

Gartner Research

# Predicts 2024: AI & Cybersecurity — Turning Disruption Into an Opportunity

Jeremy D’Hoinne, Avivah Litan, Nader Henein, Mark Horvath, Akif Khan, Robertson Pimentel, Bart Willemssen, Dennis Xu, William Dupre

4 December 2023

**Gartner**<sup>®</sup>

## Predicts 2024: AI & Cybersecurity — Turning Disruption Into an Opportunity

Published 4 December 2023 - ID G00800663 - 27 min read

By Analyst(s): Jeremy D'Hoinne, Avivah Litan, Nader Henein, Mark Horvath, Akif Khan, Robertson Pimentel, Bart Willemsen, Dennis Xu, William Dupre

Initiatives: Cyber Risk; Meet Daily Cybersecurity Needs

Gartner predicts that AI will durably disrupt cybersecurity in positive ways, but also create many short-term disillusionments. Security and risk management leaders need to accept that 2023 was only the starter for generative AI, and prepare for its evolutions.

### Overview

#### Key Findings

- Generative AI (GenAI) is the latest technology in a long line of proclaimed disruptive technologies promising to fulfill the ongoing desire for organizations to drastically increase productivity metrics for all teams via automation of tasks.
- Today, most GenAI functions built into security products are focused on adding natural language interfaces to existing products to improve efficiency and usability, but promises of full automation start to appear. Past attempts to fully automate complex security activities, including using machine learning techniques, have rarely been entirely successful and can be a wasteful distraction today, and with short-term disillusionments.
- GenAI is at peak hype, driving very aggressive predictions based on the state of the technology today. This leads to unrealistic disruption claims, but also ignores next steps in GenAI evolution, such as multimodal models and composite AI.
- The initial forays by cybersecurity vendors into generative AI offer only a limited glimpse of the technology's promise and might not be the best indication of what the future could be.

## Recommendations

Security and risk management (SRM) leaders in charge of developing cybersecurity roadmap should:

- Construct a multiyear approach for progressively integrating GenAI features and products when they augment security workflows. Start with application security and security operations.
- Evaluate efficiency gains in tandem with the cost of GenAI implementations, and refine your detection and productivity metrics to account for new GenAI cybersecurity features.
- Prioritize investments in AI augmentation of the workforce, not just task automation. Prepare for short-term increased spend and long-term skill requirements changes due to GenAI. Monitor potential shift in attack success due to GenAI.
- Account for potential privacy challenges and balance expected benefits, with risks associated with cumulative cost in the valuation of large-scale GenAI adoption in security.

## Strategic Planning Assumptions

By 2028, multiagent AI in threat detection and incident response will rise from 5% to 70% of AI implementations to primarily augment, not replace staff.

Through 2025, generative AI will cause a spike of cybersecurity resources required to secure it, causing more than a 15% incremental spend on application and data security.

By 2026, 40% of development organizations will use the AI-based autoremediation of insecure code from AST vendors as a default, up from less than 5% in 2023.

By 2026, attacks using AI-generated deepfakes on face biometrics will mean that 30% of enterprises will no longer consider such identity verification and authentication solutions to be reliable in isolation.

By 2028, the adoption of generative augments will collapse the skills gap, removing the need for specialized education from 50% of entry-level cybersecurity positions.

## Analysis

### What You Need to Know

Predictions are statements of Gartner's positions and actionable advice about the future. This research highlights Gartner Predicts relevant for security and risk management leaders who have to navigate aggressive claims that GenAI is disrupting cybersecurity. Past experiences lead to skepticism given previous "AI washing," which caused expensive investments that didn't deliver expected results.

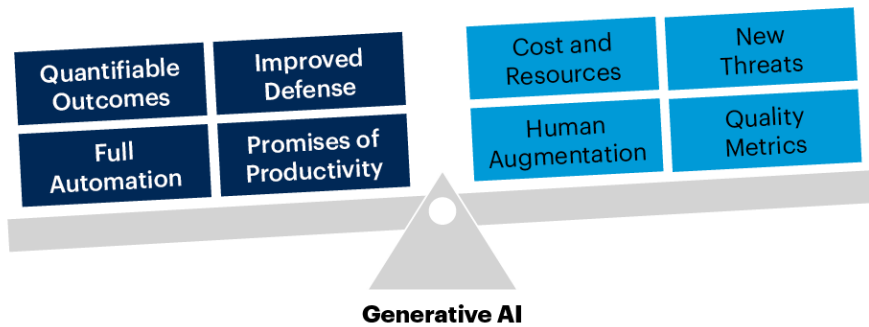
In *4 Ways Generative AI Will Impact CISOs and Their Teams*, Gartner gives recommendations on areas of immediate focus for security leaders:

- Manage the consumption of hosted and embedded GenAI applications.
- Update application security practices to AI applications, using AI trust, risk and security management (AI TRiSM) technologies.
- Assess the first wave of GenAI announcements from cybersecurity providers, and put a plan to integrate new features and products when they are more mature.
- Acknowledge that malicious actors will also use GenAI and be prepared for unpredictable changes in the threat landscape.

