RSAConference™

CISO PERSPECTIVES

How Top CISOs Are Navigating the Risks and Opportunities of GenAI

Exclusive Insights from the RSAC Executive Security Action Forum (ESAF)

A Message from the RSAC ESAF 2024-2025 Program Committee

As a forum for candid discussion among peers, ESAF sessions are confidential, invitation-only, and limited to a select group of senior-most information security and risk executives. Through the "CISO Perspectives" series of reports and blogs, we are sharing some of our hard-earned wisdom with the broader cybersecurity community. Our goal is to help all organizations improve the management of cyber risks.

For more information on ESAF and links to our previous reports, see *www.rsaconference.com/ESAF*

The RSAC ESAF 2024-2025 Program Committee includes:

Rich Agostino, Senior Vice President, Chief Information Security Officer & Infrastructure, Target
Brad Arkin, Executive Vice President, Chief Trust Officer, Salesforce
Jason Barnett, Vice President & Chief Security Officer, HCA Healthcare
Ben Brophy, Group Chief Information Technology and Security Officer, Reckitt
Sherron Burgess, Senior Vice President & Chief Information Security Officer, BCD Travel
Tim Callahan, Senior Vice President, Global Chief Information Security Officer, Aflac
Josh Davis, Group Vice President and Chief Cyber Security Officer, Toyota Motors North America
Michael Higgins, Vice President and Chief Information Security Officer, L3Harris
Gary Harbison, Senior Vice President, Global Chief Information Security Officer, Johnson & Johnson
Aaron Hughes, Senior Vice President and Chief Information Security Officer, Albertsons Companies
Susan Koski, Executive Vice President, Chief Information Security Officer, PNC Financial Services
Andy Ozment, Executive Vice President, Chief Technology Risk Officer, Capital One
Chris Porter, Senior Vice President and Chief Information Security Officer, Fannie Mae
John Scimone, President and Chief Security Officer, Dell Technologies
Carolann Shields, Senior Vice President, Global Chief Information Security Officer, 3M
Emma Smith, Chief Information Security Officer, Vodafone

View Program Comittee bios **here**.

RSAC™
Executive Security Action Forum
A Community of Fortune 1000 CISOs

About RSAC ESAF

The Executive Security Action Forum (ESAF), an RSA Conference (RSAC) community, has been a trusted forum for Fortune 1000 Chief Information Security Officers (CISOs) since 2003. Led by a program committee, the community shares information at confidential sessions throughout the year and at our annual meeting at RSA Conference, enabling security leaders at some of the world's largest enterprises to collaborate and find actionable solutions to common challenges.

Contents

The Risks of Rapid Generative AI Adoption

How are CISOs at Fortune 1000 Companies Approaching Generative AI (GenAI)?

Recently, GenAI has been one of the most intensely-discussed topics in the RSAC Executive Security Action Forum (ESAF) community of Fortune 1000 CISOs. The confidential nature of ESAF sessions enables participants to speak candidly about their firsthand experiences.

In this report, we share highlights of those conversations. For ESAF CISOs, GenAI creates three security imperatives:

1. **Protect the use of GenAI**
2. **Protect with GenAI**
3. **Protect against GenAI**

You'll read about their early strategies to help their companies deploy GenAI safely across the business, to build better defenses with GenAI, and to confront GenAI-enabled threats.

We've summarized their discussions and anonymized details to preserve confidentiality. We also share results from RSAC's recent Fortune 1000 survey on GenAI.

What is Generative AI? (GenAI)

Generative AI (GenAI) is a form of artificial intelligence that can create new content such as natural-sounding text, audio, images, and code. GenAI algorithms are trained on large amounts of data. When given a natural-language prompt, GenAI systems predict what output a human would give, based on their training data.

Business Use Cases for GenAI

At most large enterprises, GenAI is rapidly being adopted for business applications. In fact, RSAC's recent Fortune 1000 survey on GenAI showed that 100% of companies plan to implement GenAI in the next 12 months for a range of business use cases.[1]

The top use cases that companies plan to focus on are in employee productivity, software developer productivity, and customer service (see Figure 1). However, priorities could change in the coming months as they evaluate a growing number of use cases.

