42 percent), EMEA (39 percent) Asia Pac (39 percent) and LATAM (31 percent).

Organizations in North America are not adopting Agentic AI because of the lack of proper risk and security controls and data quality issues (45 percent and 36 percent, respectively). EMEA organizations and Asia Pac organizations are most concerned about the lack of proper risk and security controls and complex system integration (44 percent, 40 percent, 43 percent and 41 percent, respectively), as shown in Figure 27.

LATAM organizations' concerns are lack of proper risk and security controls and lack of in-house expertise (41 percent and 35 percent, respectively).

Figure 27. Why would your organization not adopt Agentic AI?
Two responses permitted

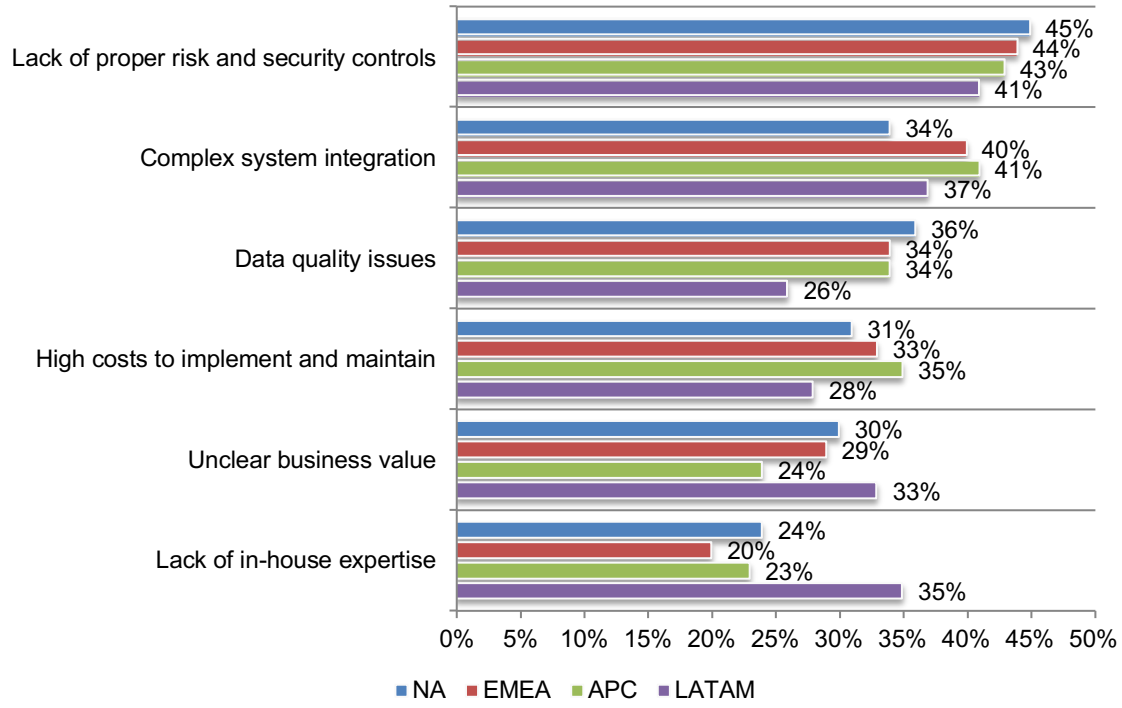

Part 3. Methodology

A sampling frame of 56,779 IT and IT security practitioners in North America, AsiaPac, EMEA, and LATAM who are knowledgeable about their organizations' use of AI for cybersecurity and business purposes were selected as participants to this survey. Table 1 shows 2,140 total returns. Screening and reliability checks required the removal of 262 surveys. Our final sample consisted of 1,878 surveys (NA 611, APAC 428, EMEA 496, LATAM 343) or a 3.2 percent response rate.

Table 1. Sample response	Freq	Pct%
Sampling frame	56,779	100.0%
Total returns	2,140	3.8%
Rejected or screened surveys	262	0.5%
Final sample	1,878	3.3%

Pie chart 1 reports the respondent's organizational level within participating organizations. More than half (56 percent) of respondents are at or above the supervisor level. The largest category at 22 percent of respondents is staff/technician.