Excessive hype damages our perception of time and balance, but roadmap planning requires that cybersecurity leaders factor in all possibilities, without a strong fact base that balances cybersecurity realities with GenAI hopes or promises (see Figure 1).

Figure 1: Balancing Cybersecurity Reality with GenAI Hopes

## Balancing Cybersecurity Reality With GenAI Hopes



Source: Gartner  
800663\_C

Gartner

The cybersecurity industry has long been obsessed with fully automated solutions. The hype surrounding GenAI already led to unrealistic promises, potentially damaging the credibility of longer-term improvements coming from future features and products.

**2023 was the year of GenAI announcements, 2024 should be the year of minimum viable products; 2025 might be the first year of GenAI integration in security workflows delivering real value.**

As stated in the Hype Cycle for Generative AI, 2023, “Several innovations have a five- to 10-year period to mainstream adoption.” This is the case for “autonomous agents” and Gartner believes that cybersecurity leaders focusing on human augmentation will achieve better results than those jumping too quickly on solutions promising full automation.

In the shorter term, we’ll observe expansions of cybersecurity use cases from experiments of multimodal GenAI (i.e., learning from more than text content) and will improve our ability to measure productivity gains (see Innovation Insight: Multimodal AI Explained).

## Strategic Planning Assumptions

**Strategic Planning Assumption:** By 2028, multiagent AI in threat detection and incident response will rise from 5% to 70% of AI implementations to primarily augment, not replace staff.

**Analysis by:** Jeremy D’Hoinne, Dennis Xu

### Key Findings:

- More than a third of the first wave of announcements on GenAI in cybersecurity relate to security operation activities. Touted capabilities range from basic interactive help prompts to new dedicated product announcements aimed at becoming the primary interface for incident response and posture assessments.
- Full automation of threat detection, alert triage and incident responses are the “reach the moon” objectives of many threat detection, investigation and response (TDIR) initiatives.
- History often repeats and GenAI sparks the same overly-optimistic hopes for security operations, similar to what unsupervised machine learning did for threat detection more than five years ago.
- Conversely, teams with a higher maturity might imprudently dismiss generative cybersecurity AI, based on the early and immature implementations of large language models (LLMs) in the form of “SOC assistants” prompts.

### Near-Term Flag:

Through 2024, less than a third of generative cybersecurity AI implementation will lead to security operation productivity improvements for enterprises, generating more spend.

By 2026, the emergence of new approaches, such as “action transformers,” combined with more mature GenAI techniques will drive semiautonomous platforms that will significantly augment tasks executed by cybersecurity teams.

### Market Implications:

Building strong security operations is difficult, even for larger and well-funded organizations. Picking the right mix of tools, services and internal staff will suffer if cybersecurity teams invest time on tools that don’t deliver to their promise of automation.

We've observed this previously when implementations of unsupervised machine learning for threat detection promised to wipe out false positives and enable automated response. It took years for the tools to mature, and for security operation teams to tune them and narrow down automated blocking to the few use cases where it worked. With LLMs today – and autonomous agents, multimodal and foundation models in the future – organizations face a similar challenge. Early claims of GenAI awesomeness divert expectations from incremental improvements and team augmentation to less likely big shifts in automation, skill requirements and staff versus tool balance.

**Gartner anticipates short-term GenAI disillusion, especially in 2024, where external pressure to increase security operation productivity will collide with low maturity features and fragmented workflows.**

Symptoms of ill-prepared GenAI integration will include:

- Absence of relevant metrics to measure GenAI benefits, combined with premium prices for GenAI add-ons.
- Difficulties to integrate AI assistants in existing collaboration workflow within the security operation teams, or when partnering with a third-party security operation provider.
- Quickly growing “prompt fatigue:” too many tools offering interactive interface to query about threats and incidents.

With time, new AI approaches – combined with other non-AI techniques where relevant – might bring security operations closer to autonomous decisions for identified use cases. Emerging AI techniques supporting this promise include:

- **Multiagent systems (MAS):** Type of AI systems composed of multiple, independent but interactive agents.
- **Action transformers:** Models that learn from human actions.
- **Autonomous agents:** Self-prompting agents that can take actions based on LLMs recipes.

