Enterprise Strategy Group™ by TechTarget

Generative AI for Cybersecurity: An Optimistic but Uncertain Future

Jon Oltsik | Distinguished Analyst and Fellow
ENTERPRISE STRATEGY GROUP

JANUARY 2024

Research Objectives

Since the introduction of ChatGPT in November 2022, generative AI (GenAI) has been described as everything from a novelty and an economic boon to a threat to humanity. As this debate continued, GenAI took center stage at the RSA Conference 2023 with the introduction of and subsequent hoopla around Microsoft Security Copilot. Many other vendors have introduced similar capabilities since. Few would argue against the idea that GenAI (and AI in general) will have a profound impact on society and global economics, but in the near term, it introduces new risks as employees connect to GenAI applications, share data, and build homegrown large language models (LLMs) of their own. These actions will inevitably expand the attack surface, open new threat vectors, introduce software vulnerabilities, and lead to data leakage.

Despite these risks, generative AI holds great cybersecurity potential. Generative AI could help improve security team productivity, accelerate threat detection, automate remediation actions, and guide incident response. These prospective benefits are so compelling that many CISOs are already experimenting with GenAI or building their own security LLMs. At the same time, security professionals remain anxious about how cybercriminals may use GenAI as part of attack campaigns, and how they can defend against these advances.

Have organizations embraced GenAI for cybersecurity today, and what will they do in the future? To gain further insight into these trends, TechTarget's Enterprise Strategy Group (ESG) surveyed 370 IT and cybersecurity professionals at organizations in North America (US and Canada) responsible for cyber-risk management, threat intelligence analysis, and security operations, with visibility into current GenAI usage and strategic plans.

This study sought to:

- Identify current usage of and plans for generative AI.
- Establish how generative AI influences the balance of power between cyber-adversaries and cyber-defenders.
- Determine how organizations are approaching generative AI governance, policies, and policy enforcement.
- Monitor how organizations will apply generative AI for cybersecurity use cases.

KEY FINDINGS:

BACK TO CONTENTS

Generative AI Has a Foothold Today and Will Be Pervasive by the End of 2024

Many Are Already Developing Proprietary LLMs to Support a Variety of GenAI Use Cases

The research indicates that 85% of organizations have a proprietary large language model in place, while another 13% report that they are in the process of developing one. It's likely that many of these are early-stage projects based on open models. Nevertheless, this data indicates robust activity, forecasting rapid development of production applications in 2024.

Which business functions are ripe for GenAI applications? Organizations are well underway toward ubiquitous use of GenAI in areas like IT operations, software development, sales, research, and others. Like "shadow IT" in the past, growing use of GenAI will introduce cyber-risks as multiple departments and individual employees interact with open GenAI applications, experiment with open source, and create their own LLMs. CISOs must assess and communicate these risks while working with executives to create the right governance models and policies for risk mitigation. Additionally, security professionals must implement compensating controls and monitor users and networks for anomalous, suspicious, and malicious behavior.

Top five current or planned GenAI use cases.

Development plans for proprietary large language models.

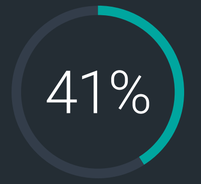

LLMs are already a part of production generative AI applications

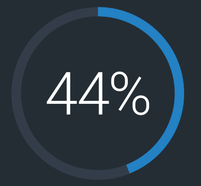

LLMs are already a part of an ongoing development project

We are just getting started developing our own LLMs

We plan to develop our own LLMs in the future

We are interested in developing our own LLMs in the future

BACK TO CONTENTS

Pervasive Security Concerns for Generative AI

Mitigating GenAI risk is not some future requirement. Rather, organizations are already dealing with, or struggling to address, generative AI risks proactively. Respondents claim that their organizations are already blocking access to generative AI sites. Security professionals are also concerned about GenAI-driven data leakage and remain unsure as to whether employees are already accessing and using GenAI applications today. CISOs must address these risks before losing control of who does what with GenAI across the enterprise.

> "Organizations are already dealing with, **or struggling to address**, generative AI risks proactively."

Cybersecurity professionals' GenAI opinions.

