




**Harvard
Business
Review**

ANALYTIC SERVICES



PULSE SURVEY

Advancing the Use of Artificial Intelligence While Mitigating the Cybersecurity Threat



Sponsored by



SPONSOR PERSPECTIVE



Nicole Herskowitz
Corporate Vice President, Microsoft
Microsoft 365

In today's rapidly evolving digital landscape, artificial intelligence (AI) stands at the forefront of technological innovation, promising unprecedented advancements in productivity and efficiency. Organizations around the world are realizing the significant impact that AI can have on transforming work, improving efficiency, and sparking innovation. Yet, with this excitement it is important to note that not everyone is waiting for their organization to formally adopt AI, as many are bringing their own AI tools to work. While this trend underscores the enthusiasm and eagerness for AI, it also reminds us of the ever-present, growing, and significant security risks that come with AI adoption.

At Microsoft, we recognize the transformative power of AI and the pivotal role it plays in shaping the future of business. However, we recognize that with this transformative power comes great responsibility, particularly in safeguarding organizations against the ever-growing threats in cybersecurity.

The intersection of AI and cybersecurity is not just a topic of significance but one of urgency. The World Economic Forum has highlighted cybersecurity as one of the top global threats, reflecting the increasing frequency and sophistication of cyber attacks. As AI systems become more integrated into our daily operations and infused into how we work, it becomes more necessary to protect these systems from cyber threats than before the AI era. This is not merely about defending against potential breaches but also about ensuring the integrity and trustworthiness of AI technologies themselves. Furthermore, it's important to highlight the necessity for robust data governance and the implementation of advanced security strategies, to minimize data oversharing and data leakage when using AI tools.

Our commitment to advancing AI responsibly is rooted in a deep understanding of these complexities. We are dedicated to developing AI solutions that not only propel business growth and drive innovation but also fortify our defenses against the evolving threat landscape. These strategies are essential in creating an AI-ready organization that can withstand the changing landscape of cyber threats in this new era.

This Harvard Business Review Analytic Services report delves into the critical balance between leveraging AI's potential and mitigating its associated risks. It underscores the importance of a pragmatic approach, one that marries enthusiasm for AI's capabilities with a resolute commitment to cybersecurity. At Microsoft, we advocate for a culture of shared responsibility where every stakeholder, from executives to frontline employees, plays a role in safeguarding our digital future.

As we move forward, it is imperative for organizations to stay ahead of the curve, continuously evolving their security measures in tandem with AI advancements. So, whether you have already implemented advanced AI solutions or are just starting on the journey, it's important to fully understand the entire threat landscape to minimize your security risks while maximizing the benefits of AI for your organization.

We are not just participants in this journey but leaders, setting the standards for responsible AI deployment and cybersecurity excellence. Our vision is clear: to embrace the transformative power of AI while ensuring a secure and resilient digital ecosystem for all.

Advancing the Use of Artificial Intelligence While Mitigating the Cybersecurity Threat

Artificial intelligence (AI) is being used by organizations at a time when businesses are under increasing attack from cybercriminals. The World Economic Forum has included cybersecurity in its top 10 list of global threats as cyber attacks grow in frequency and sophistication.¹ Nine in 10 respondents to a Deloitte Global 2023 Future of Cyber survey reported at least one compromise in their organization.²

Organizations admit they are worried about how AI and generative AI (gen AI) will impact this already challenging cyber threat landscape. A survey conducted by Harvard Business Review Analytic Services in September 2024 of 227 respondents from the *Harvard Business Review* audience, all involved in making or implementing their organization’s decisions about IT, finds that more than three-quarters (77%) of respondents say their organization is “very concerned” about cybersecurity and data privacy in the age of AI. A further 18% say their organization is “somewhat concerned.”

Despite this widespread apprehension, executives are caught up in a wave of AI excitement and investment. The use of AI tools is proliferating quickly across organizations of all sizes and in all sectors, and the executives in charge of securing the enterprise and its data—the chief information officer (CIO), chief information security officer (CISO), and chief technology officer—must now find ways to support their organizations in leveraging these tools in a managed, secure, and productive way.

