KPMG

Using generative AI to strengthen cybersecurity

How IT professionals can balance its risks and rewards

kpmg.com

Expectations for generative AI in cybersecurity

Companies have been using artificial intelligence (AI) tools, particularly machine learning, in a growing number of functions and tasks. Generative AI, which has robust natural language processing capabilities and can be used by non-technical employees, promises to make AI significantly more accessible and widely adopted—including in security applications. To better understand how business and technology leaders are thinking about generative AI, in March 2023 KPMG surveyed 300 global executives and senior managers at companies of various sizes across industries. (We conducted a follow-up survey in June and highlight notable differences where relevant.)

Respondents overwhelmingly viewed generative AI as a potentially transformative technology. Two-thirds predicted that generative AI will have a significant impact on their companies within the next three to five years; 31 percent anticipate high impact on enterprise risk management, including cybersecurity, and 64 percent expect moderate impact. And in our June survey, 81 percent of the 200 US respondents ranked generative AI as the top emerging technology that will impact their businesses over the next 18 months.

More than 70 percent of information technology (IT) professionals say they have prioritized cybersecurity application, and 63 percent say generative AI in cybersecurity will likely have the greatest impact on their organization; 64 percent expect to use generative AI in cyber in the next six to 12 months.

Business and IT leaders are also aware of the risks: 92 percent of respondents said there will be moderate to high risks in implementing generative AI. The No.1 risk cited was cybersecurity (54 percent), followed by privacy (53 percent), and liability (46 percent). But in the June survey, 77 percent of respondents said they had confidence in their ability to address and mitigate the risks.

IT professionals have high expectations for generative AI

Say that cybersecurity applications will be a top priority for generative AI adoption

See cybersecurity as the area of greatest potential for generative AI

Expect to implement generative AI applications in six to 12 months

IT/security executives are focused on risk management

Respondents point to data privacy and cybersecurity, and uncertain regulatory landscape among their top areas of concern:

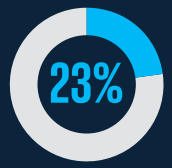

Data privacy and cybersecurity

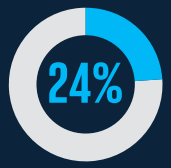

Regulatory landscape

Respondents are moving quickly to strengthen generative AI governance. In March, only 3 percent of respondents reported that they had mature responsible AI governance in place, while in June, 46 percent said they had implemented it:

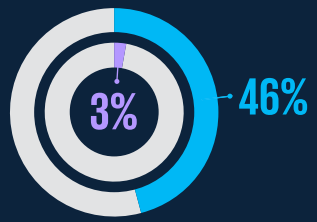

Respondents also firmly believe data security and privacy credentials to be the most important capabilities when considering partnerships:

30%
… were looking to collaborate with external partners

34%
… considered data security and privacy credentials as a key capability while partnering with or investing in other companies for generative AI

Professionals increasingly consider cybersecurity and personal data privacy as priorities for risk management:

53% ▶ 62%
Cybersecurity as a high priority

53% ▶ 63%
Privacy concerns with personal data as a high priority

But as professionals become more familiar with generative AI capabilities, they are gaining greater confidence in their ability to mitigate risks:

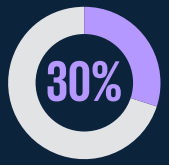

… considered risk mitigation a high concern in implementing generative AI

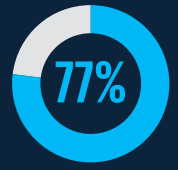

… said they are highly confident in their organization's ability to address and mitigate risks associated with generative AI

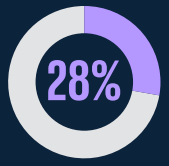

… claimed their company has or will have a comprehensive risk management approach to generative AI

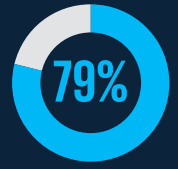

… of respondents who deployed a generative AI solution showed high confidence in their risk mitigation ability

Note: Findings are based on KPMG surveys on generative AI from March and June 2023; the comparative March survey data points are based on 225 US respondents.

March survey June survey

Opportunities for generative AI-enabled cybersecurity

Like other top leaders in their organizations, Chief Information Security Officers (CISOs) are looking into how generative AI can be effectively applied in their functions. While it is early days, we see that the new technology shows promise in four areas:

1. Forensics and response
2. Security operations
3. Identity and access
4. Third-party supply chain management

1 Forensics and response

Threat detection and rapid response

One of the biggest challenges in cybersecurity is detecting and responding to threats quickly and accurately. Generative AI can synthesize, interpret, and report on large amounts of data from multiple different sources—network logs, endpoints, and applications, for example. Where generative AI could really make a difference is its ability to rapidly gather and summarize data in response to natural language queries.