What are the top three business use cases for GenAI that your company will focus on implementing in the next 12 months? (not including security use cases)

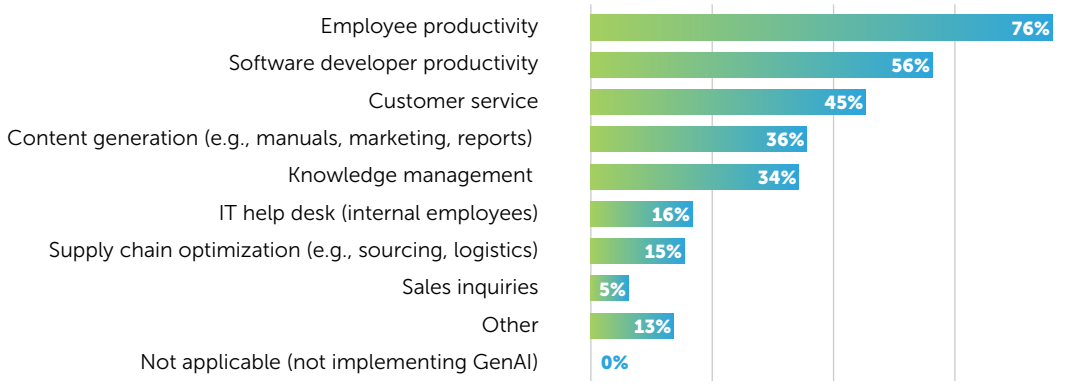

Figure 1: Top Business Use Cases for GenAI

[1] Survey of 100 Fortune 1000 CISOs conducted by RSA Conference for an internal research study in Q2 2024.

Top Security Risks of GenAI Adoption

As GenAI is being rapidly adopted, CISOs play a key role in enabling their companies to attain the benefits of GenAI while effectively managing the risks. Different use cases for GenAI introduce different types of risk. For example, without adequate controls, an IT help desk system might grant access to unauthorized users, or a customer service chatbot might disclose confidential company information.

Given the range of risks that companies are facing, what risks are the most concerning to Fortune 1000 CISOs? The recent RSAC survey shows their top concerns are data leakage and unvalidated or hallucinated data being used in decision-making (Figure 2). "Hallucination" refers to the phenomenon in which GenAI systems produce information that sounds plausible but is actually incorrect or false.

Potential for Unmanaged Risks in GenAI Adoption

In their recent discussions on GenAI, ESAF CISOs emphasized the potential for unmanaged risks. In addition to formally managed GenAI projects, GenAI is being used across an organization's ecosystem. Tools such as ChatGPT are easy for individual employees to use under the radar and GenAI features are being added to many business and consumer software applications.

Organizations also face risks as suppliers and partners adopt GenAI tools for their business operations. For example, enterprise data that is shared with third parties could be at risk of leakage if a third party puts the data into GenAI systems.

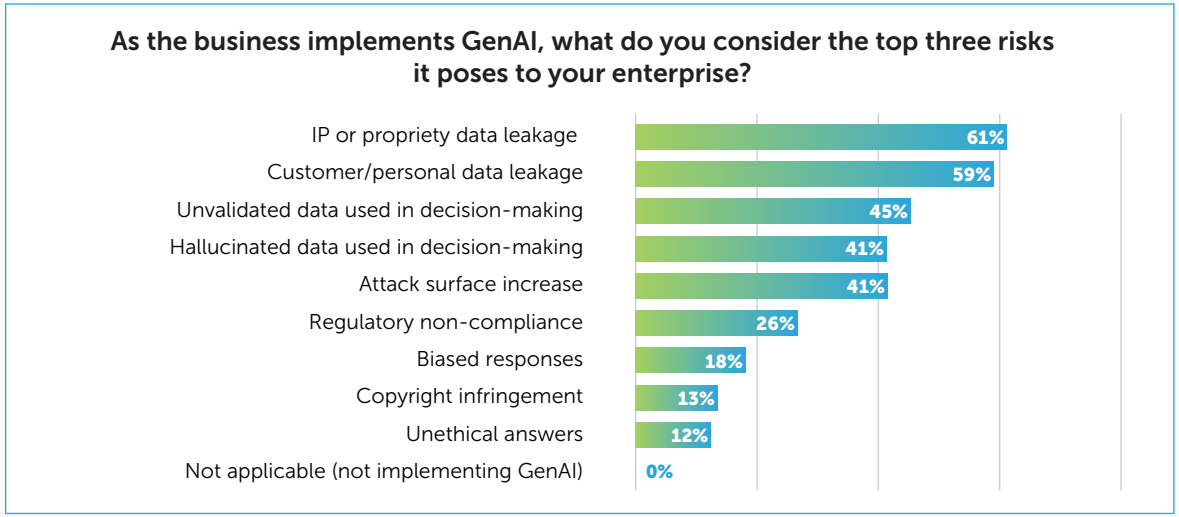

Figure 2. Top Risks as Enterprises Implement GenAI

Foundational Principles and Policies for Safe GenAI Use

Since GenAI is so easy for individual employees to start using, mass-education regarding principles and policies is critical. ESAF CISOs discussed ways to provide an environment where employees are encouraged to use GenAI while setting a clear tone around what they are—and are not—allowed to do.