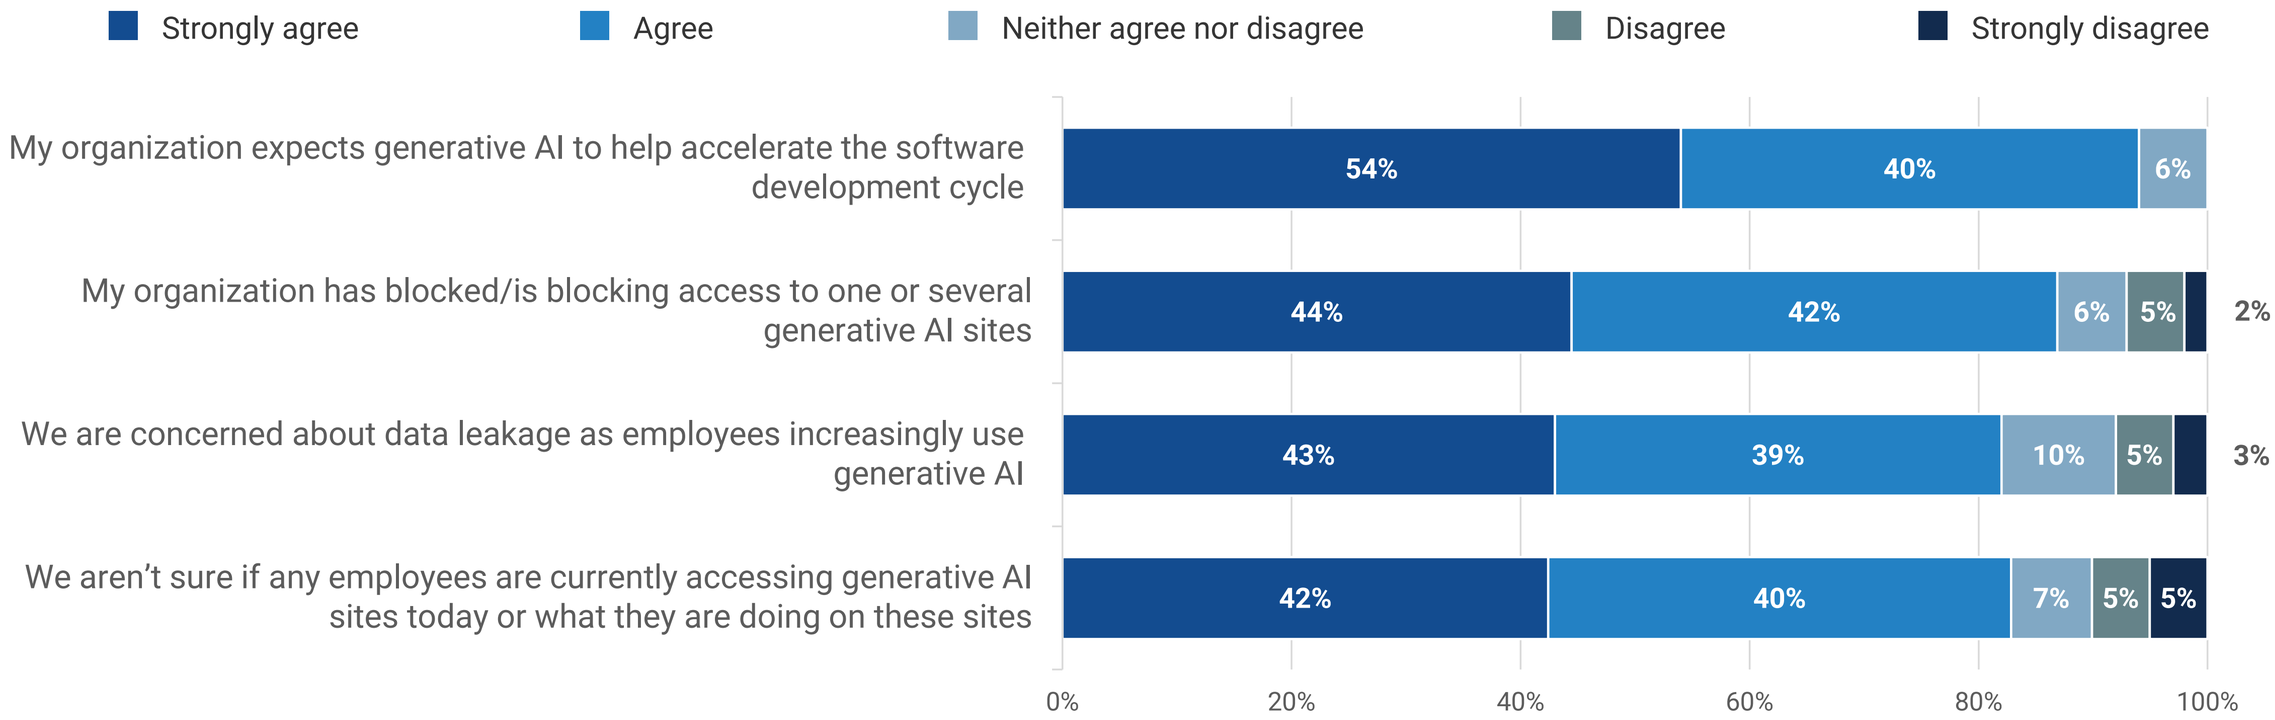

BACK TO CONTENTS

Many Are Getting Ahead of GenAI Governance, Especially Larger Organizations

Of course, most organizations have acceptable use policies and governance models in place today that may provide needed guardrails for GenAI usage. Recognizing additional risks associated with GenAI, many firms are creating specific governance models. Larger enterprise organizations, with the most at risk, are taking the GenAI governance lead, while their smaller counterparts lag. CISOs at these organizations must act as a catalyst to accelerate the creation of governance models, policies, and policy enforcement controls. Given the pace of GenAI innovation, delays will only lead to rapidly increasing risk.

Plans for a specific GenAI governance structure.

We already have a specific GenAI governance structure

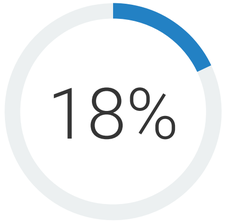

We are developing a specific GenAI governance structure

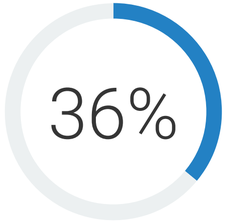

We are interested in developing a specific GenAI governance structure

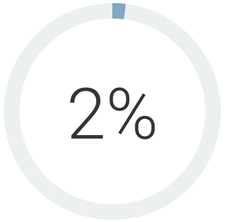

We believe our existing governance and acceptable use policies apply to generative AI

Percentage of organizations that already have a specific governance structure for the use of generative AI by company size.

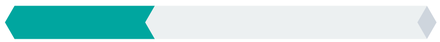

1,000 to 2,499 employees

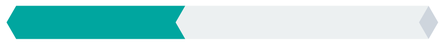

2,500 to 4,999 employees

5,000 or more employees

BACK TO CONTENTS

Weakest Areas of GenAI Governance

While many organizations have created or are creating GenAI governance models, security professionals admit that these models have abundant weaknesses, including third-party risk management (29%), data leakage protection/ rights management (27%), regulatory compliance (26%), and enforcing GenAI access policies (26%). Based on the data, it is safe to assume that most organizations have a myriad of these weaknesses in place today. Closing these gaps should be a 2024 priority.

Weakest areas of generative AI governance and policy enforcement.