Such a mandate is challenging because securing the enterprise is not straightforward and the current levels of complexity and hype around AI only complicate things further, according to Omar Khawaja, vice president of security and field CISO at San Francisco-based Databricks Inc., a global data and AI company. “What I am seeing at the moment is many organizations having a

HIGHLIGHTS



77% of respondents say their organization is “**very concerned**” about **cybersecurity and data privacy** in the age of artificial intelligence (AI).



71% agree with the statement, “My organization **recognizes the risks of AI**, but **believes the benefits** to the business outweigh the risks.”



57% of respondents at organizations moving forward with AI say their organization is **training employees** on how to use AI tools safely.

Due to rounding, some figures in this report may not add up to 100%.



“AI is going to change the world in ways we can’t even imagine yet. The organizations that ignore this revolution will share the fate of those organizations that ignored the internet at the start of the dot-com boom,” says Mikko Hypponen, chief research officer at WithSecure Oyj.

deep appetite for AI and a lack of patience to consume it, which is not a good combination from a risk perspective,” he says.

For Khawaja, this mix calls for a pragmatic assessment of business value and risks and a level-headed approach to managing the risks so cybersecurity and business teams can move forward together. “We need people to move away from the edges of the spectrum of extreme fear or extreme excitement to a place where we accept there are some risks associated with AI but also amazing positive outcomes,” he asserts. “And we need to identify controls to mitigate those risks so we can turn the balance of the equation toward the positives.”

There are many tools and approaches at the disposal of today’s information, security, and technology senior managers. Safeguarding privacy and customer identities, securing endpoints, protecting data, and implementing zero trust security, where access to IT systems and tools is verified for every user at every point of entry, are now foundational elements of creating an AI-ready organization. Organizational culture plays a critical role in ensuring employees understand how to use AI tools safely, too.

This report explores the evolving cyber threat and data privacy risks that companies are facing today and the kinds of approaches organizations are taking to benefit from the best of AI technology while also protecting themselves from those risks when using it.

A Mix of Excitement and Fear

AI models are fundamentally learning systems that rely on data to learn and then produce outputs. In theory, this process heightens certain cybersecurity- and data privacy-related risks—for example, the possibility of cybercriminals seeking entry into AI systems to get ahold of sensitive data or to interfere with data coding and output. Exabeam Inc., a global cybersecurity company headquartered in Foster City, Calif., describes cybersecurity in the context of AI as “the measures and technologies designed to protect [AI] systems from cyber threats and ensure their secure operation. This includes safeguarding the data AI systems are trained on, protecting the integrity of AI algorithms, and ensuring that AI applications are not used for malicious purposes.”³

While many experts believe the threat of cybercriminals hacking into AI systems could be on the horizon, the nearer-term worry for businesses is criminals using AI to create more convincing and sophisticated attacks. Dr. Keri Pearson, executive director of the research group Cybersecurity at MIT Sloan (CAMS) at the MIT Sloan School of Management in Cambridge, Mass., says the threats associated with AI generally fall into two broad categories: “One, the bad guys can use these tools to create faster, better, more effective attack vectors. And the other threat is that our people don’t really understand how these tools work and use them in a way that is helpful to them but might introduce some potential security vulnerabilities.”

Despite anxiety about such cybersecurity challenges, AI is becoming ubiquitous, with most organizations and workers now using, or at least exploring, AI within their business or to support their day-to-day work. For the most part, the benefits of AI for organizations are too good to pass up. Promises of improved productivity and better employee and customer experiences, for example, have motivated companies to forge ahead with their AI plans. Nearly two-thirds (65%) of respondents whose organizations are moving forward with AI say they are seeing improved productivity or efficiency as a result. Forty-one percent say they are seeing increased innovation. Only 10% of survey respondents say their organization is not moving ahead with AI, while 56% have live AI use cases or pilots in place. A further 13% have a plan for using AI but have not deployed it, while 21% are in the exploration phase but have no plans yet.

This speed of adoption underscores the levels of enterprise enthusiasm for AI technology. Mikko Hypponen, chief research officer at WithSecure Oyj, a Helsinki-based global enterprise cybersecurity solutions company, believes the true impact of AI will be greater than many appreciate. Disregarding this upheaval, he predicts, is just as risky as implementing AI without due care. “We’ve all seen the incredible things these tools can do over the past two years, but this is just the very beginning,” Hypponen explains. “AI is going to change the world in ways we can’t even imagine yet. The organizations that ignore this revolution will share the fate of those organizations that ignored the internet at the start of the dot-com boom.”