Although the myth of fully automated response and self-healing organizations might never truly turn into reality, Gartner believes that the combination of other techniques with multiagent approaches will have a big impact on security operations and security in general. Deployments aimed at both augmenting human tasks and adding precision and speed to human investigations will be more effective than single-technique AI analytics driving fully autonomous responses, such as automated containment for the foreseeable future.

## Recommendations:

- Navigate the chaos of newly announced GenAI features in security products by introducing business value-driven AI evaluation frameworks, which measure impact on tangible metrics such as speed, accuracy and productivity.
- Run GenAI pilots primarily for incident response and exposure management use cases that are not real time in nature. Set realistic short-term objectives, such as false positive reduction or opportunities to extend staff recruitment to slightly less specialized profiles.
- Protect the security operation team as much as possible from mandates originating outside of the security team to fully automate response and vulnerability treatment process. This will help avoid resistance when you need to implement promising GenAI techniques later.
- Be lucid about security providers' strategy to use GenAI as a claimed differentiator to promote large platforms leading to vendor lock-in.
- Don't neglect provider evaluation requirements to address privacy, copyright, traceability and explainability challenges.

## Related Research:

[4 Ways Generative AI Will Impact CISOs and Their Teams](#)

[Hype Cycle for Artificial Intelligence, 2023](#)

[Hype Cycle for Generative AI, 2023](#)

[Busting 4 Myths to Unlock More Cybersecurity Value](#)



**Strategic Planning Assumption:** Through 2025, generative AI will cause a spike of cybersecurity resources required to secure it, causing more than a 15% incremental spend on application and data security.

**Analysis by:** Avivah Litan, Jeremy D’Hoinne

## Key Findings:

- Gartner research shows that most enterprises have not yet formalized acceptable use policies for GenAI, so security and risk managers do not yet have a framework for instituting technical controls. <sup>1</sup>
- Integrating large language models (LLMs) and other types of models, such as foundation models in enterprise applications, bring new risks in three categories: content anomalies, data protection and AI application security.
- Almost 90% of enterprises are still researching or piloting GenAI, and most of those have yet to put AI TRiSM (trust risk and security management) technical controls or policies in place.
- Vendors hosting GenAI models do not always provide a complete set of controls that mitigate these risks. Instead, users need to acquire solutions that augment hosting vendors’ limited controls.
- IT leaders must rely on hosting LLM vendors with protection of their data, without the ability to verify their security and privacy controls.

## Market Implications:

The use of third-party hosted LLM and GenAI models unlocks many benefits, but users also must contend with new unique risks, requiring new security practices in three primary categories:

- Content anomaly detection
  - Unacceptable or malicious use
  - Unmanaged enterprise content transmitted through prompts or other methods, resulting in compromise of confidential data inputs
  - Hallucinations or inaccurate, illegal, copyright-infringing and otherwise unwanted or unintended outputs that compromise enterprise decision making or can lead to brand damage
  
- Data protection
  - Data leakage, integrity and confidentiality compromises of both content and user data in hosted vendor environment
  - Inability to govern privacy and data protection policies in externally hosted environments, or even contract service providers as data processors
  - Difficulty conducting privacy impact assessments and complying with various regional regulations, due to the black box nature of the third-party models and the mostly absent possibility to officially contract these model providers as data processors, following privacy legislative requirements
  
- AI application security
  - Adversarial prompting attacks, including business logic abuses and direct and indirect prompt injections
  - Vector database attacks
  - Hacker access to model states and parameters

Our recent survey of over 700 webinar attendees on what GenAI risks they are most concerned about validated these risk categories – and highlighted that privacy and data loss are the top risks from IT leaders. <sup>1</sup>