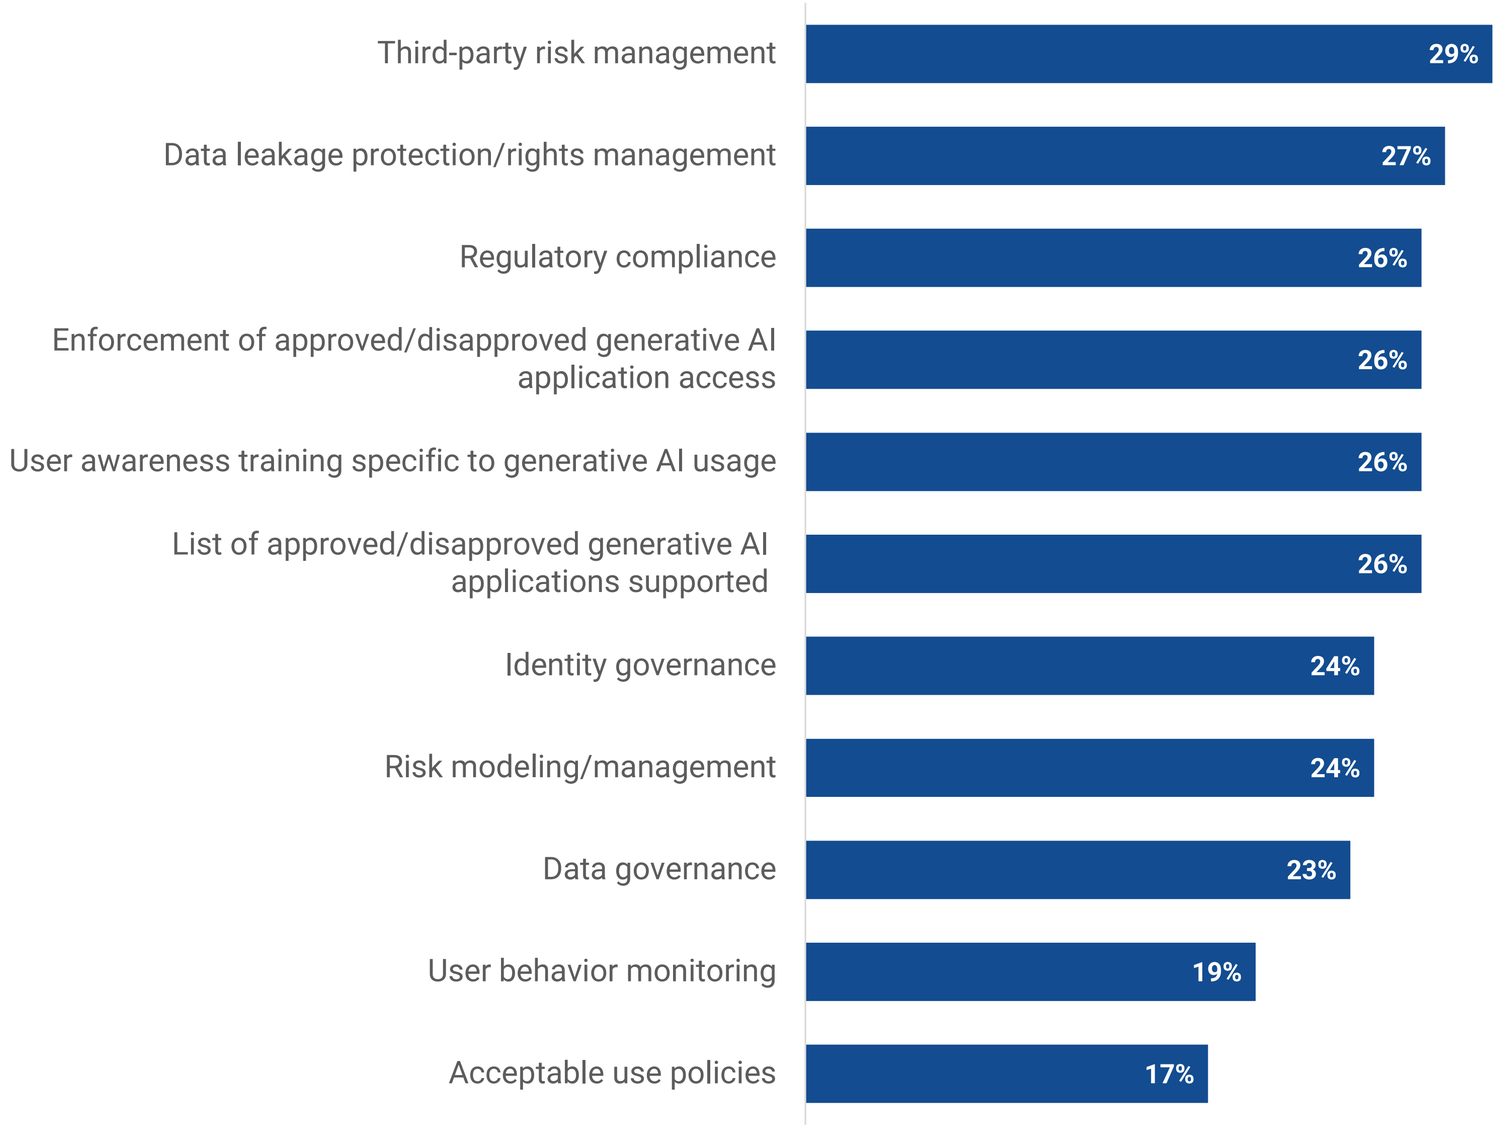

BACK TO CONTENTS

Organizations Anticipate Generative AI Risks

Generative AI Risks Associated With LLMs

As organizations experiment with GenAI applications and build their own LLMs, they face several application security challenges such as trust boundary risks (i.e., where data comes from or when it moves to an untrusted source), data management risks (e.g., data leakage or corruption), inherent underlying model risks, and general security risks. Security professionals are especially apprehensive about data management risks, while IT professionals' concerns focus on trust boundary risks. Organizations must reinforce their application security efforts with oversight, tooling, and testing to address specific GenAI application security.

Biggest risk to LLM and generative AI application development by role.

BACK TO CONTENTS

GenAI Balance of Power Skews Toward Cyber-adversary Advantage

Of course, cyber-adversaries also have access to open GenAI applications and have the technical capabilities to develop their own LLMs. WormGPT and FraudGPT are early examples of LLMs designed for use by cybercriminals and hackers. Will cyber-adversaries use and benefit from LLMs? More than three-quarters of survey respondents (76%) not only believe they will, but also feel that cyber-adversaries will gain the biggest advantage (over cyber-defenders) from generative AI innovation.

Alarmingly, most security professionals believe that cyber-adversaries are already using GenAI and that adversaries always gain an advantage with new technologies. Respondents also believe that GenAI could lead to an increase in threat volume, as it makes it easier for unskilled cyber-adversaries to develop more sophisticated attacks. Security and IT pros are also concerned about deep fakes and automated attacks.

Group expected to gain biggest advantage from generative AI innovation.

Top five reasons cyber-adversaries will likely gain the biggest advantage from GenAI technology.

Cyber-adversaries are already using generative AI technology to their advantage

Cyber-adversaries always gain an early advantage with new technologies

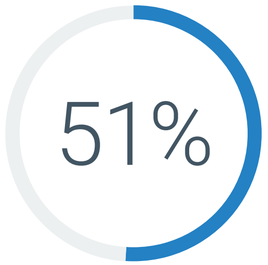

Generative AI makes it easier for unskilled cyber-adversaries to develop more advanced cyberattacks, so we anticipate greater volumes of attacks

Generative AI can help cyber-adversaries develop realistic deep fakes based on someone's voice or writing style, making it harder to detect malicious behavior

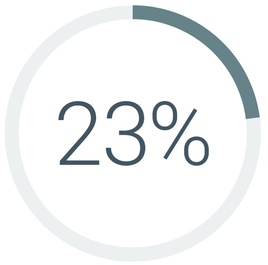

Generative AI can help cyber-adversaries automate attacks

BACK TO CONTENTS

GenAI Broadens the Attack Surface and Expands Social Engineering Tactics

Security professionals already have lots of ideas about how adversaries will use GenAI as a component of cyberattacks. Proprietary and partner-based GenAI systems will expand their attack surface, leading to unavoidable vulnerabilities and potential exploitation. Adversaries will use GenAI within phishing, business email compromise (BEC), and other types of deep fakes. There is also a concern about GenAI creating new types of malware (note: researchers have already created polymorphic malware using ChatGPT).

Given the unknown tactics, techniques, and procedures (TTPs), unclear balance of power, and potential GenAI threat vectors, cybersecurity professionals must stay vigilant by monitoring threat intelligence for signs of GenAI-based adversary TTPs and implementing the right controls for policy enforcement.

Most concerning use cases for cyber-adversaries using GenAI.