Pie chart 1. Current position within the organization

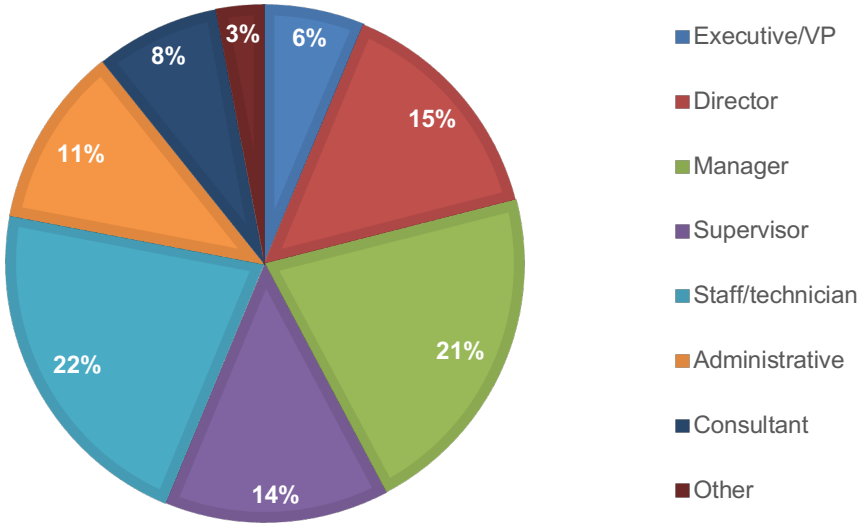

As shown in Pie chart 2, 20 percent of respondents are within the IT security/InfoSec department, 19 percent of respondents are in the engineering department and 16 percent of respondents are in IT Operations and another 16 percent are in infrastructure.

Pie chart 2. Respondents' department or team

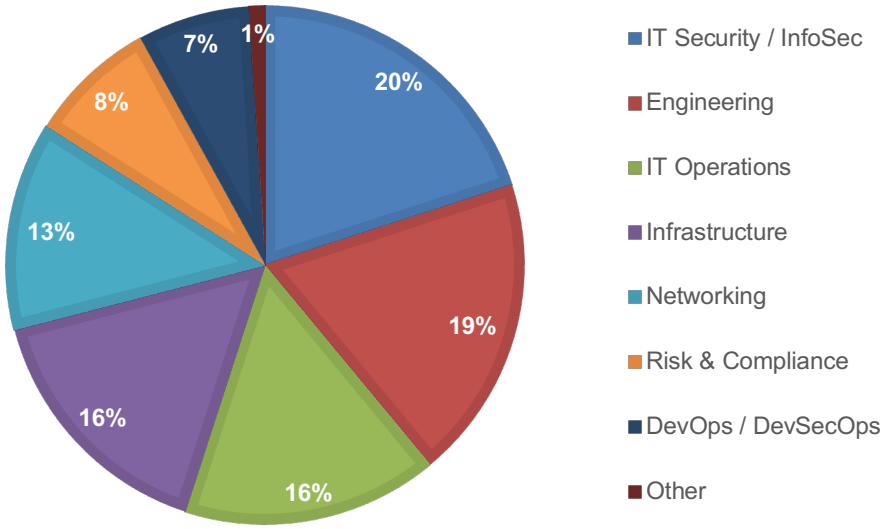

As shown in Pie chart 3, more than half (53 percent) of respondents are from organizations with a headcount of more than 5,000 employees