If not addressed, fears related to cybersecurity and data privacy could prevent companies from unlocking the power of AI—and transforming along with it. In fact, when asked what barriers stand in the way of expanded AI adoption at their organization, “cybersecurity/privacy concerns” was the number one barrier selected by survey respondents (45%). **FIGURE 1**

However, while concerns about cybersecurity and data privacy risks are high, most organizations are quite optimistic about the overall impact of AI on their business. More than two-thirds (71%) of respondents either strongly agree or somewhat agree with the statement, “My organization recognizes the risks of AI, but believes the benefits to the business outweigh the risks.”

The Evolving Threat Landscape

As organizations weigh the benefits and risks of AI, they will need to be alert to threats from both outside their company and within it. While deliberate employee misuse of AI is uncommon, Databricks’ Khawaja says there is a visibility risk related to teams using these tools in an unsanctioned way. “There are cases where business units are adopting AI in the shadows and not letting the rest of the organization know that they are doing it,” he explains. “And it is intentional, because they know there’s a right and a wrong way of doing it from an IT and a legal point of view and they want to avoid the internal bureaucracy.”

These visibility issues aside, businesses largely appear alert to the two very different types of threats AI introduces: cybercriminals using AI for attacks and employees misusing it unintentionally. Some 71% of respondents say their organization is “very concerned” about cybercriminals using AI to launch more sophisticated attacks such as phishing (another 22% are “somewhat concerned” by this). Sixty-four percent say their organization is “very concerned” about employees inadvertently feeding sensitive data into public AI tools, such as free chat-style tools that return content based on user prompts (with a further 27% “somewhat concerned”). Another 64% say their organization is “very concerned” about cybercriminals extracting sensitive information from their AI models (22% are “somewhat concerned”).

In other words, cybercriminals may be just as excited about AI as senior management is, since the technology is not only driving productivity and efficiency within organizations but also giving wrongdoers new opportunities to exploit cybersecurity vulnerabilities or extract sensitive information from businesses.

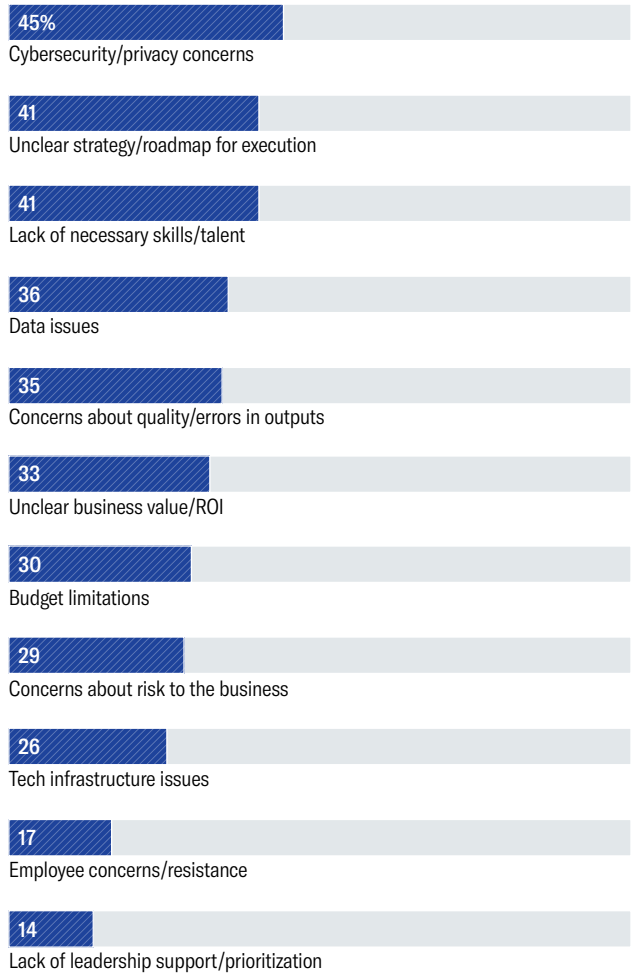
Here, WithSecure’s Hypponen explains that cybercriminals generally target businesses in one of three main ways: ransomware attacks, where important company data assets are held ransom by criminals who demand payment for their

FIGURE 1

Artificial Intelligence Anxiety

Cybersecurity/privacy concerns are the most common barrier to expanded AI adoption

What barriers stand in the way of expanded AI adoption at your organization, if any? *Select all that apply.*



Base: 227 respondents. Not shown: 5% Other, 0% None, 0% Don't know.