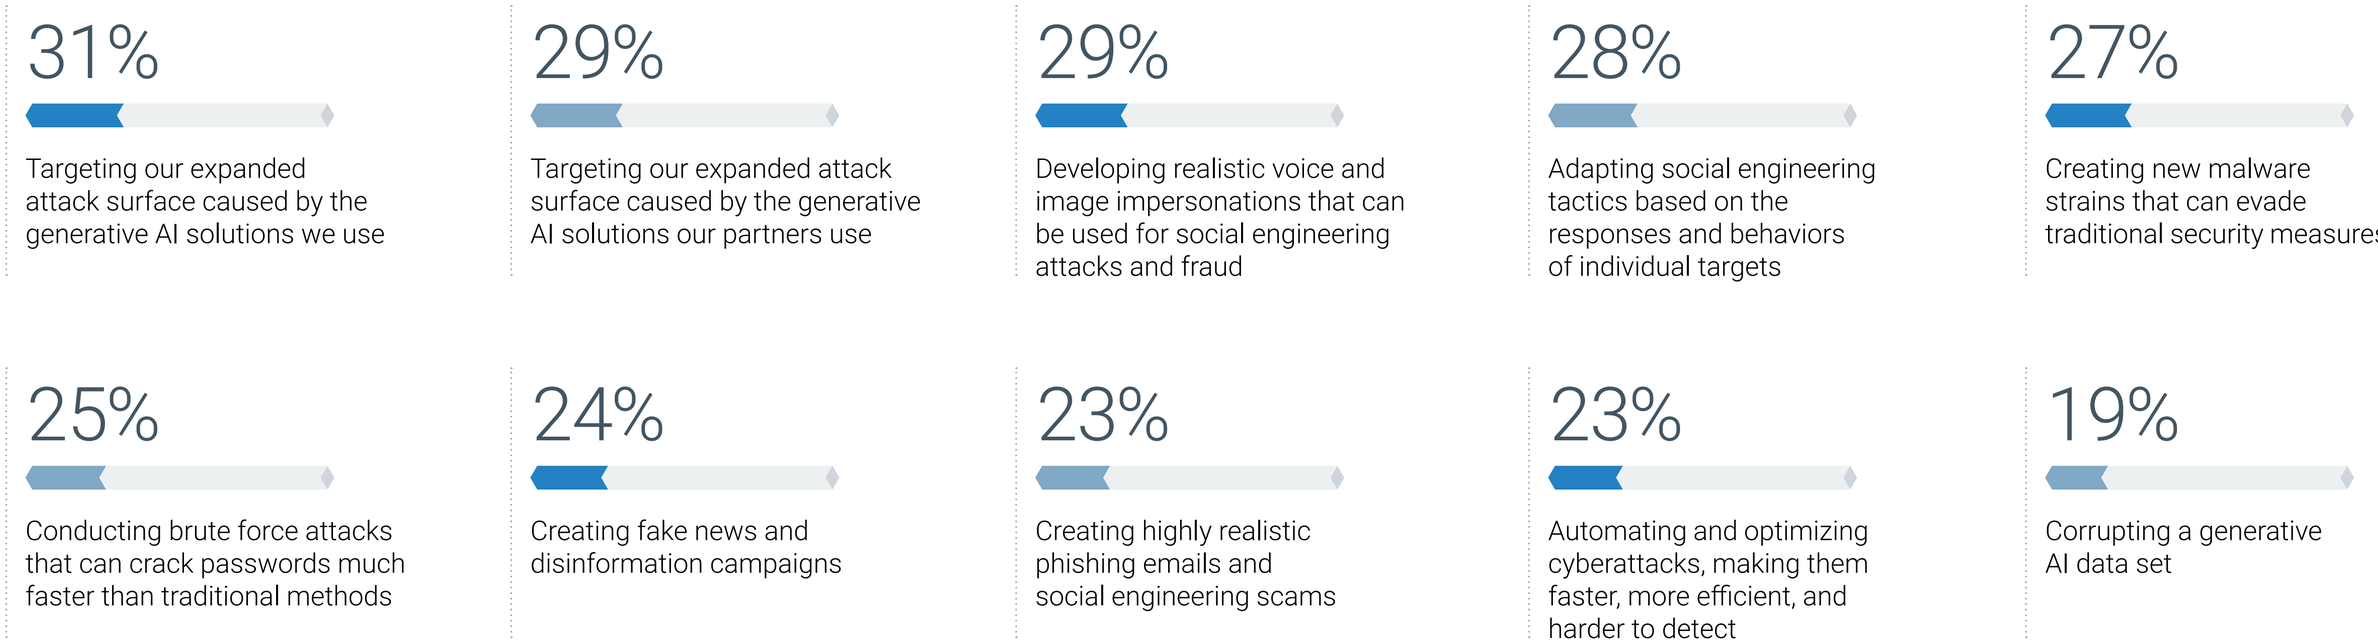

BACK TO CONTENTS

Important Security Tools for GenAI Policy Enforcement

While GenAI introduces new types of threats and vulnerabilities, organizations aren't completely defenseless. In fact, existing tools like CASB, SIEM, application security technologies, and identity and access management systems can block threats and enforce policies. For example, CASB systems can allow access to designated web-based GenAI applications for credentialed users while denying access to others. IAM systems can be used to create these credentials based on roles and specific use cases. Application security technologies can scan source code and create threat models that look for trust boundary or data usage violations. SIEM systems can log activity and implement detection rules to alert on any anomalous behavior.

These tools can provide a foundation for GenAI security, but they will need to be configured and customized appropriately to adhere to organizational policies, regulatory compliance, and threats targeting industries, regions, and individual firms.

Most important security controls for GenAI policy enforcement.

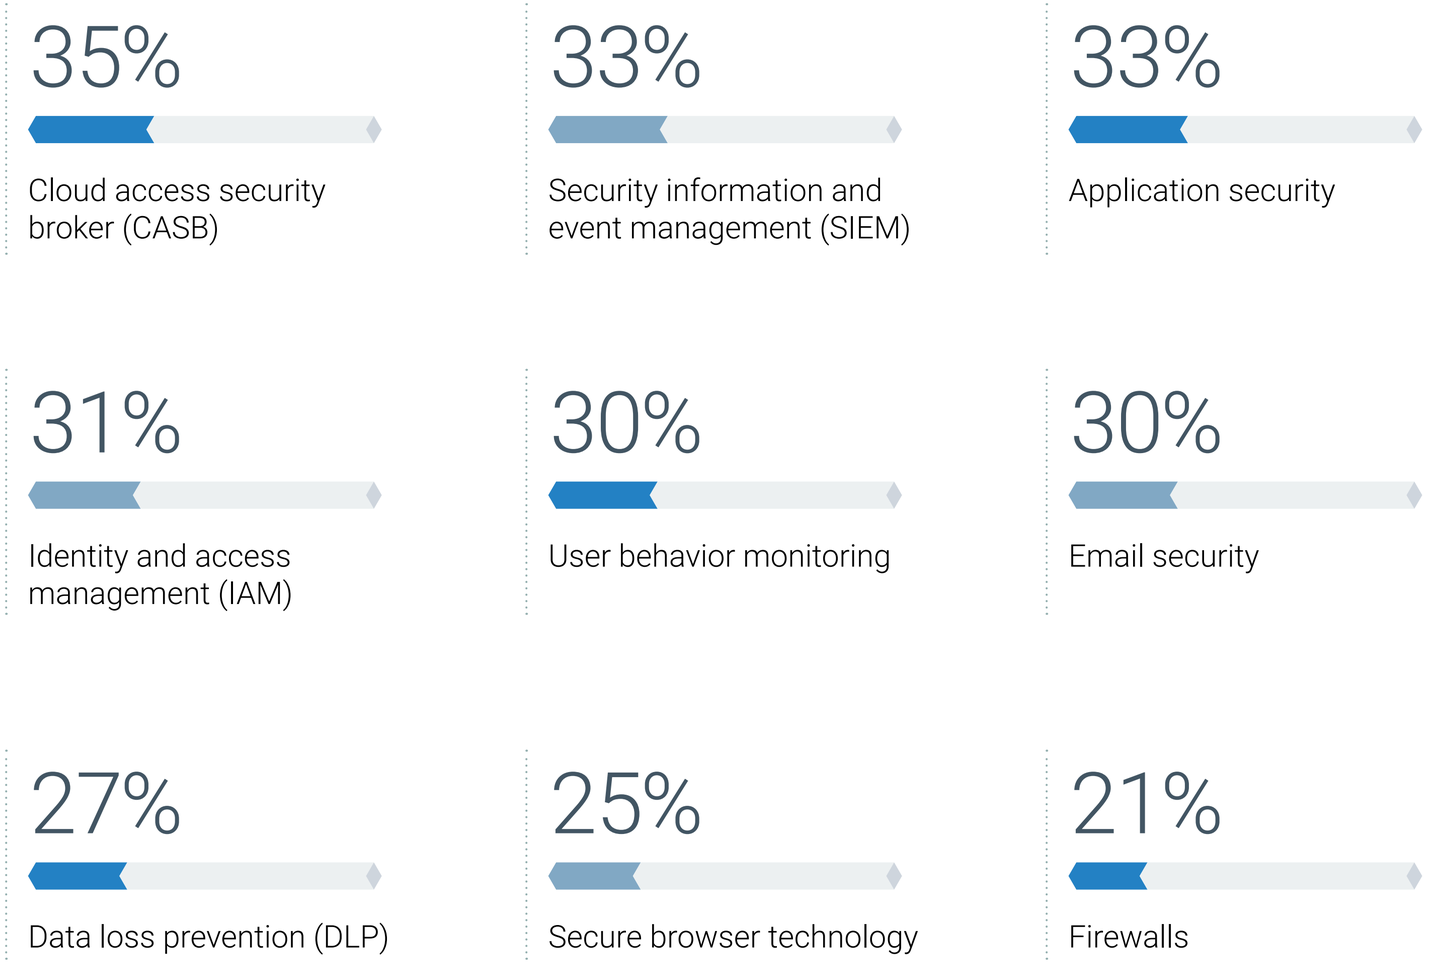

BACK TO CONTENTS

Security Professionals Are Cautiously Optimistic About Generative AI's Potential for Cybersecurity

Sentiments Toward GenAI and Cybersecurity

GenAI for cybersecurity elicits an emotional response from security and IT professionals, ranging from optimistic and excited to reserved and skeptical. Beyond these emotions, survey respondents had some concrete opinions: 93% believe generative AI could help them improve their knowledge, skills, and abilities, while at the other end of the spectrum, 80% remain skeptical about the value GenAI can bring to cybersecurity. Many (92%) believe their organization would be willing to replace existing technologies based on the GenAI capabilities of another similar product.

Level of agreement with statements related to using GenAI to fortify existing cybersecurity tools and processes.