Pie chart 3. Worldwide headcount

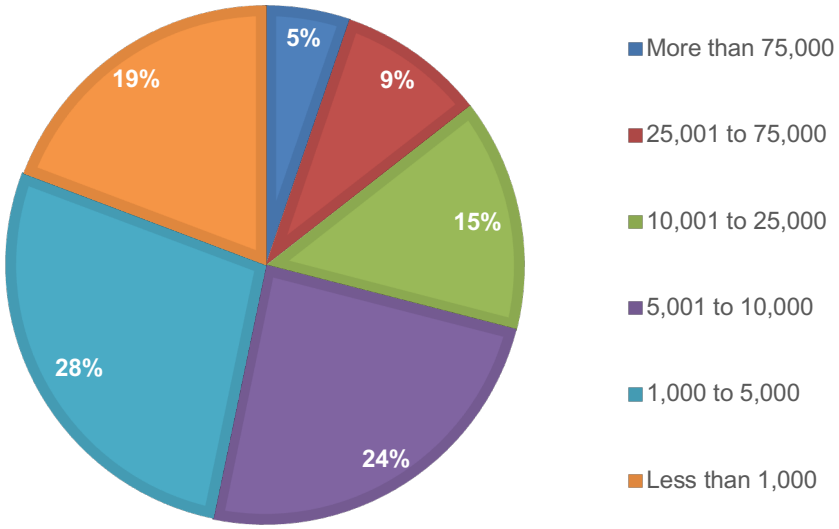

Pie chart 4 reports the industry classification of respondents' organizations. This chart identifies financial services (15 percent of respondents) as the largest industry focus, which includes banking, investment management, insurance, brokerage, payments, and credit cards. This is followed by services (12 percent), industrial and manufacturing (11 percent), technology and software (11 percent), and energy and utilities (10 percent).

Pie chart 4. Primary industry classification

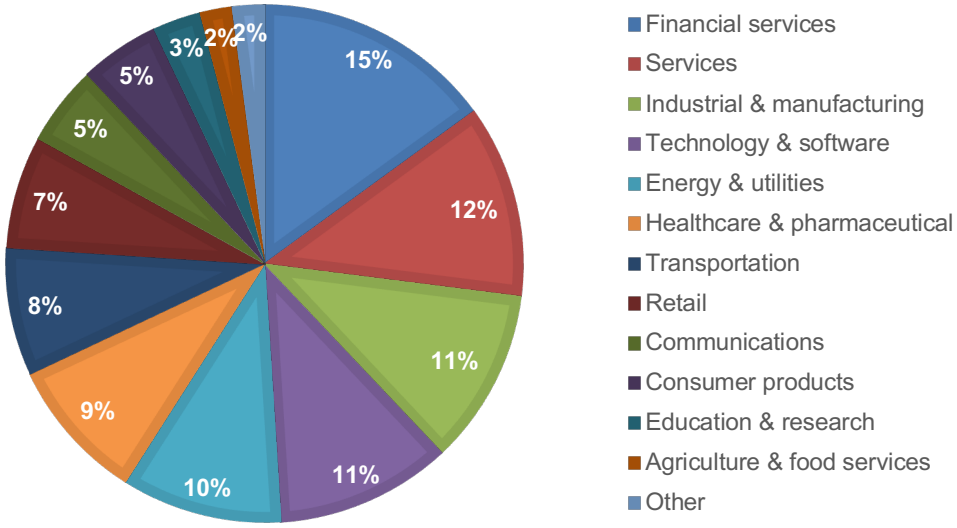

Part 4. Caveats to this study

There are inherent limitations to survey research that need to be carefully considered before drawing inferences from findings. The following items are specific limitations that are germane to most web-based surveys.

- Non-response bias: The current findings are based on a sample of survey returns. We sent surveys to a representative sample of individuals, resulting in a large number of usable returned responses. Despite non-response tests, it is always possible that individuals who did not participate are substantially different in terms of underlying beliefs from those who completed the instrument.

- Sampling-frame bias: The accuracy is based on contact information and the degree to which the list is representative of individuals who are IT and IT security practitioners. We also acknowledge that the results may be biased by external events such as media coverage. Finally, because we used a web-based collection method, it is possible that non-web responses by mailed survey or telephone call would result in a different pattern of findings.

- Self-reported results: The quality of survey research is based on the integrity of confidential responses received from subjects. While certain checks and balances can be incorporated into the survey process, there is always the possibility that a subject did not provide accurate responses.

Appendix: Detailed Survey Results

The following tables provide the frequency or percentage frequency of responses to survey questions. All survey responses were captured in November 2025.