Early steps that their companies are taking include:

- **Developing a document on the ethical use of AI** which states the principles that all employees are responsible for following, i.e., AI usage should be beneficial, equitable, transparent, responsible, and accountable.
- **Creating clear communications and training** around acceptable use of public GenAI tools, for example:
  - Requiring that no proprietary, confidential, or sensitive information should be submitted to public GenAI tools without corporate approval.
  - Outlining which large language models (LLMs) to use for which purposes and in some cases blocking access to LLMs that have not been assessed.
- **Leveraging big names in the company** to get the messages across, for example:
  - Having the CIO and CTO send a joint email to all employees outlining the policies.
- **Implementing ways to display a reminder message** about appropriate use when employees visit public GenAI websites and keep logs of all usage.

Evaluating Third-Party Use of GenAI

ESAF CISOs also discussed augmenting third-party risk assessments to evaluate AI capabilities in products and services. If a vendor has added AI capabilities, they should:

- Provide sufficient documentation to assess the AI's capabilities and limitations.
- Explain the provenance of the data used to train the AI.
- Demonstrate the effectiveness of the risk management measures they put in place.

Frameworks for assessing AI in third-party products are emerging. One example is the US Federal Government's Office of Management and Budget Policy on AI which includes a section on "Managing Risks in Federal Procurement of AI."

Third-party risk assessments should also verify that if enterprise data is shared with a third party, no enterprise data will be put at risk through the third-party's use of GenAI tools.

New and Emerging AI Regulations

Another fundamental aspect of managing the risks of GenAI is ensuring the company's use of GenAI complies with regulations. The global regulatory landscape is evolving quickly. Examples of legislative initiatives that large global companies are tracking include:

Jurisdiction	Legislative Initiatives	Timeline
Brazil	Proposed AI regulation: Bill of Law 2338	Introduced May 2023.
Canada	Proposed Artificial Intelligence and Data Act	Introduced June 2022. Amended November 2023.
China	Various regulations, e.g., Generative AI Measures and Ethical Review Measures	Effective August 2023 and December 2023.
EU	Artificial Intelligence Act	Entered into force in August 2024.
US Federal	Executive Order on AI	Signed in October 2023.
US States	17 states have enacted bills	Proposed mostly in 2024.

Key Takeaways

- CISOs' top concerns with the business's use of GenAI are data leakage and unvalidated or hallucinated data being used in decision-making.
- Fast-paced GenAI adoption in the extended enterprise raises the potential for unmanaged risks. Besides formally managed GenAI projects, GenAI tools are easily used by individual employees and third parties.
- Foundational elements of GenAI risk management include training employees on acceptable use of GenAI, assessing third party use of AI, and paying close attention to emerging AI regulations.

Insights on Generative AI Governance

What are Fortune 1000 Companies Doing to Govern the Business's Use of GenAI?

In this section, we look at what it takes to govern the rapid, and often chaotic, adoption of GenAI in the enterprise. A leading CISO shared how their company is assessing and prioritizing GenAI use cases.

Corralling All GenAI Proposals through a Review Board

At most large enterprises, the CISO and the security team are a key part of a cross-functional effort to oversee the company's GenAI initiatives. In fact, according to the recent RSAC survey of Fortune 1000 CISOs, the security team is involved in GenAI policy and governance at 88% of companies and security team members are on the GenAI oversight committee at 74% of companies.

What's involved in GenAI oversight? As the CISO described, a key element is establishing a review board to assess GenAI proposals. Their company's GenAI use case review board has members with expertise in a wide range of functions including privacy, ethics, environmental sustainability, human resources, legal, and multiple aspects of security, and is led by a Chief AI Officer.