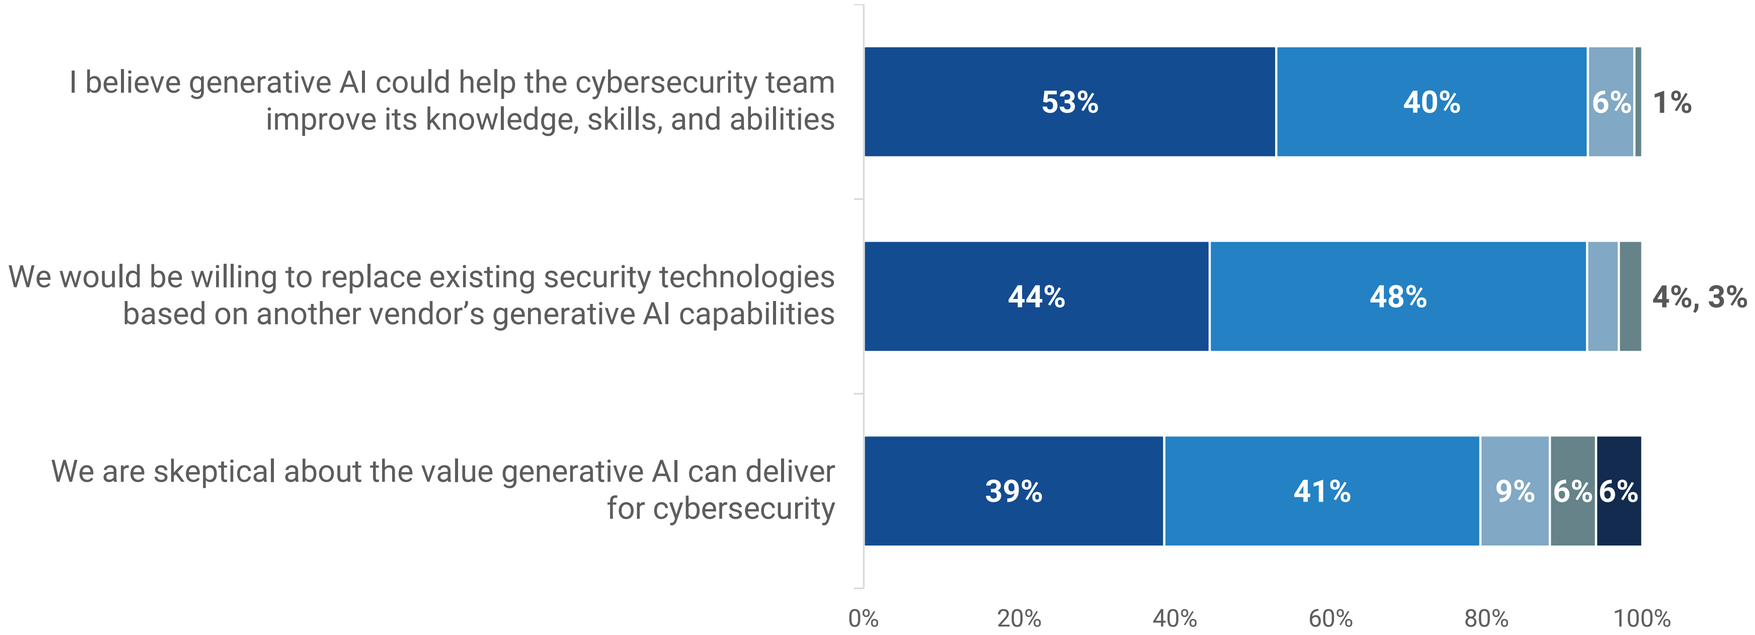

"GenAI for cybersecurity **elicits an emotional response from security and IT professionals**, ranging from optimistic and excited to reserved and skeptical."

Jon Oltsik | Distinguished Analyst and Fellow
ENTERPRISE STRATEGY GROUP

BACK TO CONTENTS

"Given these misgivings, **CISOs should champion GenAI for cybersecurity efforts while remaining judicious in its adoption**."

Clearly, opinions vary widely, even between IT and security respondents. For example, security professionals tended to be more reserved (28% versus 19% for IT pros), fearful (23% versus 15% for IT pros), and pessimistic (20% versus 12% for IT pros). Alternatively, IT pros were more neutral than their security colleagues (26% versus 20% for security pros). Given these misgivings, CISOs should champion GenAI for cybersecurity efforts while remaining judicious in its adoption.

Sentiments of cybersecurity defenders regarding the potential impact of GenAI.

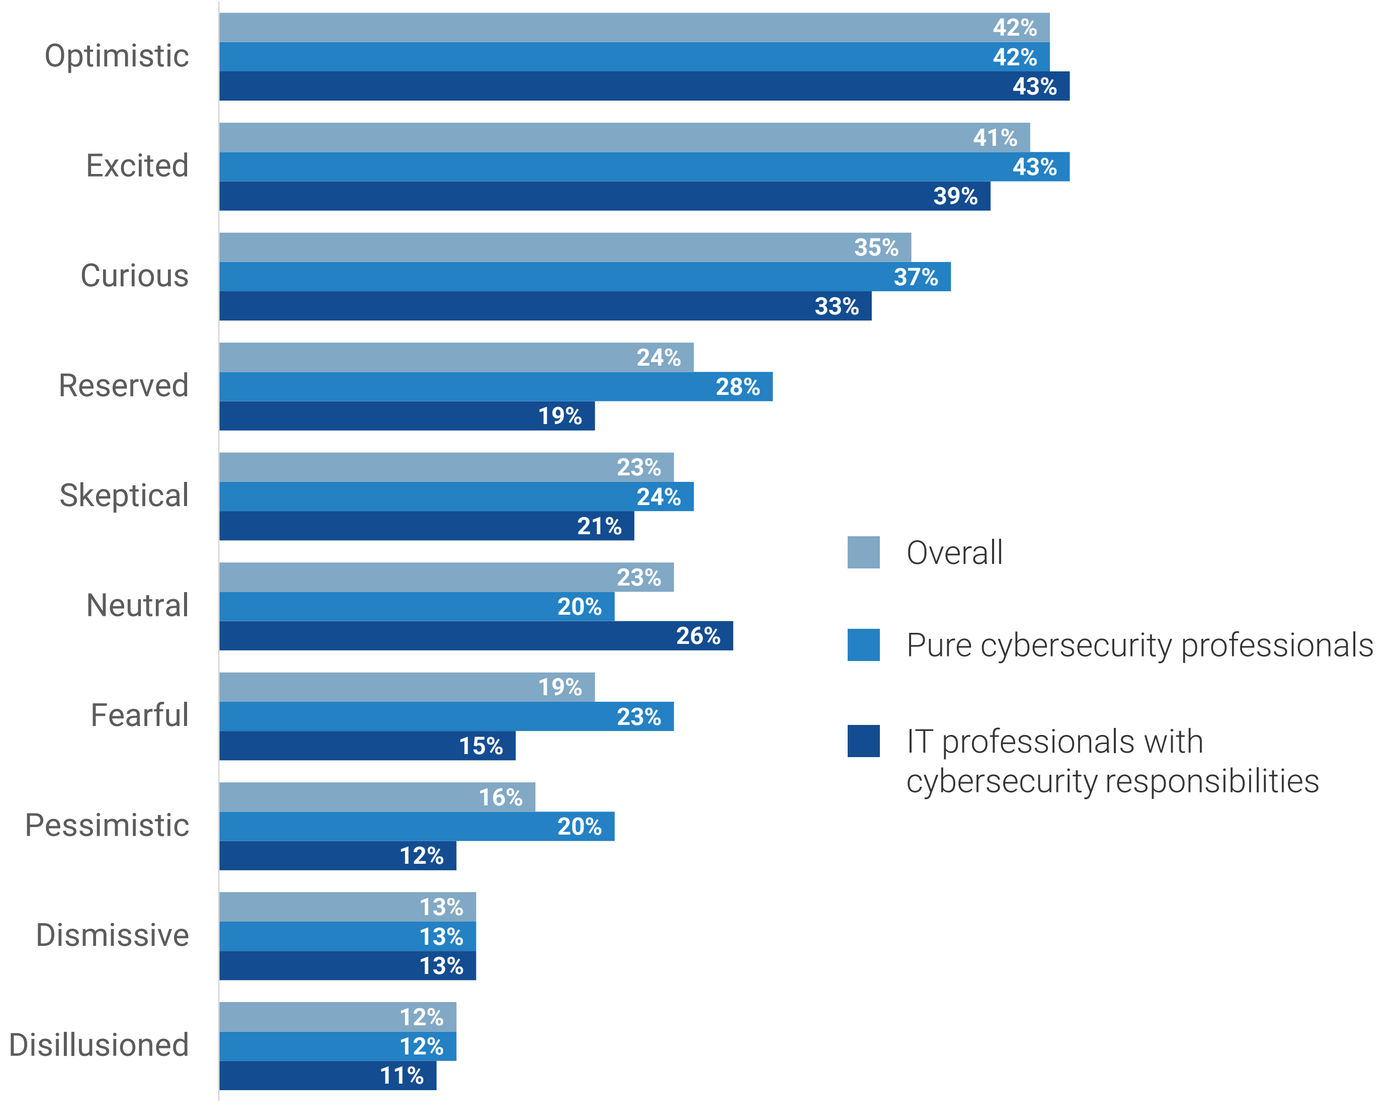

BACK TO CONTENTS

Security Professionals' Opinions on Machine Learning

While generative AI may be a shiny new cybersecurity object, other AI technologies, like machine learning, have been around for nearly a decade. With all these years of experience, 92% of respondents agree that machine learning has improved the efficacy and efficiency of cybersecurity technologies. Still, this benefit didn't come without baggage. It's apparent from this research that machine learning-based cybersecurity technologies created false positive alerts and needed tuning and customization. And while they've become commonplace today, it took many years for vendors to deliver valuable machine learning (ML) for cybersecurity. It's worth noting that 83% of respondents agree that looking back, the hype around machine learning for cybersecurity was excessive. Little wonder then why cybersecurity pros are reserved, pessimistic, and fearful of machine learning technology. With this experience, CISOs should take a prudent approach to GenAI implementation, looking for quick wins and incremental value rather than force multiplying results.