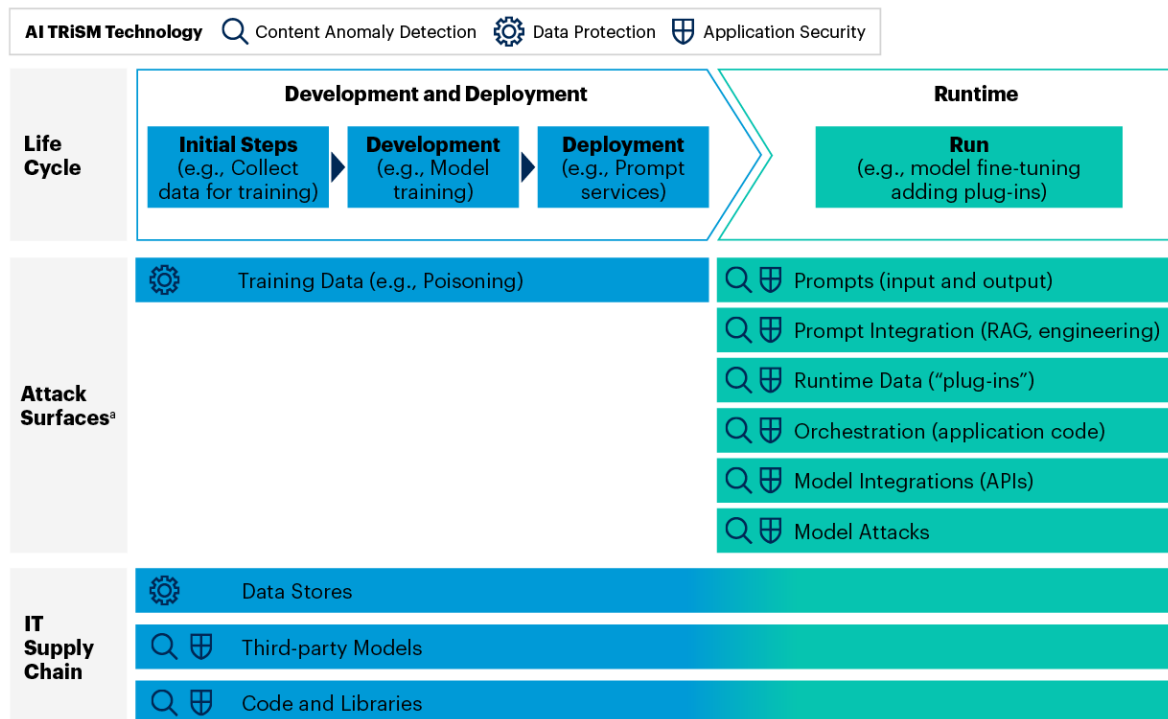
These risks are exacerbated when using externally hosted LLM and other GenAI models, as enterprises lack capabilities to directly control their application processes and data handling and storage. However, the risks still exist in on-premises models hosted and directly controlled by the enterprise – especially when security and risk controls are lacking.

These three categories of risks confront users during runtime of AI applications and models. Figure 2 shows how these three risks affect AI model development and deployment, the AI model at runtime, plus the effect from AI risks in the IT supply chain. This includes training data, third-party models, code and libraries, and prompt and model integrations.

These new attack surfaces will drive enterprise security departments to spend time and money implementing GenAI security and risk management controls, such that application and data security spending will increase at least 15% through 2025.

**Figure 2: Generative AI Attack Surfaces Across the AI Life Cycle**

## Generative AI Attack Surfaces Across the AI Life Cycle



Source: Gartner

<sup>a</sup>Main sample attack surfaces only; others not shown

796422\_C

Gartner expects that many enterprises will initially acquire solutions that mitigate input/output risks through anomaly detection or secure AI applications to gain visibility into enterprise use of GenAI applications and models. This includes use of off-the-shelf applications, such as ChatGPT or interactions through other integration points like plug-ins, prompts or APIs. Getting their arms around enterprise interactions with GenAI is the first priority for organizations, and these products can provide a good map of those interactions. Once the map is established, core functions of mitigating risks and security threats can be gradually deployed. This all has major implications on security staffing and budgets; hence our prediction that security budgets will increase.

## Recommendations:

- Organize within and across your enterprise to manage new GenAI risks and security threats. Once organized, establish acceptable GenAI use policies for your enterprise, and enforce them on a continual basis in part using AI TRISM technology.
- Set up proofs of concept to test emerging AI TRISM products, specialized in GenAI in the three new risk and security categories to augment your security controls, and apply them to production applications once they perform as required.
- Use content anomaly detection products that mitigate input and output risks to enforce acceptable use policy, and prevent unwanted or otherwise illegitimate model completions and responses from compromising your organization's decision making, safety and security.
- Perform user awareness training to remind users to always validate the output of GenAI products for accuracy before incorporating them into business workflow.
- Evaluate the use of AI application security products to protect your organization from hackers who exploit new GenAI threat vectors to damage your organization and its assets.
- Continue to use known security controls to protect sensitive information, application stacks and assets, but recognize they don't mitigate risks unique to LLMs, such as inaccurate, inflammatory or copyrighted outputs in responses.