Assessing the Risks of GenAI Use Cases

In looking at the proposed GenAI use cases, the review board assesses risks from many angles including:

Risk area	Examples
Security and privacy	• Infrastructure vulnerabilities • Model stealing • Data poisoning • Intellectual property theft • Privacy rule constraints
Societal and ethical	• Ethical dilemmas • Dehumanization • Misinformation and deepfakes
Regulation	• Regulatory lag (investing in projects that later face restrictions) • Transparency and accountability
Business Resiliency	• Overdependence on AI
Economic and employment	• Job displacement
Environmental	• Resource consumption • Electronic waste
Psychological	• Mental health

The CISO noted that AI project assessment is broader than traditional assessments, and the security team must be trained to evaluate the additional facets. Besides risk, their GenAI review board assesses projects based on:

- **Business Value**: Looking at cost, benefits, scope, and urgency.
- **Process**: Must be clearly defined and a good candidate for AI.
- **Resource**: Sufficiency of AI implementors, subject matter experts, and data source owners.
- **Data**: Examining availability, quality, completeness, accuracy, and classification of data.
- **Technology**: Must have adequate technological infrastructure.
- **Explainability**: Must be transparent across processes and functions to build end-user trust.

Prioritizing GenAI Use Cases

When the company mandated that all GenAI use cases must first be cleared by the review board, proposals came flying in from across the business including operations management, R&D, manufacturing, marketing, customer service, HR, and finance.

In reviewing 800+ proposals, the review board identified an "unhealthy degree of chaos," with company resources being expended without coordination or strategic direction. People were spending time and money on projects that were not feasible or too risky.

Ultimately, the review board prioritized the projects that were most aligned to the company's core strategic areas, about 130 projects. Resources were reallocated to ensure effective and efficient implementation.

After the initial wave of proposals, the average turnaround time to review a project is now about a week. Almost always, the review board says yes, with provisions. It provides detailed, prescriptive feedback on what must be done to make the use case policy compliant.

AI Data Governance

Since the company considers data integral to the success of any AI strategy, they have paired use case oversight with data governance. The use case review board works in tandem with a data stewardship board, which sets data standards and taxonomy for the company's key datasets.

AI Security Standard

Another aspect of the company's risk management strategy is ensuring GenAI system developers adhere to the corporate AI security standard. Their document is based on reference architectures and critical vulnerabilities of GenAI that are published in sources such as the OWASP Top 10 for LLMs list.

The security standard document must be frequently updated as new techniques for attacking and misusing GenAI systems arise. The CISO emphasized that the security team requires the developers to continually reference the security standard, and to self-attest to compliance with the *latest* version.

> *"We didn't know what to expect, because we really had no central process or visibility inside of the company to have a sense of what was happening GenAI-wise. And we just got flooded with use cases...Even in departments that one would expect to have low technology acumen, there was somebody doing something with GenAI, exploring how they could make their business process or function better."*
>
> **— ESAF CISO**

Tips for Securing Generative AI Systems

How do Fortune 1000 Companies Build Security into GenAI Systems?

Here we look at the experience of a company that got out of the gate early and has already rolled out successful GenAI initiatives. Their CISO shared how they built security into a new internal customer service bot which serves their global workforce of over 45,000.

System Overview

The internal customer service bot was developed using a commercial GenAI chatbot platform and trained on company-specific IT workflows and processes. It assists employees with tasks such as approval requests for access to resources and helps resolve technical issues such as trouble with software applications or computer hardware. The system runs in a managed cloud and integrates over 20 enterprise resources and 100 external knowledge resources.

No personal, regulatory, or customer data is allowed in the system. However, because the system can be used for sensitive tasks such as changing forgotten passwords, strong security is critical.

Generative AI Assessment

The CISO emphasized that, in addition to verifying the GenAI system is fit for use, it's necessary to assess its *trustworthiness*. The project team evaluated the GenAI platform against the company's AI ethical framework, a set of criteria for assessing AI. Their questions included:

- Are the answers that the system provides understandable, transparent, and auditable?
- To what degree are the answers accurate? Is the level of hallucination acceptable for the task?
- Is it resilient? Does it perform consistently in unpredictable scenarios?
- Does it behave responsibly, not insulting anyone or showing bias?
- Is the training dataset secure against data poisoning attacks?
- Considering how critical the decisions are and how much the system can be trusted, would human-in-the-loop or human-on-the-loop controls be appropriate for certain workflows?

Using an ethical AI framework provides:

- Structure for responding when board members and customers ask CISOs how they know the company's use of GenAI is trustable.
- Methodology for assessing GenAI against increasingly stringent government regulations.