Source: Harvard Business Review Analytic Services survey, September 2024

release; business email compromise attacks, like phishing, where employees are tricked into disclosing sensitive information to criminals; and denial of service attacks, which are similar to ransomware attacks but are often carried out for protest or political reasons and can involve, for example, an organization’s entire web presence being taken down.

Each of these attack types has existed for many years, and the question for many cybersecurity experts now is whether AI,



“There have been a few high-profile cases of employees feeding tools with a company’s intellectual property [IP] only for that IP to then be used in the replies given to other users. The company has exposed something it didn’t want to expose, and now the question is ‘Who owns the answer?’” says Dr. Keri Pearlson, executive director of the research group Cybersecurity at MIT Sloan (CAMS) at the MIT Sloan School of Management.

and specifically gen AI, has made these threats any worse. In most cases, Hypponen says, there is little evidence to indicate they have—but that is unlikely to remain the case for long. “For ransomware and denial of service attacks, we are not seeing any real overlap where AI is shaping these,” he notes. “However, we are starting to see the use of gen AI-created deepfakes in business email compromise.”

Cybercriminals can now create convincing deepfakes—audio or video imitations of real people—to enhance the social engineering they carry out as part of their phishing attacks. Hypponen says such attacks are still rare and businesses have a window of opportunity to improve their cyber defenses. “I have only dealt with two cases with real evidence that deepfakes were used in a business email compromise scam, and both used deepfake audio,” he explains. “But it’s quite clear it is going to get worse, and the fact that it isn’t worse yet is nice, as companies still have some time to get ready.”

Khawaja acknowledges that it’s difficult to know whether criminals are using AI tools. “I was recently speaking with a group of intelligence leaders from all over the world, and this was one of the topics of conversation. The reality is, we don’t know the extent to which AI is being used at present. But absolutely, we’d expect attackers to use it,” he says. “Take phishing as an example. Before, you had to speak the language, and you had to do all this reconnaissance and knowledge gathering to create an email that was specific and targeted. So you had to have a lot of resources, and it was not really scalable. Now all that has evaporated because, in theory, you can do targeted attacks at scale using gen AI.”

The other broad risk related to AI—inadvertent misuse by employees, which involves such things as feeding sensitive information into tools and being unaware of the risks of doing so—is the more likely clear and present threat. For businesses moving forward with AI, the lack of employee awareness about the risks of using AI is the number one security or data privacy challenge they are facing in their use of AI, selected by 51% of respondents. **FIGURE 2**


Many public AI tools, which are freely available and can generate content based on user inputs and prompts, are of

FIGURE 2

The AI Enemy Within

Lack of employee awareness is the most common security or data privacy-related challenge

What security or data privacy challenges has your organization experienced in its use of AI? *Select all that apply.*



Base: 205 respondents at organizations moving forward with AI (i.e., exploring or evaluating potential use cases, planning or piloting use cases, or have fully operational use cases). Not shown: 9% Don't know, 8% None, 1% Other.

Source: Harvard Business Review Analytic Services survey, September 2024

high concern precisely because they are free to use, and the content that users feed into these tools can become part of the data set AI learns from. “There have been a few high-profile cases of employees feeding tools with a company’s intellectual property [IP] only for that IP to then be used in the replies given to other users,” CAMS’ Pearlson says. “The company has exposed something it didn’t want to expose, and now the question is ‘Who owns the answer?’”

Given these risks, it’s of little surprise that several organizations appear worried about employees using these tools. Forty-four percent of respondents say their organization is “very concerned” about employees’ use of public AI models or tools, and another 44% are “somewhat concerned.”