## Related Research:

[4 Ways Generative AI Will Impact CISOs and Their Teams](#)

[Innovation Guide for Generative AI in Trust, Risk and Security Management](#)

## Generative AI Policy Template

### Microsoft Azure OpenAI vs. OpenAI: Comparing GenAI Trust, Risk and Security

### Quick Answer: How to Make Microsoft 365 Copilot Enterprise-Ready From a Security and Risk Perspective

**Strategic Planning Assumption:** By 2026, 40% of development organizations will use the AI-based auto-remediation of insecure code from AST vendors as a default, up from less than 5% in 2023.

**Analysis by:** Mark Horvath

#### Key Findings:

- Although 80% of vendors offering Application Security Testing (AST) have some form of suggesting fixes to code based on security problems, (autoremediation), less than 5% of development organizations use it – in part because the solutions it offers are generally examples, rather than actual code fixes.
- Developers complain that autoremediation suggested by code security tools (AST tools) often have adverse side effects on other aspects of their code, like performance and reliability. Because most developers have KPIs around these code aspects – and less stringent ones around security – they view these suggestions negatively.
- Developers can feel overloaded by the number of plug-ins to their developer environment – each offering advice on a specific parameter (e.g., code quality assessments, performance and optimization suggestions, etc.). Any new additions to the Integrated Development Environment (IDE) will need to synthesize suggestions based on the input of more than one autocorrection tool.

**Near-Term Flag:** While many AI-based secure code assistants are planned or are in development, their adoption by real-world production teams in 2024, as opposed to pilots or proofs of concept (POCs), will be a leading indicator that they offer an advantage over existing systems.

#### Market Implications:

Currently, the application security testing market is centered around a handful of core tools used for determining elements of code security risk (e.g., SAST, DAST, IAST, SCA, IaC, etc.). Although they interface with developers on a daily basis, they are primarily security tools and were designed to be used by, and for, security professionals working with developers. They are often heavy in terms of technical security jargon and assume that developers have an understanding of the data, and are able to action it to reduce security risk. However, there is often a considerable gap between the security training that developers receive, and real-world code security issues that often don't look like the examples they are taught.

Remediation guidance from standard AST tools is usually in the form of autocorrection, which works in ways similar to a spell checker (e.g., is this line formatted correctly)? Guidance to the developers is usually specific only to security, and only to the line or lines in question. It fails to provide a more comprehensive analysis of different aspects of the code in a larger context. This results in fairly generic advice, usually reflecting the OWASP top 10 as the basis of repair.

Large language models (LLMs) have the advantage that they are not only able to more easily deal with multiple code metrics like security, quality and reliability, they are very flexible in the way they can present the data and suggestions to developers. LLMs have the promise of being able to convert security jargon into an easier to understand format, leading to a better understanding of the issue and a more effective fix. The current generation of code security AIs offer a developer a choice of several different suggestions for addressing vulnerabilities, putting the developer in charge of picking the type of remediation that best fits into the application, thus preserving the “own your code” philosophy. This has several advantages:

- AIs and people often work better together than either one alone. The AI assistant offers a broader (and potentially deeper) view of a vulnerabilities' security posture, while the human understands the application's context, goals and workflows. The AI assistant allows a better selection of possible remediations, while keeping the application's function in mind.
- By presenting multiple options to the developers, they can more easily recognize and filter out misidentifications/hallucinations from the AI assistant.

- None of the auto remediation options available from AST tools effectively include parameters like performance, code quality, reliability etc., which are both important to development teams and well-correlated with security findings. New AI-based code assistants can optimize several variables beyond just security to give developers more motivation in line with their development KPIs.

## Recommendations:

- Most enterprises should not use generic LLMs like ChatGPT for code generation, code security scanning or secure code review, due to the higher error rates of tools not specific to security. Instead, rely on tools that offer enterprise grade security and governance controls for assisting developers with technical tasks like security.
- Pilot no more than two or three different AI security code assistants to compare and contrast their capabilities. Though products are recently becoming commercially available, the market still has a long way to go before these are common tools. The current generation has strengths and weaknesses in different areas, so have development teams test them out to determine the most effective ones for your organization.
- Maintaining the existing developer experience is critical to the successful adoption of any developer focused tools. Changes in workflow, experience or testing works against the “muscle memory” of developers and generates friction, which will frustrate developers who will then avoid using the tools.
- Remember that these tools use an LLM, which will need periodic retraining. When choosing a vendor, ask specifically about privacy, data retention and retraining details to protect your IP. Ask about indemnification around IP loss, licensing issues with some code or accidentally re-using another company’s IP.
- AI Coding Assistants are rapidly becoming a popular way for developers to write better code at a faster rate. Be sure to run Static Analysis (SAST) and Software Composition Analysis (SCA) on code that has been generated by AI. This will help ensure code quality, protect IP rights and cut down on AI mistakes and misrepresentations.