Securing Generative AI Data Flows and Access

Aside from the ethical issues, the CISO said securing GenAI systems is similar to securing SaaS systems. However, security teams should pay particular attention to the following areas:

Access Controls	To reduce the risk of prompt injection attacks, use data classification and access controls to ensure that the system cannot give out information that the user does not have access to.
API Security	GenAI systems involve many connection paths, any of which can be attacked. Secure and test all API connections.
Authentication	Attackers are directly targeting GenAI systems and are bypassing MFA. Ensure integrations are configured correctly. Monitor for out-of-bounds behavior, e.g., API calls coming from unexpected places.
Data Protection	Design the system to minimize data persistence.
Logging	Log events from every integration point and correlate them in your SIEM. Look especially for out-of-band and out-of-bounds requests and reuse requests. If necessary, work with your vendors to capture the necessary logs. Make sure you can see all data that is sent to the vendor.
Testing	The security team found a critical security flaw in the chatbot platform that made it vulnerable to clickjacking attacks. They alerted the vendor and implemented compensating controls until the vendor fixed the problem. The CISO emphasized, "This is why you do the work. The devil is in the details."

Ongoing Model Tuning

The models require constant tuning to ensure they are effective and trustworthy. Otherwise, they tend to become less accurate over time, a phenomenon known as *model drift*.

To enable tuning, the CISO recommends regularly collecting data on how the chatbot is being used and how accurate answers are. This includes collecting and analyzing the users' prompts to the chatbot and their feedback ratings on the answers and providing feedback to the chatbot platform vendor so that the vendor can assist with model and algorithm enhancements.

Project Outcomes and Future Directions

Full ROI metrics for the GenAI-powered chatbot will require more data, however a positive sign is a 40% reduction in help desk calls in the first four months after rollout. Accuracy for the most common tasks is above 90%. The company plans to expand the system from 20 to a total of 70 integrations.

According to the CISO:

> *"We have to continue to stay on the forward-leaning side of this. Otherwise, the demand will outstrip our ability to do this in a safe manner. We need to manage both the technical baseline and the expectations of responsible use."*
>
> **— ESAF CISO**

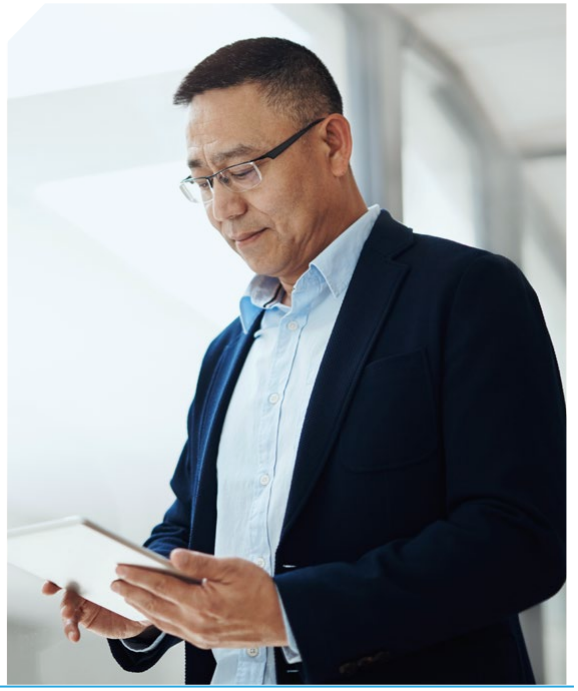

Transforming Security with Generative AI

How do Fortune 1000 Security Teams Intend to Use AI?

This section explores how leading CISOs plan to use GenAI to transform security. We consider what use cases are feasible now, and what use cases might realistically emerge in the next few years. And we've got some more insights from RSAC's Fortune 1000 survey on GenAI.

Enterprise Adoption of Generative AI for Security

According to a recent RSAC survey of Fortune 1000 CISOs, 72% have already implemented GenAI for security and 98% plan to in the next 12 months. However in interpreting these results, ESAF CISOs emphasized that in the next 12 months, the scale and scope of most implementations will be limited. Most security organizations will still be in the experimentation phase with GenAI.

Threat detection is currently the most common use case and prioritizing vulnerabilities is expected to be the fastest-growing use case (Figures 3 and 4). Given how quickly the GenAI space is evolving, these plans may change in the coming months as security teams continue to evaluate use cases.