Survey response	Global
Sampling frame	56,779
Total returns	2,140
Rejected surveys	262
Final sample	1,878

Screening questions

S1. Has your organization adopted AI? Please select one choice only.	Global
Yes, AI is fully adopted	22%
Yes, AI is partially adopted or plan to adopt	60%
We have no plans to adopt AI at this time (stop)	18%
Total	100%

S2. How familiar are you with your organization's use of AI for cybersecurity and business purposes?	Global
Very familiar	20%
Familiar	37%
Somewhat familiar	20%
No knowledge (stop)	23%
Total	100%

Part 1. Background on organization's security posture and AI adoption

Q1. How would you describe your organization's IT security posture in terms of its effectiveness at mitigating risks, vulnerabilities and attacks across the enterprise on a scale from 1 = not effective to 10 = very effective?	Global
1 or 2	14%
3 or 4	18%
5 or 6	16%
7 or 8	22%
9 or 10	30%
Total	100%

Q2. Has your organization experienced one or more security incidents, including cyberattacks, in the past 12 months?	Global
Yes, one	39%
Yes, more than one	48%
No	10%
Unsure	3%
Total	100%

Q3. If yes, what best describes the type of security incidents, including cyberattacks, experienced by your organization? Please select all that apply.	Global
Advanced malware	21%
Zero-day attacks	28%
AI-generated attacks	27%
APIs	16%
Cloud malware injection attacks	17%
Cloud cryptomining	21%
Phishing / social engineering	40%
Denial of service	33%
Account takeover	17%
Credential theft	21%
Ransomware	34%
Web application attack	30%
Web-based attack	22%
Compromised / stolen devices	26%
Malicious insider	26%
Negligent insider	20%
Other (please specify)	4%
Total	402%

Q4. Which of the following functions are involved in AI governance and policies in your organization? Please check all that apply.	Global
Chief Executive Officer	7%
Chief Technology Officer	15%
Chief Information Security Officer	21%
Chief Data Officer	17%
Head of data science	15%
IT security director/manager	17%
Chief Information Officer	32%
Security architect	15%
Security analyst	14%
Chief Compliance and Privacy Officer	29%
Chief Legal Officer	23%
Chief Financial Officer	6%
Internal Audit	21%
Chief Risk Officer	20%
Total	250%

Q5. Does your organization **have or will it have** a Chief AI Officer to ensure accountability for AI systems?	Global
Yes	34%
No	46%
Unsure	20%
Total	100%

Q6. What best describes the maturity of your organization's use of AI? Please select **one choice** only.	Global
Early adoption stage – We have defined what the organization's cybersecurity AI strategy is, investments are planned and partially deployed (please skip to Q8)	42%
Full adoption stage – AI technologies are mostly deployed. The program has C-level support and adequate budget. (please skip to Q8)	37%
Mature stage – AI in cybersecurity activities is fully deployed and security risks assessed Effectiveness of AI is measured with KPIs and C-level executives are regularly informed about AI's ability to prevent and reduce cyberattacks.	21%
Total	100%

Q7. Please select the **top three** organizational or governance challenges that your organization experienced when deploying AI-based security technologies within your organization.	Global
It requires too much staff to implement and maintain AI-based technologies	50%
There is not enough time to integrate AI-based technologies into security workflows	44%
We can't recruit personnel experienced in AI-based technologies	26%
We don't have the internal expertise to validate vendors' claims	32%
There is insufficient budget for AI-based technologies	46%
There is insufficient supervision and oversight of AI learning and decision-making	39%
IT and IT security functions are not aligned on the organization's AI strategy	30%
The business case is not clear	26%
Other (please specify)	7%
Total	300%

Q8. Following deployment of AI, which of the following are barriers to the effectiveness of AI-based security technologies used by your organization today? Please select the **top three** factors.	Global
AI tools/technology we need are not available	31%
We can't apply AI-based controls that span across the entire enterprise	34%
We can't create a unified view of AI users across the enterprise	39%
There are errors and inaccuracies in AI decision rules	45%
There are errors and inaccuracies in data inputs ingested by AI technology (engine)	40%
There is a heavy reliance on legacy IT environments	38%
There are Interoperability issues among AI technologies	37%
There is a lack of mature and/or stable AI technologies	30%
Other (please specify)	6%
Total	300%

Q9. Does your organization take a risk-based approach to its AI strategies and initiatives?	Global
Yes	43%
No	50%
Unsure	7%
Total	100%