## Related Research:

Quick Answer: Mitigating the Top Five Security Risks of AI Coding

Emerging Tech: Generative AI Code Assistants Are Becoming Essential to Developer Experience

Magic Quadrant for Application Security Testing

Hype Cycle for Application Security, 2023

Innovation Guide for AI Coding Assistants

**Strategic Planning Assumption:** By 2026, attacks using AI-generated deepfakes on face biometrics will mean that 30% of enterprises will no longer consider such identity verification and authentication solutions to be reliable in isolation.

**Analysis by:** Akif Khan, Robertson Pimentel

### Key Findings:

- In a little over a decade, there have been several inflection points in fields of AI that can be used to generate synthetic images. <sup>2,3,4,5</sup> These capabilities have been made available via accessible tools that enable relatively unskilled users to create synthetic images of real people's faces ("deepfakes"), with obvious consequences for abuse by threat actors to subvert biometric processes.
- Identity verification is commonly carried out by users on their mobile device, involving taking a picture of their photo identity document and a selfie of themselves, which are then biometrically compared. Mobile application development has been democratized, with advanced developer tools easily available to all. This has enabled bad actors to find new ways of attacking these mobile-based identity verification processes.
- Identity verification and authentication processes using face biometrics today rely on presentation attack detection (PAD) to assess the user's liveness. Current standards and testing processes to define and assess PAD mechanisms do not cover digital injection attacks using the sophisticated deepfake images of faces that can be created today. <sup>6,7,8</sup>



- Informal discussions with identification verification vendors indicate that deepfakes are now involved in around 15% of detected fraudulent identity presentations. Of course, the number of undetected attacks is an unknown factor. The majority of vendors cite presentation attacks as still being the most common attack vector, but that injection attacks have increased over 200% during the first nine months of 2023.

## Market Implications:

Gartner currently tracks over 70 identity verification vendors in the market. These vendors verify identity via a process in which a user takes a picture in real-time of their government-issued photo identity document. This is then assessed for authenticity using computer vision techniques, sometimes augmented by human analysts. The user is then prompted to take a selfie of their face. This is first assessed for liveness by means of PAD, and then biometrically compared to the picture in the identity document. Identity verification is now a business critical security control for a range of organizations (see Market Guide for Identity Verification), and should not be confused with use of on-device biometrics, which are not bound to an identity.

Attempts to subvert the face biometrics have typically focused on presentation attacks, in which attackers have worn latex masks, or pointed device cameras at pictures or videos (either original content or synthetic deepfake) of the person being impersonated. Most vendors are focused on such presentation attack methods, given that is the focus of current standards and testing for liveness detection.

Attackers, however, are looking beyond presentation attacks and have leveraged the diversity and accessibility of emulation and application development tools available to make more sophisticated digital injection attacks. These consist of bypassing the camera on the device and injecting the deepfake image directly into the path between the sensor and the verifier. Toolkits specifically for injecting deepfakes are freely available to all. <sup>9</sup>

Preventing such attacks will require a combination of PAD and also injection attack detection (IAD) and image inspection. Active PAD (“liveness detection”) techniques, such as asking a user to blink, turn their head or repeat a random phrase, are becoming of diminishing use given tools that create synthetic imagery in real time that can mimic an attacker’s head movements. <sup>10</sup> Passive liveness detection techniques offer more resilience, given that an attacker may not know which features of the image are being used to assess liveness. However, with the rapid advances taking place in the field of deepfakes, the longevity of passive detection techniques such as analysis of perspective, reflection of illumination or micro-inspection of blood flow should not be assumed. Some vendors are also using GenAI to create synthetic data upon which to train their systems to better detect deepfakes.

In a crowded market that includes many smaller regional vendors and niche players, not all will be able to remain at parity with the attackers in this rapidly evolving arms race. There will be client attrition and those vendors who can keep ahead of the curve and demonstrate efficacy will reap the benefits. Similar challenges are present in other biometric fields – for example, the creation of synthetic voices already poses problems to the integrity of voice biometric processes today.

#### Recommendations:

- Favor vendors who can demonstrate that they have features and a strategy that looks beyond current standards and are tracking, categorizing and quantifying the novel attacks they experience today.
- Establish a minimum baseline of controls by working with vendors that not only focus on PAD (with passive liveness detection, rather than active), but have specifically invested in mitigating the latest deepfake-based threats using IAD coupled with image inspection.
- Further increase the chances of detecting attacks on your identity verification processes by adding additional risk and recognition signals, such as device identification and behavioral analytics. For authentication use cases, always use additional signals or credentials beyond just face biometrics.
- Improve resilience by accepting that some attacks will likely succeed and invest in postverification monitoring of account activity to help detect bad actors who may have defeated the identity verification or authentication process, but may reveal themselves through subsequent anomalous actions.