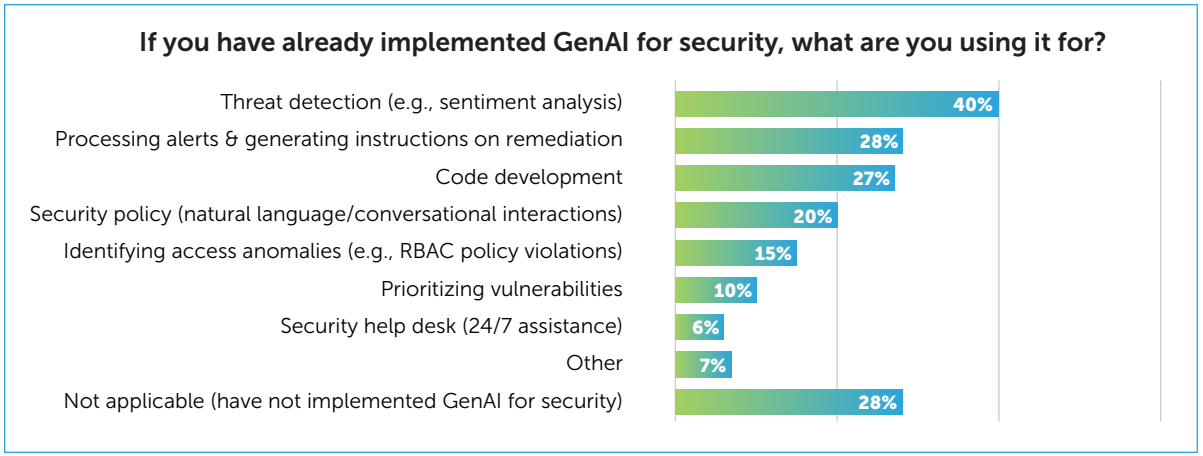

Figure 3. Current Security Use Cases

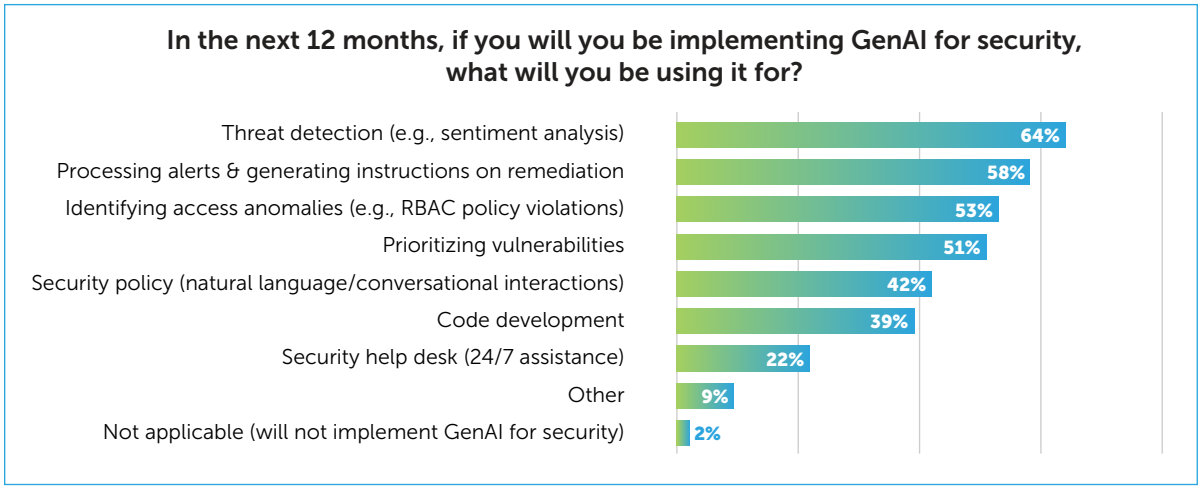

Figure 4. Upcoming Security Use Cases

Identifying Priority GenAI Use Cases

Since security workflows vary considerably between companies, each security team will need to determine the best GenAI use cases for their own organization. For example, one of the CISOs from the ESAF community shared how their company's security team undertook a research project to identify and prioritize use cases.

Over five months, the team conducted interviews across the security organization, asking questions such as, "What problems are so difficult and resource-intensive that we're not doing them well today?" and "Realistically, which of these processes could we make more efficient with GenAI?"

The team also examined the rate of technology development in the GenAI space and made rough estimates of when various technologies might be ready for adoption.

Based on the results, they project that in the next five years, AI could have the largest impact on their security organization in areas that include:

Areas in Security	Examples of GenAI Use Cases
Business Operations	Produce content such as training materials and handle basic project management tasks.
Incident Monitoring	Help perform faster threat detection and automation of common tasks.
Security Assessment	Assess compliance with requirements.
Security Remediation	Determine the potential implications of vulnerabilities and add context to support remediation guidance.

They identified the **highest-priority use case** as "giving real-time security advice and information to employees." Employee assistance chatbots are a common early GenAI project.