Q10. If yes, what steps does your organization take to minimize AI risks? Please select all that apply.	Global
Create policies and procedures for AI usage, ethical considerations and compliance with regulations	61%
A dedicated oversight committee to ensure accountability for AI systems	51%
Identify potential risks such as bias in training data or model drift and assess their likelihood and impact	59%
Continuously monitor the speed of AI performance, evaluate decisions and update models to prevent performance degradation or the emergence of harmful behaviors	53%
Implement secure data lifecycle management for AI training and inference data	36%
Restrict access to sensitive datasets used in AI training (e.g. through granular access controls	50%
Conducting ongoing evaluations of the efficacy of the cybersecurity tool chain in securing AI training data and systems	45%
Other (please specify)	7%
Total	362%

Q11. Does your organization use the NIST AI risk management framework and/or ISO framework for AI systems using machine learning to design and deploy trusted AI applications? Please select **one choice** only.	Global
Yes, the NIST AI risk management framework	27%
Yes, the ISO framework for AI systems	30%
Yes, both the NIST AI risk framework and the ISO framework	22%
No	21%
Total	100%

Q12. Does your organization have tools or processes in place to investigate and respond to AI-generated or AI-assisted threats such as the misuse of tools that result in data leakage and prompt injection model manipulation?	Global
Yes	50%
No	44%
Unsure	6%
Total	100%

Q13. Using the following 10-point scale, please rate your organization's focus on the governance of AI systems from 1 = low focus to 10 = very high focus.	Global
1 or 2	10%
3 or 4	16%
5 or 6	21%
7 or 8	28%
9 or 10	25%
Total	100%

Q14. Using the following 10-point scale, please rate the priority of having well-established processes, accountability structures or tools to manage non-human identities from 1 = low priority to 10 = very high priority	Global
1 or 2	11%
3 or 4	16%
5 or 6	22%
7 or 8	25%
9 or 10	26%
Total	100%

Q15. Using the following 10-point scale, please rate how important explainability is in AI-powered solutions from 1 = low importance to 10 = highly important	Global
1 or 2	12%
3 or 4	14%
5 or 6	27%
7 or 8	32%
9 or 10	15%
Total	100%

Q16. Using the following 10-point scale, please rate the effectiveness of your organization's use of AI in threat detection and hunting **to reduce the time to detect anomalies, new patterns, and emerging threats that might otherwise go undetected** from 1 = low effectiveness to 10 = high effectiveness	Global
1 or 2	11%
3 or 4	20%
5 or 6	18%
7 or 8	28%
9 or 10	23%
Total	100%

Q17. Using the following 10-point scale, please rate the effectiveness of your organization's use of AI in threat detection and hunting to **provide deeper contextual insights and reducing the manual workload** from 1 = low effectiveness to 10 = high effectiveness	Global
1 or 2	13%
3 or 4	18%
5 or 6	21%
7 or 8	30%
9 or 10	18%
Total	100%

The following questions focus on the difficulty of minimizing specific AI risks.

Q18. Using the following 10-point scale, please rate the difficulty in minimizing such data risks as error propagation and misleading and harmful content caused by low-quality data used to train generative AI models from 1 = low difficulty to 10 = extremely difficult.	Global
1 or 2	12%
3 or 4	14%
5 or 6	13%
7 or 8	27%
9 or 10	34%
Total	100%

Q19. Using the following 10-point scale, please rate the difficulty in minimizing such prompt or input risks such as misleading, inaccurate or harmful responses due to unsophisticated prompts or questions being provided to the AI model from 1 = low difficulty to 10 = extremely difficult.	Global
1 or 2	10%
3 or 4	13%
5 or 6	19%
7 or 8	28%
9 or 10	30%
Total	100%

Q20. Using the following 10-point scale, please rate the difficulty in minimizing user risks such as the unintended consequences due to users becoming unwitting parties to the creation of misinformation and other harmful content from 1 = low difficulty to 10 = extremely difficult.	Global
1 or 2	15%
3 or 4	17%
5 or 6	12%
7 or 8	24%
9 or 10	32%
Total	100%