Level of agreement with statements related to machine learning and cybersecurity technologies.

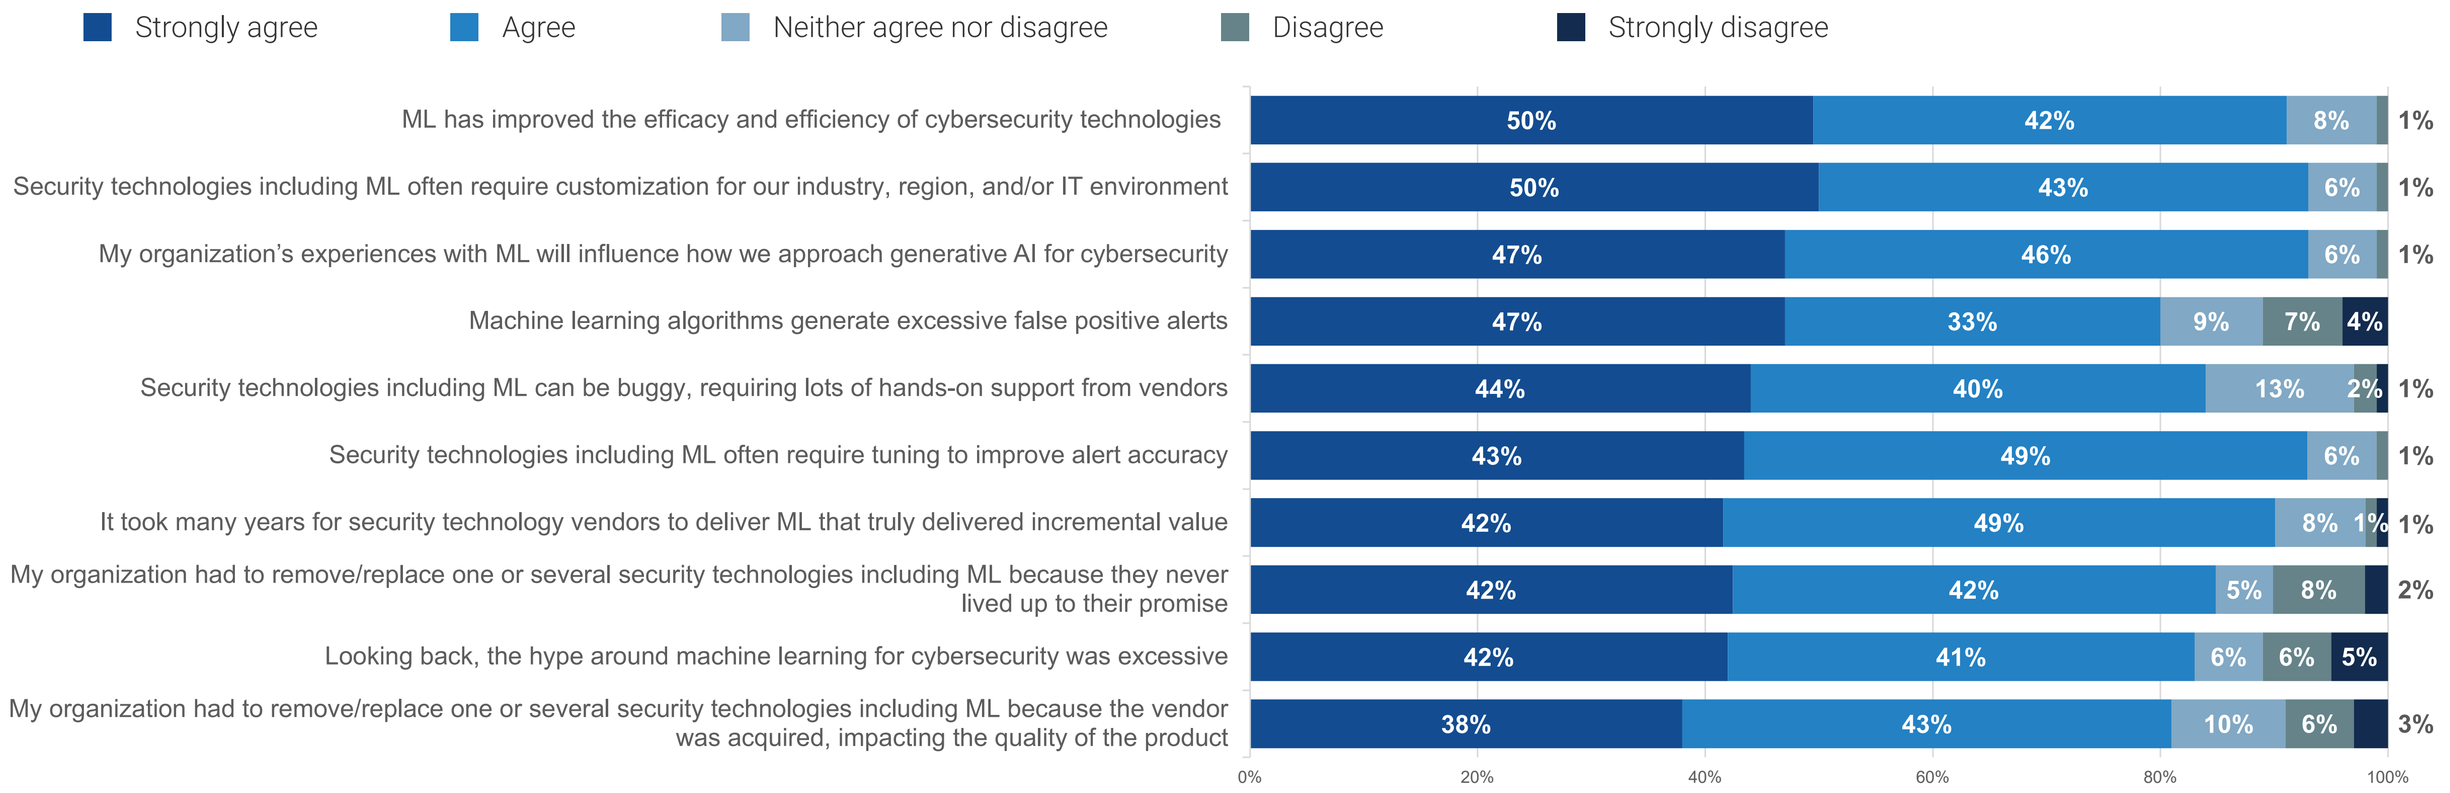

BACK TO CONTENTS

Promising GenAI Cybersecurity Use Cases

Enterprise organizations will likely upgrade existing tools like EDR, SIEM, and threat intelligence platforms (TIPs) with new versions containing GenAI functionality. Some of the more advanced organizations will create their own security LLMs or customize existing models to meet organizational, industry, and regional needs. As they experiment with open AI engines, their goals must include improving security efficacy, efficiency, and staff productivity.

To that end, survey respondents see a number of promising use cases, including improving security hygiene and posture management, guiding the cybersecurity staff with recommended actions, consolidating security technologies, and training junior employees. Along with these, ESG has seen practical use and strong results in two other areas cited: Security professionals are saving time by using GenAI to help them create summary reports (i.e., threat intelligence reports, incident response reports, cyber-risk status reports, etc.). Additionally, security teams are starting to rely on GenAI to help them better analyze threat intelligence germane to their organization, industry, and region. Security pros will also use GenAI for natural language queries, improving homegrown software and creating detection rules in the near future.