AI Impact on Future Jobs in Security

The company's research shed light on a highly emotional aspect of GenAI: the future of jobs. For most security roles, they expect GenAI to have little impact in the medium term, mostly in relatively narrow areas. The CISO felt it was important to communicate this message to allay fears around job security. Given expected rates of GenAI technology development, the research team projected reduced costs of 10% across their converged security program.

Another CISO shared their security staffing strategies vis-a-vis AI. Their team is counting on AI to help ease the talent shortage in security. Even though they expect the volume of work to go way up in the next few years, they plan to keep headcount at current levels and use technology to increase the team's capabilities.

That CISO also shared two examples of use cases they are exploring:

- Red teaming: Using AI to come up with novel ways to penetrate the enterprise perimeter and suggest these tactics to the red team.
- Change management: Using AI to speed up the approval process for system design changes. With AI analysis, they hope to reduce the turnaround time from 30 days to one day.

Challenges to GenAI Adoption

In ESAF discussions, CISOs are generally optimistic about the potential for AI to make security operations more powerful over the next 5 to 10 years. But currently, GenAI adoption in security is hindered by significant barriers such as immature tools, lack of skills, and risk of data leakage, as indicated by the RSAC survey (Figure 5).

Multi-Year Journey

Most ESAF CISOs see the implementation of GenAI as a multi-year journey, as the tools mature and security teams integrate GenAI into their operations. Few CISOs have indicated they are drastically changing their security strategies in the short term. However, board members and other corporate leaders often convey a strong sense of urgency to do more with AI.

CISOs may be under pressure to make major strategic changes immediately even when they feel evidence supports more gradual change. CISOs must be prepared to communicate about the levels of certainty in predictions around AI, justify their security budget allocations, and manage expectations.

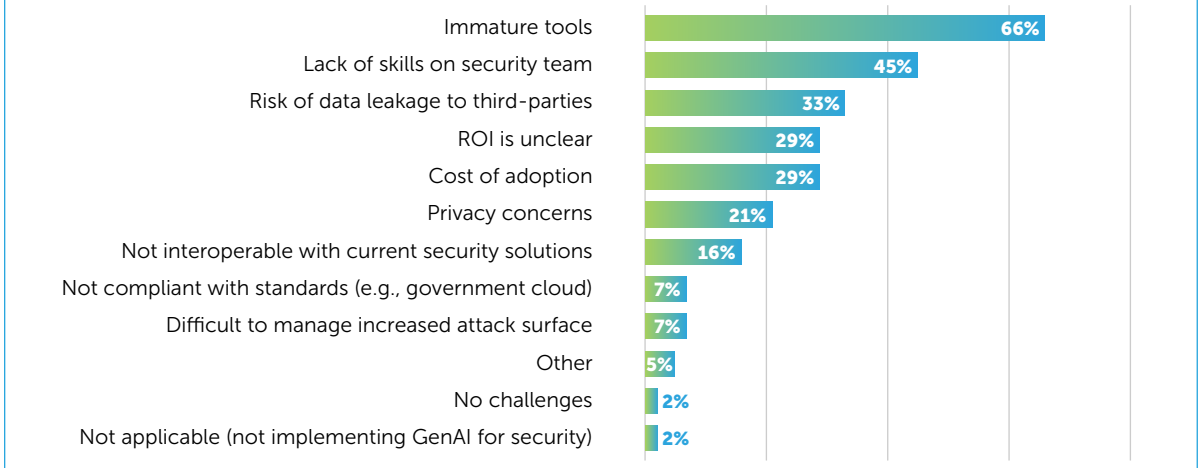

Figure 5. Top Challenges to GenAI Adoption for Security

Tackling the Rise of AI-Powered Cyber Attacks

What are Fortune 1000 CISOs Doing about the Increasing Use of AI in Cyber-Attacks?

CISOs are having a lot of conversations about this topic. It has even become a staple at board meetings, according to the RSAC ESAF community. In this section, we share how CISOs at the forefront of AI plan to protect against AI-enabled threats by building AI-powered defenses.

Common AI Cybersecurity Threats

Attackers are already using AI. In a recent survey of Fortune 1000 CISOs, RSAC found that 72% said they have already seen threat actors use GenAI against their enterprise (Figure 6).

For the CISOs in the ESAF community, they know that 5 to 10 years from now, they will face a vastly different threat landscape because of AI-enabled threats. It is difficult to predict precisely what will change, or when, therefore they must plan amidst uncertainty.