## Related Research:

- Market Guide for Identity Verification
- Buyer's Guide for Identity Proofing
- How to Mitigate Account Takeover Risks

**Strategic Planning Assumption:** By 2028, the adoption of generative augments will collapse the skills gap, removing the need for specialized education from 50% of entry-level cybersecurity positions.

**Analysis by:** Nader Henein, William Dupre

## Key Findings:

- The current technical skills gap for cybersecurity workers has grown to a chasm. Today, the deficit stands at 3.4 million open positions, <sup>11</sup> up 26% year over year with no indication that this trend will slow down.
- Even though a university degree may not be required for many entry-level cybersecurity workers, this is often replaced by alternative credentials, such as industry certifications or bootcamps.
- When AI capabilities are deployed in-line with user tasks to boost worker capabilities beyond what the average human could achieve, this is referred to as an “augment.” The developments in GenAI have given birth to generative augments that can be built to support specific tasks and roles.

## Near-Term Flag:

- General-purpose augments, such as Microsoft's Security Copilot or Google's Duet AI integrated into Mandiant Threat Intelligence, will make limited inroads to complement knowledge workers' workflows – but adoption and usage will be limited by their conversational/chat interface.
- A new industry of development houses will emerge, focused on building purpose-built augments for organizations investing in products such as Azure OpenAI, AWS Bedrock, Google's Vertex AI or NVIDIA NeMo.

## Market Implications:

When there is a fundamental change in the tooling available to an industry, this is closely followed by an equally fundamental shift in how that industry resources and operates. The advent of GenAI augments will change how organizations hire cybersecurity workers looking for the right aptitude, just as much as the right education.

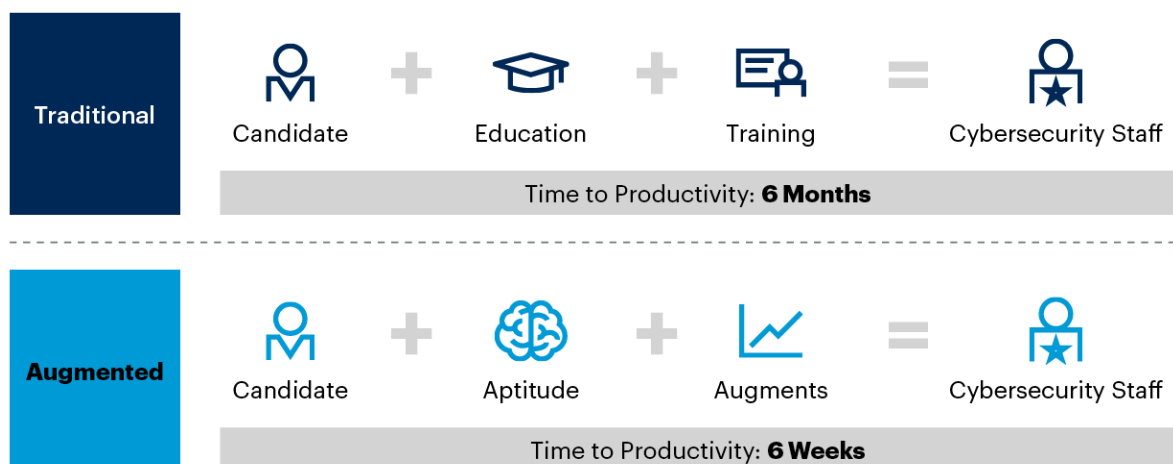
Augments are AI-based agents that are deployed in-line with user tasks to boost worker capabilities beyond what the average person could achieve. Generative Augments, as the name suggests focuses on using GenAI capabilities within the augment acting as an agent on behalf of the user (for a detailed illustration see How to Deploy Generative AI Capabilities Behind the Firewall to Augment Your Workforce).

Today, Gartner sees many general-purpose augments appearing in office productivity tools that are used on a day-to-day basis. For security practitioners this pattern repeats, mainstream platforms such as Microsoft and Google all offer augmentation. Though most of these are still conversational, their next iteration will be “baked into” the practitioners’ workflow. In parallel, we see the rise in purpose-built augments, developed within organizations, first focused on specific tasks. Later on, those task-specific augments will be combined to support roles and facilitate recruitment for those roles (see Figure 3).