Areas in which GenAI is believed to hold the most promise for cybersecurity.

BACK TO CONTENTS

"These actions could scale staff productivity, but skeptical cybersecurity teams **aren't quite ready for autonomic processes just yet.**"

Response Actions for Low-risk GenAI Recommendations

Generative AI also has the potential to automate security processes like patching vulnerable software, quarantining a system, or disabling a user account. These actions could scale staff productivity, but skeptical cybersecurity teams aren't quite ready for autonomic processes just yet. If their organization were given a low-risk recommendation from a GenAI application, nearly three-quarters (71%) of respondents say that a staff member would review the recommendation before taking the remediation action manually, while 25% claim a staff member would review the recommendation and, if they approve it, allow the generative AI application to execute the remediation action. Only 4% would let the application execute the remediation action on its own without human intervention.

Approach to handling low-risk recommendations from a GenAI application.

71% A staff member would review the recommendation before taking the remediation action manually

25% A staff member would review the recommendation and, if they approve it, allow the generative AI application to execute the remediation action

4% The generative AI application would execute the remediation action on its own, without any human intervention

BACK TO CONTENTS

Generative AI and Security Process Automation

While security teams remain cautious, they do pinpoint where GenAI-based process automation could be helpful in areas like security investigations, orchestration across heterogeneous security controls, case management, and alert enrichment. Anecdotally, ESG has heard similar themes from CISOs. They'd like to automate alert triage processes to bolster the productivity of tier one/junior analysts. They believe GenAI could increase security analysts' throughput if it could string together individual alerts into attack chains and then enrich these attack chains with threat intelligence and put them in a MITRE ATT&CK context.

End-to-end GenAI-based process automation may be a scary thought for security teams. Perhaps they will get there over the next few years, but in the meantime, GenAI should be applied sensibly to low-risk, time-consuming processes that act as bottlenecks for security teams.

Areas in which security operations process automation supported by GenAI would be most helpful.

BACK TO CONTENTS

Generative AI Will Become a Purchasing Consideration, Though Plans Are Still Developing

GenAI Solutions Drive Cybersecurity Spending

It appears that GenAI cybersecurity technologies are viewed as a generational change. This conclusion is based on the fact that more than three-quarters (78%) of organizations plan on increasing spending for security solutions with generative AI capabilities. And while the bulk of new spending (73%) will focus on adding new generative AI products or features, CISOs will also dedicate budget dollars to security technologies to mitigate risks and enforce security policies specific to general GenAI initiatives. This spending balance will likely vary widely based on industry, regulations, and organizations' GenAI application usage. The more GenAI in use, the more organizations will need appropriate compensating controls for risk mitigation.

Impact of security solutions with GenAI capabilities on security budget.

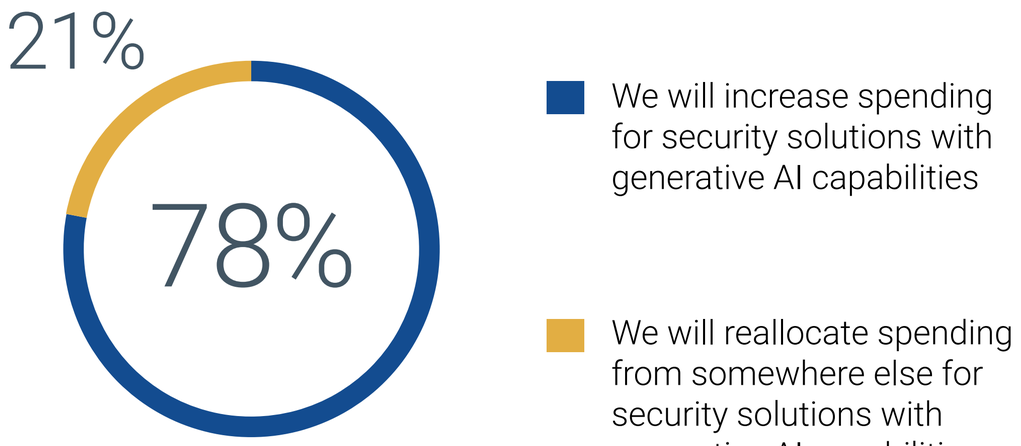

Priorities for spending on GenAI over the next 12-24 months.

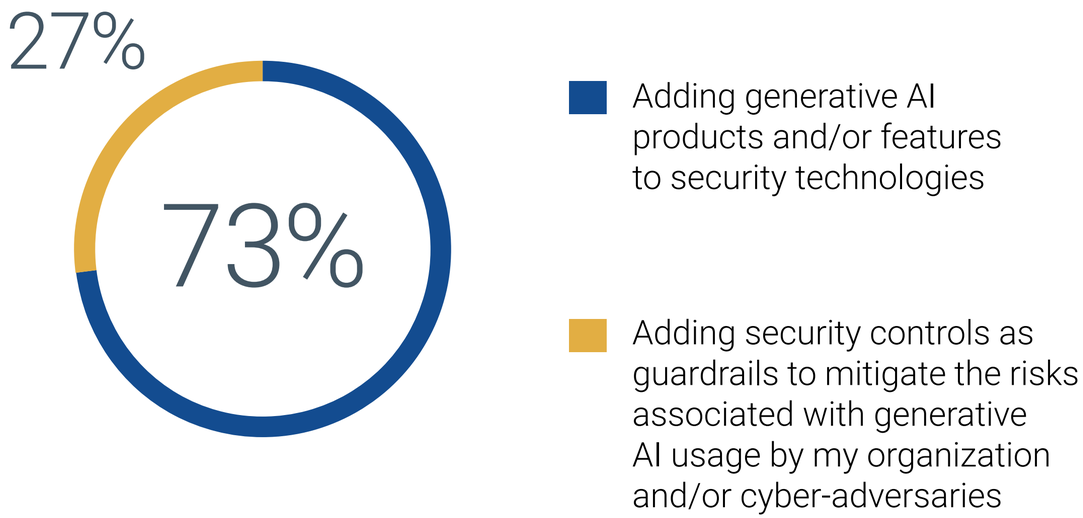

BACK TO CONTENTS

Generative AI and Security Vendor Selection

At this point, more than half (54%) of organizations are leaning toward choosing a single vendor and technology for generative AI security tools, but this strategy may change with time and rapid technology innovation. In fact, 31% of organizations will choose a primary GenAI vendor/technology but selectively use others, and 15% will use GenAI capabilities across numerous or a few select vendors/technologies.

Since GenAI-based security technologies are in their genesis phase, organizations will have to learn how to evaluate products and capabilities as part of their RFI/RFP processes. As of now, the top purchasing considerations include vendors having a formal documented process for data security/privacy as part of development and operations, capabilities built on top of multiple LLMs (including open source versions), an open architecture, and lots of experience with AI/ML-based products. As previously stated, CISOs will likely use the lessons learned from their previous experiences with machine learning to evaluate tools and set the right expectations within the organization.

Most important considerations for GenAI capabilities offered by security vendors.

38%

A formal documented process for data security/privacy as part of development and operations

36%

Capabilities built on top of multiple LLMs, including open source versions

34%

An open architecture that can accomodate data from other security vendors as part of the LLM

33%

A long history of offering products based on AI/ML

29%

Capabilities built on top of existing LLMs like GPT 4

29%

Capabilities offered by one of the vendors my organization currently works with

28%

An LLM based on security content (i.e., threat intelligence, CVEs, malware analysis, etc.)

28%

All LLM data being stored within my country's borders

Likely strategy for augmenting security tools with generative AI products and/or features.

54%

Choose a **single vendor/ technology** for generative AI

31%

Choose a **primary vendor/ technology** for generative AI but **selectively use other vendors/technologies**

8%

Use generative AI capabilities **across numerous vendors/ technologies**

7%

Use generative AI capabilities equally across **a few select vendors/technologies**

BACK TO CONTENTS

Thoughts on GenAI Pricing for Cybersecurity

Generative AI capabilities will likely come as additional features in new cybersecurity product revisions. Vendors are expected to charge a premium for GenAI capabilities, so ESG asked security professionals which pricing models they preferred in this scenario. The data suggests a bit of uncertainty, which is expected for this type of new technology. Nevertheless, survey respondents' preferences range from pricing based on the number of security personnel using GenAI through pricing baked into product pricing to pricing based on the total number of users within an organization.