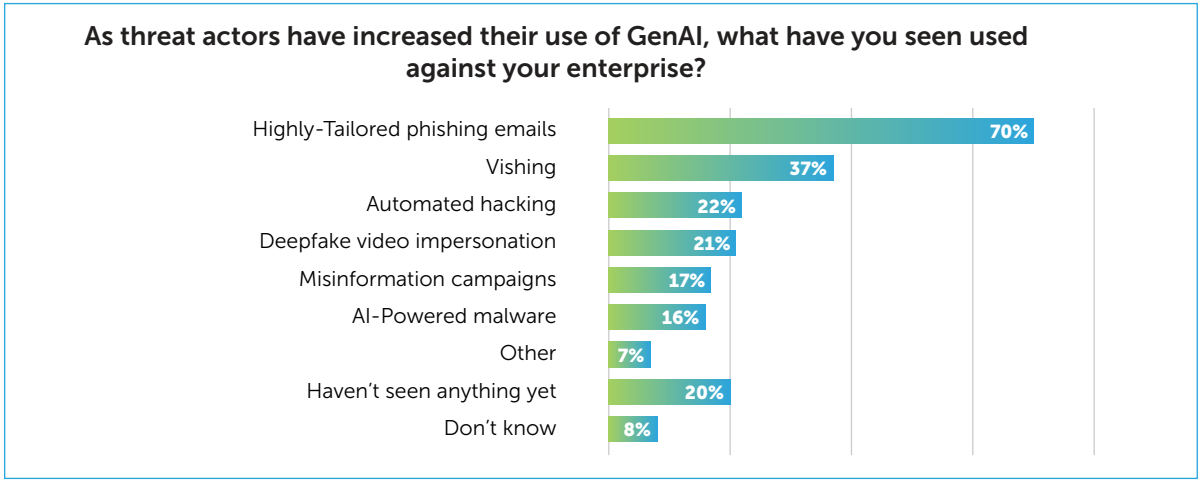

Figure 6. AI-Enabled Threats

Top Concern is AI-Powered Automated Hacking

Of the emerging attack techniques, a top concern is automated hacking. Powered by AI, attacks could unfold much faster than human security teams could respond.

Currently, sophisticated attacks are generally conducted in stages that can take days, weeks, or months. This gives security teams a window of time to detect an attack and contain it. At some point in the future, with AI-powered automated hacking, an attack might be completed in milliseconds.

To defend against this kind of attack, the end goal is to deploy defenses that respond instantaneously to match the speed of the malicious actor's AI-driven attack. Human response is too slow; security teams will need an autonomous system working independently and reacting instantly to threats. Getting there will require building increasingly more powerful AI-enabled defenses over time.

Autonomous AI

A machine-versus-machine matchup requires security teams to achieve autonomous AI systems. One CISO said they consider AI systems to have four levels of capability: analysis, assistance, augmentation, and autonomy. In the near term, they see most security teams deploying AI systems with analysis and assistance capabilities. Higher-level capabilities are longer-term goals.

Building an AI-Enabled Security Organization

Defending against AI-enabled threats requires building AI-enabled defenses. ESAF CISOs are taking steps to apply AI across security processes including:

Research

Look at your security processes to determine where you should apply AI first and where you will get the greatest benefits.

AI Training

Provide basic AI training for everyone and advanced training for specific technical teams. One CISO described "tons of training" as the number one way to prepare the security team for its mandate of becoming an AI-enabled organization.

> *"Getting people ready is going to be the most difficult and the most impactful thing to get right. This cultural transformation needs to happen to get people excited, confident, and enabled to operate in this environment."*
>
> **— ESAF CISO**

- Training for the whole team includes building a basic level of proficiency around AI. It also addresses the fear around job security, since it lays the groundwork for developing skills that will enable workers to transition into new roles.
- Training for team members involved in securing AI includes threat modeling and secure design principles for securing AI data platforms, ML models, and AI-enabled applications.
- Training for team members involved in implementing AI includes data management, data pipelines for ML; and building, training, testing, and deploying ML.

Security Data Preparation

Standardize and optimize data formats and implement modern governance and data pipeline management so data such as security telemetry data can be used in AI tools. One CISO contends that data preparation is even more important than training. Be aware, CISOs who have begun data preparation at scale say the task is a bigger impediment than they had expected.

AI Assessment

Develop a robust assessment methodology to evaluate AI tools that you build or consume. Ensure that AI tools are trustworthy so you can apply AI to your most sensitive security processes and decision support systems.

Read more from the RSAC ESAF community of Fortune 1000 CISOs in the ***CISO Perspectives series***.