The generative AI program can incorporate knowledge of a vast range of third-party application terms (such as error codes) as well as internal organization-specific definitions and process protocols (such as how a "log-in" is defined or what signal constitutes a particular level of risk). It can access information, research, and generate reports more efficiently, rapidly, and accurately than any manual method, freeing analysts to further interpret data, drive better responses, and improve security. A human with experience, intuition, and insider knowledge is still essential for picking out hard-to-detect anomalies. But generative AI provides a huge head start.

In addition, generative AI can be used for alert triage. For example, based on past incidents, a generative AI application could create a predefined checklist that an analyst could review before deciding the best course of action. And the application could function as a security advisor, pointing out how such threats have been dealt with in the past by others, whether inside or outside the enterprise. Finally, generative AI could be used to validate manual responses. It is already being integrated into leading endpoint detection and response technologies that largely automate the efforts of security operating center (SOC) analysts.

Phishing and fraud prevention

Generative AI can analyze and screen points of entry, such as emails, to identify and flag suspicious patterns, including unusual sender addresses or content that includes signals suggesting that they are not authentic. Generative AI can also analyze user behavior and detect fraudulent transactions in real time. Another application could be analyzing accounts payable records to prevent unauthorized payments.

Insider threat detection

Generative AI can be used to monitor employee behavior and report anomalies that may indicate a potential insider threat. By auditing and quantifying user activity, such as login times, file access patterns, and network activity, generative AI can learn what is normal for each user or class of user, and detect and flag any unusual behavior that may indicate a threat.

2 Security operations

Vulnerability management

Various AI methodologies can be deployed to identify vulnerabilities in software and systems through which cybercriminals can gain access. Generative AI can then add value as a digital advisor on the best ways to respond. For example, if a vulnerability ticket is assigned for remediation, you may rely on vendor tools to identify appropriate fixes, patches, or configurations to their applications. But what is often missing is the organization-specific context—what groups must provide approval? Is there a designated URL destination for generating an exception request? Which internal team should be engaged based on the location of the vulnerability? Generative AI can function as a chat bot advisor, capable of adding the organization-specific content to the vendor details, bringing all the pertinent internal and external signals together in a prescriptive way. It can thereby recommend tailored best-practice remediation steps.

Attack surface management

Monitoring the perimeter—firewalls, web servers, transactions within the publicly accessible cloud, and so on—is a huge and constantly changing challenge. Generative AI can be deployed to parse structured data sets, report on what is included in the perimeter surface, identify the biggest potential surface threats, and assess the degree of risk in the attack surface at any given time. Generative AI can be used to establish a baseline and flag what has changed. It can also determine what information is lacking to help identify and fix external vulnerabilities, misconfigurations, and other weaknesses.

Metrics, dashboards, reporting

Generative AI can be very useful in turning data into insights. A risk analyst, for example, can write a prompt such as: "What are the three areas that have changed in my environment this week from last year at this time?"; or "Is my risk posture improving?" Generative AI can then compile internal and external information to provide the update as well as recommendations and guidance. For example, it could compare how the company performs event logging with best practices published by the National Institute of Standards and Technology and identify gaps. It can also identify which security tools are missing (according to industry best practices) and recommend where to invest in new tools. It can answer questions about performance over time, compare performance over specific intervals, and generate a monthly CISO scorecard. Overall, generative AI can serve as a security advisor for detection and response to cyberattacks or cyber incidents, and make recommendations for improving processes. Finally, generative AI can help create dashboards with graphical interpretations of data for rapid assimilation of information and can generate reports synthesizing data from multiple internal and external sources.

3 Identity and access

Least-privilege security

Least-privilege security ensures that access to systems is limited to individuals doing a particular job at a particular time. Despite significant investment, least-privilege provisioning remains largely manual, highly discretionary, fairly static, and not well monitored. Generative AI has the potential to automate this process. For example, a governance platforms with generative AI could be trained on large models to assess job function, title, physical location, etc. to establish the baseline for preliminary provisioning, amended provisioning, or deprovisioning access by individual or class. Behavior anomalies can be flagged for human intervention and follow-up. Security teams can now manage access privileges for tens of thousands of identities, across hundreds of applications and multiple types of access, and implement continuous access management on a huge scale.

Identity protection

Impersonation and other misuse of digital identities can be forestalled through generative AI applications that research and report patterns of behavior, identifying signals that point to fraud and even specific types of cyberattacks or vulnerabilities. Suspicious behavior and anomalies in normal patterns of engagement can be flagged for investigation.

4 Third-party supply chain management

Inherent risk profiling

Organizations may do business with thousands of third-party providers—security program providers among them—potentially sharing sensitive data in the process. To ensure that their organization is protected from third-party risk it is essential to understand the risk posture associated with any supplier, which has been a cumbersome and time-consuming task. Due to time and budget constraints, security teams typically have created robust security profiles for a small share of partners.

With generative AI automating the process, companies can go back and perform risk assessments on all partners and use the tool to vet new ones. Additionally, security and risk profiles that have been deemed satisfactory can be extrapolated to create security checklists for future profiling as well as to inform contractual terms and conditions for engagement.