Q21. Using the following 10-point scale, please rate the difficulty in minimizing model and bias risks such as the breach of ethical and responsible AI principles in the language model development leading to discriminatory or unfair outputs from 1 = low difficulty to 10 = extremely difficult.	Global
1 or 2	11%
3 or 4	14%
5 or 6	13%
7 or 8	32%
9 or 10	30%
Total	100%

Part 2 Privacy, security and ethical considerations

Q22. On the 10-point scale, how difficult is it to safeguard confidential and personal data used in your organization's AI from 1 = not difficult to 10 = highly difficult?	Global
1 or 2	12%
3 or 4	13%
5 or 6	14%
7 or 8	27%
9 or 10	34%
Total	100%

Q23. On the 10-point scale, does the use of AI make it difficult to comply with privacy and security regulations and mandates from 1 = not difficult to 10 = highly difficult?	Global
1 or 2	11%
3 or 4	12%
5 or 6	18%
7 or 8	29%
9 or 10	30%
Total	100%

Q24a. Does your organization have data privacy policies specifically for the use of AI?	Global
Yes	41%
No	54%
Unsure	5%
Total	100%

Q24b. If yes, what is included in the privacy policies? Please select all that apply.	Global
Data stewardship requirements to mitigate uses of personal information that are adverse or unfair to individuals the data relates to	59%
Establish data transparency and/or disclosure rules and the rights of individuals to access information relating to them	49%
Appoint a privacy officer to oversee the governance of AI privacy policies	44%
Conduct regular privacy impact assessments	59%
Work with vendors to ensure “privacy by design” in the use of AI technologies	42%
Apply differential privacy or other privacy-preserving techniques to AI models	44%
Define protocols for securing AI training and inference data throughout its lifecycle such as mandatory encryption and access control policies for sensitive data	58%
Other (please specify)	5%
Total	360%

Q25. In the future, it will be possible to have AI systems reason and make autonomous decisions based on ethics, regulations and law to avoid misuse.	Global
Strongly agree	20%
Agree	28%
Unsure	12%
Disagree	21%
Strongly disagree	19%
Total	100%

Q26. Our AI models can learn robust norms and make safe decisions autonomously.	Global
Strongly agree	20%
Agree	27%
Unsure	12%
Disagree	21%
Strongly disagree	20%
Total	100%

Q27. Human oversight is needed in AI governance because of the speed in which attackers can adapt.	Global
Strongly agree	21%
Agree	30%
Unsure	12%
Disagree	21%
Strongly disagree	16%
Total	100%

Part 3. Adoption of GenAI

Q28a. Has your organization adopted GenAI?	Global
Yes, fully deployed	18%
Yes, partially deployed	34%
Will adopt in the future (please skip to 32a)	23%
No plans to adopt	25%
Total	100%

Q28b. If no plans to adopt, why? Please select the **top two** reasons only.	Global
Lack of in-house expertise	55%
High cost of development and implementation	47%
Concerns about accuracy, bias and misinformation	52%
Business case unclear	46%
Total	200%

Please skip to Q32a

Q29. How much of your organization's work tasks have been automated by GenAI?	Global
Less than 10 percent	20%
10 percent to 25 percent	30%
26 percent to 50 percent	20%
More than 50 percent	19%
Unsure	11%
Total	100%
Extrapolated average	25%

Q30a. Does your organization take steps to reduce security risks pertaining to the usage of GenAI?	Global
Yes	50%
No (please skip to Q32a)	50%
Total	100%

Q30b. If yes, does your organization take any of the following steps? Please select all that apply.	Global
Conduct training programs to familiarize everyone with the risks and rewards of GenAI	55%
Use experts to validate rough draft generative AI outputs	50%
Automate, update and upgrade cyber countermeasures	37%
Continuously assess access privileges on a user-by-user basis to identify probable attack vectors	48%
Leverage behavioral analytics to detect anomalies that may signal novel or emerging threats	43%
Continuously evaluate the AI models vulnerabilities to adversarial attacks in different domains using emerging evaluation toolkits	40%
Other (please specify)	6%
Total	278%

Q31. Using the following 10-point scale, please rate GenAI's effectiveness in generating actionable insights to support better security decisions from 1 = low effectiveness to 10 = highly effective	Global
1 or 2	11%
3 or 4	21%
5 or 6	25%
7 or 8	22%
9 or 10	21%
Total	100%