Despite these preferences, GenAI pricing is likely to be a quagmire for the foreseeable future as the market develops. Security and purchasing managers should monitor different pricing models and compare notes with other organizations in their industry as they negotiate with vendors.

Most appropriate pricing models for security technologies that include GenAI capabilities.

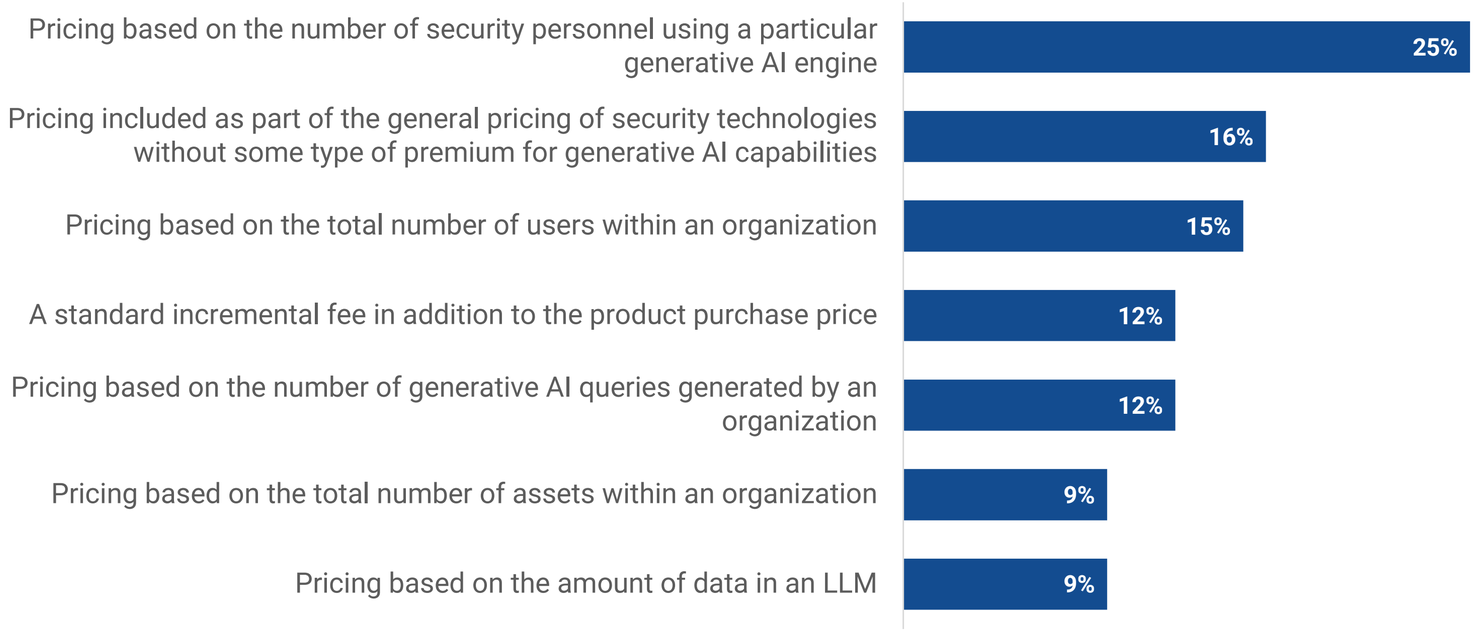

"GenAI pricing is **likely to be a quagmire for the foreseeable future** as the market develops."

BACK TO CONTENTS

RESEARCH METHODOLOGY AND DEMOGRAPHICS

To gather data for this report, TechTarget's Enterprise Strategy Group conducted a comprehensive online survey of IT and cybersecurity professionals from private- and public-sector organizations in North America between November 6, 2023 and November 21, 2023. To qualify for this survey, respondents were required to be involved with supporting and securing, as well as using, generative AI technologies. All respondents were provided an incentive to complete the survey in the form of cash awards and/or cash equivalents.

After filtering out unqualified respondents, removing duplicate responses, and screening the remaining completed responses (on a number of criteria) for data integrity, we were left with a final total sample of 370 IT and cybersecurity professionals.

Respondents by number of employees.

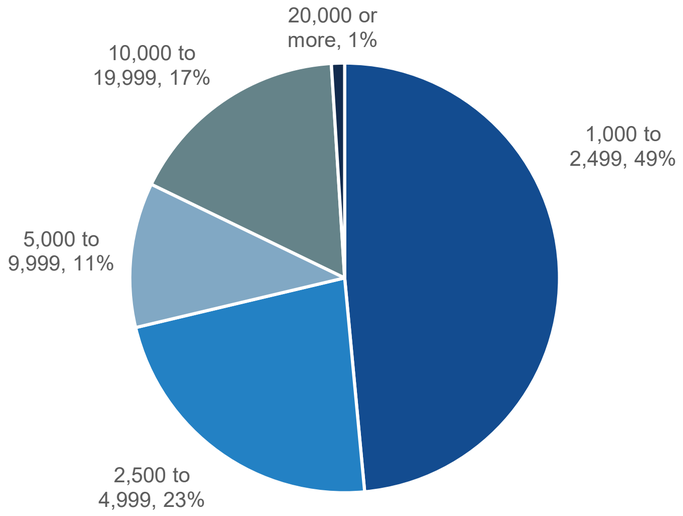

Respondents by age of organization.

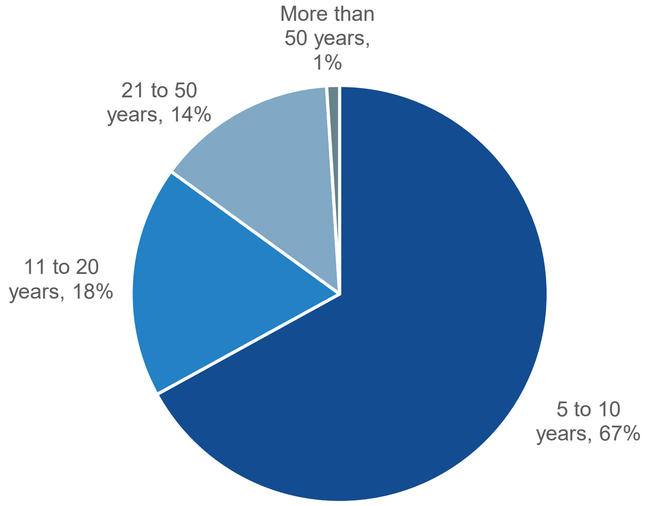

Respondents by industry.

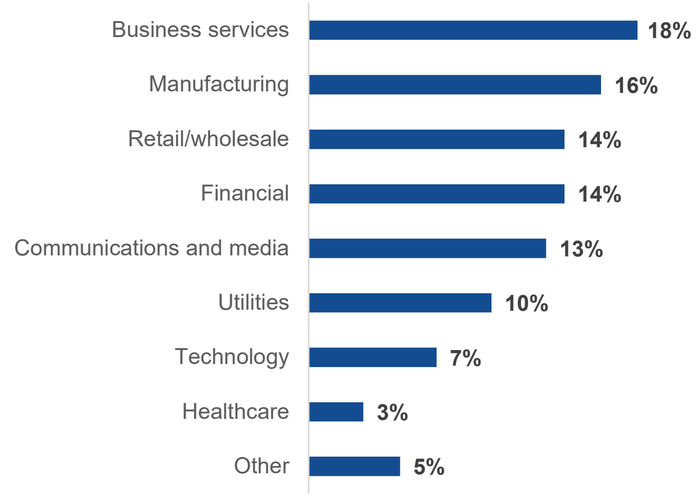

BACK TO CONTENTS

Enterprise Strategy Group™
by TechTarget

Enterprise Strategy Group is an integrated technology analysis, research, and strategy firm providing market intelligence, actionable insight, and go-to-market content services to the global technology community.