Ongoing risk management

Generative AI tools can be used to monitor compliance with agreed security protocols, the status of any required mitigation strategies, and the impact of changes that alter the risk profile. Has the supplier suffered a breach? Is a partner being acquired? Are its platforms undergoing significant alteration? Have they added new technology partners? All these what-ifs affect the risk profile and compliance. Automated, ongoing due diligence can track what is changing with every vendor in the supply chain and assess potential risk. In our June survey, those who had already deployed generative AI solutions had a greater confidence in their ability to mitigate risk.

Barriers to adoption

While the promise of generative AI technologies is persuasive, CISOs and other business leaders recognize the significant barriers to its adoption. As a relatively new technology, there are many unknowns. The biggest concerns at this early stage are exposure of proprietary and customer data, loss of control of intellectual property, inaccurate results, and intentional—even criminal—misuse—to produce deepfakes to commit fraud, for example. As generative AI tools and applications are adopted widely across the organization, the risks multiply. Responsible use of AI, clear guidelines and governance, and accountability are essential—and CISOs can play a part in making sure the guidelines are understood and followed.

Secure a safe start

As we have shown, generative AI will likely have many applications in cybersecurity. We believe that CISOs that use generative AI to improve security processes will have an opportuniuty to better protect their organizations. However, the benefits of being an early adopter must be weighed against the risks and unknowns.

KPMG has identified five steps for successful deployment of generative AI applications for cybersecurity (once the overall strategic approach has been determined). Following these can help ensure that the organization derives the most value from generative AI investments.

1 Align on specific use case(s)

Define desired outcomes of the pilot and prioritize specific area(s) of focus

2 Gather relevant data

Identify relevant data sources and validate that threat data is accurate, complete, and up to date

3 Train the model

Train the AI model(s) to understand the data and perform the desired analysis and tasks

4 Test and refine

Test performance and refine algorithms by comparing outputs to those of humans or existing analytical models

5 Integrate, monitor, and scale

Integrate workflows via APIs, monitor performance, and evaluate opportunities to scale

How KPMG can help

KPMG can help you apply the five-step approach to get your generative AI implementations off to a sound start and fine tune initiatives already underway. We understand the multidimensional security challenges faced by enterprises across their technology, people, processes, and partners. Our teams of technology professionals are skilled and experienced with a diverse array of leading technologies, platforms, and modern development practices, including generative AI applications, data science, and cybersecurity. We are well-versed in the organizational changes that are required to extract maximum advantage from generative AI while protecting against its potential downsides. And we can help curate risk mitigation initiatives from conceptual solution to deployment roadmap, providing a nonstop solution for cyber resiliency.

Whether it's helping you lead a cybersecurity, business continuity, or digital transformation, KPMG creates tailored, data-driven solutions that help you safeguard security, deliver value, drive innovation, and build stakeholder trust.

Contact us

Matt Miller
Principal, Advisory
Cyber Security Services
matthewpmiller@kpmg.com

Katie Boswell
Managing Director, Advisory
Cyber Security Services
katieboswell@kpmg.com

Jim Wilhelm
Principal, Advisory
Cyber Security Services
jameswilhelm@kpmg.com

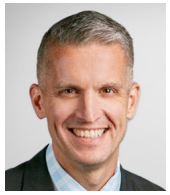

Douglas LaGore
Principal, Advisory
Cyber Security Services
dlagore@kpmg.com

David Nides
Principal, Advisory
Cyber Security Services
dnides@kpmg.com

Ryan Budnik
Director, Advisory
Cyber Security Services
rbudnik@kpmg.com

Brian Marks
Director, Advisory
Cyber Security Services
bjmarks@kpmg.com

Diana Keele
Director, Advisory
Cyber Security Services
dkeele@kpmg.com

Related thought leadership:

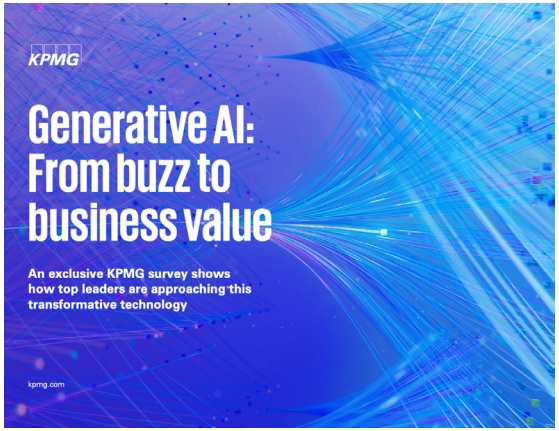

Learn how KPMG can help make your generative AI implementation successful, and explore how we can help you adopt AI in a safe, trustworthy, and ethical manner.

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

DASD-2023-12932.

August 2023

kpmg.com/socialmedia