The lack of understanding of how the tools work and the security, compliance, and privacy standards each app abides by is creating a dangerous situation where the controls put in place could be ineffective or, worse, amplify the threats. According to Pearlson, “It’s the Wild West, and people are starting to do things because they are scared. Some companies have banned the use of public AI tools, but people can still go home and use them there. Or they can use the tools on their phones, and then the company really doesn’t know what information is going out there.”

Encouraging Responsible AI Use

Tony Anscombe, chief security evangelist at ESET, spol. s r.o., a global cybersecurity company headquartered in Bratislava, Slovakia, shares the opinion that banning the use of public AI tools is not the solution. “It’s better to have controlled access than to ban access, and you need to teach employees what they should and shouldn’t be using the tools for,” he says.

However, a number of organizations are employing blunt tactics such as prohibiting the use of public AI tools. While the most common measure organizations are taking to manage employee use of public AI tools is advising or cautioning those workers against using such tools for work, selected by 47% of respondents, nearly a quarter (24%) have simply shut down access to them altogether. **FIGURE 3**

Pearlson’s research with CAMS focuses on managerial, organizational, and strategic issues in cybersecurity. She asserts that one of the most important things to realize about cybersecurity is that an organization can never be 100% secure. “You could spend large amounts of resources and never really reach total security, so we need another way of thinking about it, and the research we are doing is around the idea of resilience,” she says. “Protection is keeping the bad guys out; resilience thinking assumes they are going to get in and prepares for that so the organization can respond, recover, and get back to operations as quickly as possible with as little damage as possible.”

Building trust by clarifying shared responsibility and creating an organizational culture, in her view, are among the most powerful defenses companies have in terms of reducing the different risks associated with AI use.

“One company we are studying uses the concept of trust as they build processes, structures, and relationships so they are ahead of the potential issues that might arise as we figure out the risks and impacts of AI,” says Pearlson. “Their idea is that everyone is doing the best they can, with limited information, so they need to work together to secure the demarcation points of responsibility. For example, if the data used to train an AI model is poisoned by a hacker, who is responsible for that data poisoning? Instead of blaming the person who owns the data or the team who manages it, or even the team who built the AI system, they will have other ways to respond and recover from the incident.”

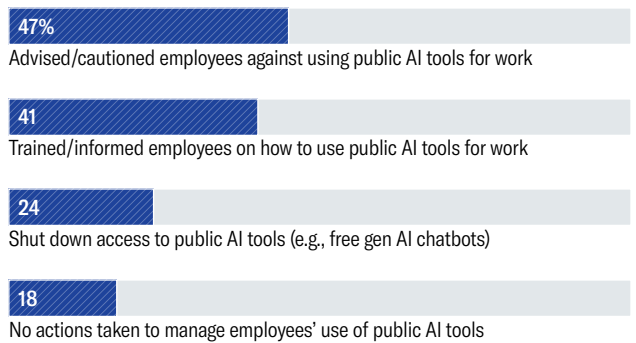
On the matter of organizational culture, Pearlson looks beyond training and awareness programs. “Training is an important component of culture since it sets a baseline, but training alone is not sufficient to reduce human risk. If it was, we wouldn’t still be seeing so many cybersecurity issues,” she says. “What we need to do is infuse everyone in the organization with the feeling that they have some role-appropriate responsibility for helping the company stay secure. To do that we have to change people’s values, attitudes, and beliefs.” Organizations need to take the fear out of having a cybersecurity conversation. Instead, they need to model

FIGURE 3

Controlling Public AI Tools’ Usage

Organizations mostly issue advice or provide training to employees

What measures, if any, has your organization taken to manage employees’ use of public AI tools? *Select all that apply.*



Base: 227 respondents. Not shown: 4% Other, 3% Don’t know.

Source: Harvard Business Review Analytic Services survey, September 2024



47%

**say their
organization has
advised and/
or cautioned
employees against
using public AI
tools for work.**

and reward the right behaviors, from the very highest levels of leadership down, she says.

Protecting the Enterprise

Beyond cultivating a good culture and a sense of shared responsibility, there are many ways organizations can make it harder for criminals to succeed in an attempted attack. Indeed, despite the high levels of organizational anxiety, the executives in charge of cybersecurity are largely excited by AI and willing to make it work, claims Hypponen. “Every time I have these discussions with CISOs or CIOs or other people in charge of security in large organizations, everybody is excited,” he notes. “Everyone sees the benefits, and cybersecurity people themselves say they are using these tools several times every day.”