Part 4. Adoption of Agentic AI

Q32a. Has your organization adopted Agentic AI? Please select one choice only.	Global
Yes, fully adopted	15%
Yes, partially adopted	23%
Will adopt in the future (please skip to Part 5)	43%
No plans to adopt	19%
Total	100%

Q32b. If no plans to adopt Agentic AI, why? Please select the **top two** choices only.	Global
High costs to implement and maintain	32%
Unclear business value	28%
Complex system integration	38%
Data quality issues	33%
Lack of proper risk and security controls	43%
Lack of in-house expertise	26%
Total	200%

Please skip to Part 5

Q33. What percentage of your employees use AI agentic agents that autonomously perform sequences of tasks such as coding, email response or data queries?	Global
Less than 5 percent	21%
5 to 10 percent	23%
11 to 20 percent	17%
21 to 50 percent	19%
More than 50 percent	20%
Total	100%
Extrapolated average	23%

Q34. Please rate the effectiveness of Agentic AI in automating manual tasks for security teams with limited human intervention from 1 = low effectiveness to 10 = highly effective.	Global
1 or 2	11%
3 or 4	19%
5 or 6	26%
7 or 8	23%
9 or 10	21%
Total	100%

Q35. Please rate the effectiveness of Agentic AI in removing human error by systematically retrieving comprehensive threat intelligence from 1 = low effectiveness to 10 = highly effective.	Global
1 or 2	18%
3 or 4	21%
5 or 6	22%
7 or 8	21%
9 or 10	18%
Total	100%

Q36. To what degree will the malicious use of AI agents increase the risk of data theft? Please select **one** choice only.	Global
AI agents significantly increase data theft	29%
AI agents will moderately increase data theft	26%
AI agents only minimally increase data theft	22%
AI agents do not increase the risk of data theft	18%
Unsure	5%
Total	100%

Q37. To what degree will the malicious use of AI agents make detection of an intrusion more difficult? Please select **one** choice only.	Global
AI agents significantly make detection of an intrusion more difficult	41%
AI agents will moderately make detection of an intrusion more difficult	25%
AI agents only minimally make detection of an intrusion more difficult	20%
AI agents do not make detection of an intrusion more difficult	12%
Unsure	2%
Total	100%

Part 5. Organization and respondent's demographics

D1. What best describes your position level within the organization?	Global
Executive/VP	6%
Director	15%
Manager	21%
Supervisor	14%
Staff/technician	22%
Administrative	11%
Consultant	8%
Other (please specify)	3%
Total	100%

D2. What best describes your department or team?	Global
IT Security / InfoSec	20%
IT Operations	16%
Infrastructure	16%
Engineering	19%
Networking	13%
Risk & Compliance	8%
DevOps / DevSecOps	7%
Other (please specify)	1%
Total	100%

D3. What range best describes the full-time headcount of your global organization?	Global
Less than 1,000	19%
1,000 to 5,000	28%
5,001 to 10,000	24%
10,001 to 25,000	15%
25,001 to 75,000	9%
More than 75,000	5%
Total	100%

D4. What best describes your organization's primary industry classification?	Global
Agriculture & food services	2%
Communications	5%
Consumer products	5%
Education & research	3%
Energy & utilities	10%
Financial services	15%
Healthcare & pharmaceutical	9%
Industrial & manufacturing	11%
Retail	7%
Services	12%
Technology & software	11%
Transportation	8%
Other (please specify)	2%
Total	100%

For more information about this study, please contact Ponemon Institute by sending an email to research@ponemon.org**.**

Ponemon Institute
Advancing Responsible Information Management

Ponemon Institute is dedicated to independent research and education that advances responsible information and privacy management practices within business and government. Our mission is to conduct high quality, empirical studies on critical issues affecting the management and security of sensitive information about people and organizations.

We uphold strict data confidentiality, privacy and ethical research standards. We do not collect any personally identifiable information from individuals (or company identifiable information in our business research). Furthermore, we have strict quality standards to ensure that subjects are not asked extraneous, irrelevant or improper questions.