Businesses are improving their cyber defenses by increasing employee understanding and awareness of the risks and ensuring their data is secure and well managed. Specifically, the two most common efforts organizations moving forward with AI are undertaking to ensure they use AI in a secure way are training employees and improving data governance, each chosen by 57% of respondents. **FIGURE 4** Rolling out or improving cybersecurity awareness campaigns for employees came in next, at 46%.

Beyond training and awareness efforts or work focused on improving data governance, a lot of the common cybersecurity tactics and defenses are not new. Most have been in place for some time; the upsurge of AI has simply amplified the business case for investing in them.

ESET’s Anscombe points out that companies have been on a journey to improve their cyber defense capabilities for years—and not just because of AI. “There are so many other drivers that are pushing the boundary of cybersecurity,” he asserts. “There’s regulation, like data protection laws, or disclosure laws that force public companies to declare when they’ve had a data breach. And you’ve got cyber insurers pushing for higher standards. My point is, there are a lot of factors that have pushed companies to take cybersecurity far more seriously, whether digital transformation, regulation, or new technologies.”

While AI introduces new types of vulnerabilities associated with the tools themselves, such as data poisoning (where a hacker could tamper with the data used to train an AI model and therefore impact the kinds of answers it returns) or model manipulation (where the code of the AI model itself is manipulated for malicious purposes), Anscombe says cybercriminals still need a way to access company systems. In other words, AI models need to be protected in the same way as any of the other data assets in an organization.

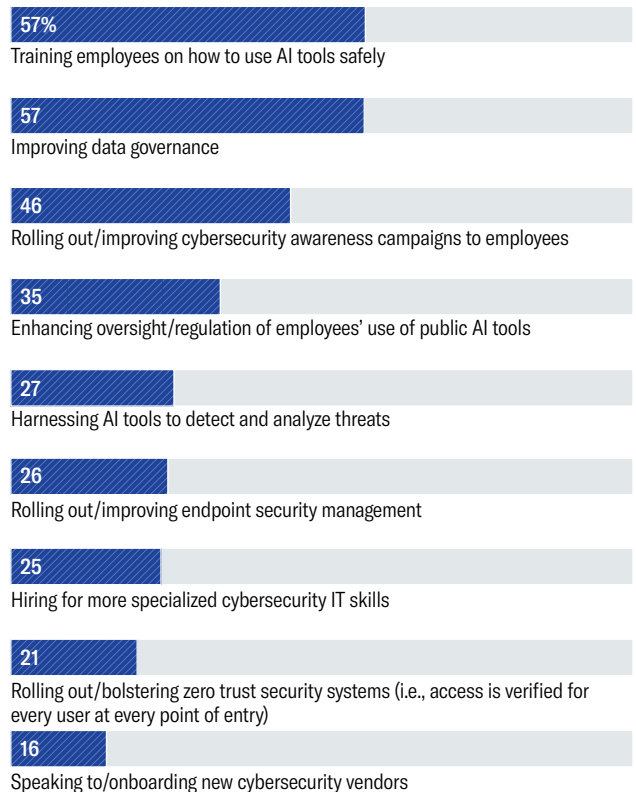
Security strategies like zero trust as well as processes such as endpoint management and identity and access management

FIGURE 4

Different Approaches to Secure Usage

With AI tools, organizations focus mostly on training and improving data governance

What efforts, if any, is your organization putting resources toward/working on to ensure it uses AI in a secure way? *Select all that apply.*



Base: 205 respondents at organizations moving forward with AI (i.e., exploring or evaluating potential use cases, planning or piloting use cases, or have fully operational use cases). Not shown: 8% Don't know, 5% None, 1% Other.

Source: Harvard Business Review Analytic Services survey, September 2024

are, therefore, increasingly important as the numbers of endpoints—physical or virtual devices that connect to a network, such as phones and web-based apps—increase and employees can log in to company systems from anywhere to work remotely.

Here, Anscombe says one of the simplest but most effective approaches to controlling access is multifactor authentication (MFA) or two-factor authentication. “This has been around for years, but we still see companies not deploying it for the wrong reasons, like users not accepting it as a pain point of logging in,” he grouses. “There is no excuse for not having it. A



“Enlightened technology leaders understand it is not technology first; it is business first. We need to find a way to safely use these technologies and invite business colleagues in to be partners in working out how to move forward responsibly together,” says CAMS’ Pearlson.

major data breach that happened in February 2024 happened because the company didn’t have MFA. The total loss of that breach will be over \$1 billion.”

Of course, cybersecurity professionals can use AI itself to improve enterprise defenses, particularly for real-time threat detection as the network environment expands and becomes more complex. For this mission, Hypponen explains that most organizations will need to rely on their security vendors. “You have to work with vendors [that] understand the threats and are getting ready for the future threat landscape,” he explains. “And the most important capability is their reaction time. The faster they can detect an anomaly, the quicker [they] can act.”

The Way Forward

As businesses find a way to move forward with AI—capturing the best aspects of this technology and protecting against the new risks it may introduce—forward-looking cybersecurity executives would do well to remember their role in helping the business create value, says Pearlson. “Enlightened technology leaders understand it is not technology first; it is business

first,” she says. “We need to find a way to safely use these technologies and invite business colleagues in to be partners in working out how to move forward responsibly together.”

The threat landscape is constantly evolving, but there is a real opening for businesses to improve their cyber defenses. The next technology wave beyond AI will surely introduce new types of threats, so knowing how to build the right culture of shared responsibility and heightened risk awareness is critical for equipping organizations for the momentous AI opportunities of today and the technological developments of the future.

“Businesses exist to take risks,” says Khawaja. “The important thing is to not get caught up thinking about every possible risk there is, but rather to focus on the ones that are relevant to your particular AI use case. Once you have identified the specific risks, you can home in on the relevant controls you can implement to mitigate those risks. Understanding the risks in a general sense is about as useful as knowing the average temperature across the country on a given day.”

Endnotes

- 1 World Economic Forum, The Global Risks Report, 2024, January 2004. https://www3.weforum.org/docs/WEF_The_Global_Risks_Report_2024.pdf.
- 2 Deloitte, 2023 Global Future of Cyber survey. https://www.deloitte.com/content/dam/assets-shared/docs/services/risk-advisory/2023/gx-deloitte_future_of_cyber_2023.pdf.
- 3 Exabeam Inc., “AI Cyber Security: Securing AI Systems Against Cyber Threats.” <https://www.exabeam.com/explainers/ai-cyber-security/ai-cyber-security-securing-ai-systems-against-cyber-threats/>.

METHODOLOGY AND PARTICIPANT PROFILE

Harvard Business Review Analytic Services surveyed 227 members of the *Harvard Business Review* audience via an online survey fielded in September 2024. Respondents qualified to complete the survey if they were involved in making or implementing their organization's decisions about IT (information technology), including around topics like AI and cybersecurity.

Size of Organization

32%
10,000 or more employees

30%
1,000–9,999 employees

8%
500–999 employees

30%
50–499 employees

Seniority

31%
Executive management/
board members

39%
Senior management

19%
Middle management

11%
Other grades

Key Industry Sectors

15%
Manufacturing

13%
Financial services

11%
Technology

10%
Government/
not-for-profit

All other sectors
less than 9% each

Job Function

26%
General management

18%
IT

All other functions
less than 7% each

Regions

40%
North America

21%
Europe

21%
Asia Pacific

9%
Middle East/Africa

9%
Latin America

Figures may not add up to 100% due to rounding.



Harvard Business Review

ANALYTIC SERVICES

ABOUT US

Harvard Business Review Analytic Services is an independent commercial research unit within Harvard Business Review Group, conducting research and comparative analysis on important management challenges and emerging business opportunities. Seeking to provide business intelligence and peer-group insight, each report is published based on the findings of original quantitative and/or qualitative research and analysis. Quantitative surveys are conducted with the HBR Advisory Council, HBR's global research panel, and qualitative research is conducted with senior business executives and subject-matter experts from within and beyond the *Harvard Business Review* author community. Email us at hbranalyticservices@hbr.org.

hbr.org/hbr-analytic-services