**Figure 3: Illustration of a Generative Augment Workflow vs. Traditional**

### Illustration of a Generative Augment Workflow vs. Traditional

Workflows for Knowledge Worker Onboarding



Source: Gartner  
800663\_C

In the not so distant future, these augments will start learning by observing their user becoming not only task or role specific, but trained to be user specific.

This will change how we teach, how we hire and how we specialize across cybersecurity as the industry receives a much needed boost through generative augments.

## Recommendations:

- For internal enterprise use cases, invest in generative augments that support the users as they work, rather than conversational bots that require the user to stop and chat. This gives the organization more fine-tuned control over how generative capabilities are used and a better return on investment. For example, an augmented security operation tool could improve incident response by automatically generating focused and contextual playbooks and threat models based upon various telemetry sources and a properly grounded LLM.
- Track homegrown and vendor-provided augments at the task and role level to better gauge the evolution in the skillsets you require in entry-level cybersecurity workers, and coordinate accordingly with your HR partners.
- Utilize this toolkit from Gartner (see [Identifying Adjacent Talent for Key Cybersecurity Roles](#)) to identify adjacent talent for some of the more critical cybersecurity roles. These prospects would make ideal candidates for augmentation.

## Related Research:

[How to Deploy Generative AI Capabilities Behind the Firewall to Augment Your Workforce](#)

[Plan for Generative AI's Impact on Jobs](#)

[Podcast: Planning for Generative AI's Impact on Jobs](#)

## A Look Back

*In response to your requests, we are taking a look back at some key predictions from previous years. We have intentionally selected predictions from opposite ends of the scale – one where we were wholly or largely on target, as well as one we missed.*

*This report is too new to have on-target or missed predictions.*

## Evidence

<sup>1</sup> Generative AI: A Look at Emerging Governance Practices

<sup>2</sup> Generative Adversarial Nets, Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, Yoshua Bengio.

<sup>3</sup> Deep Unsupervised Learning Using Nonequilibrium Thermodynamics, Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, Surya Ganguli.

<sup>4</sup> Improving Language Understanding by Generative Pre-Training, Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever.

<sup>5</sup> NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis, Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng.

<sup>6</sup> Information Technology Biometric Presentation Attack Detection Part 3: Testing and Reporting, ISO.

<sup>7</sup> ISO 30107-3 Presentation Attack Detection Test Methodology And Confirmation Letters, iBeta.

<sup>8</sup> Face Analysis Technology Evaluation (FATE) PAD, NIST.

<sup>9</sup> The Deepfake Offensive Toolkit, GitHub.

<sup>10</sup> Real-Time DeepFake Streaming With DeepFaceLive, UniteAI.

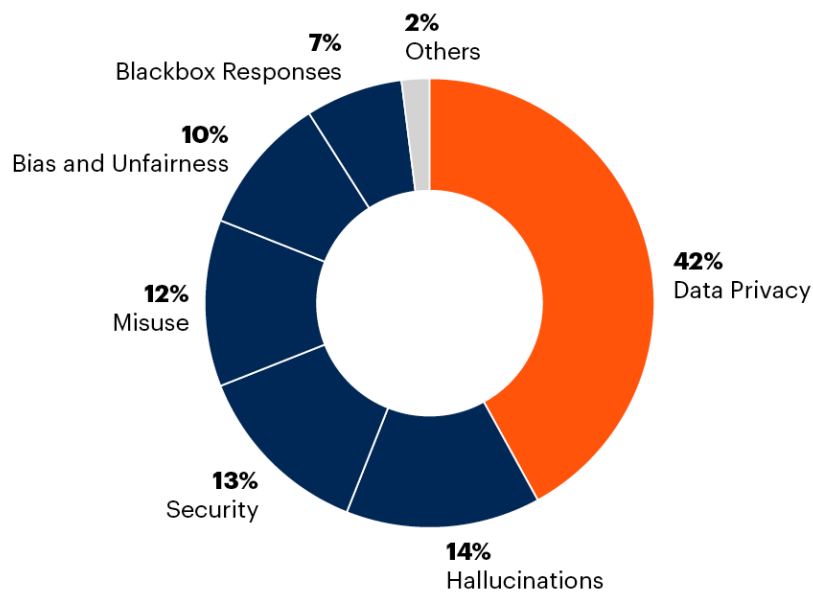
<sup>11</sup> ISC2 Cybersecurity Workforce Study, ISC <sup>2</sup>.

## Note 1: IT Executive Poll

Gartner conducted many surveys with IT and security executives on GenAI. The graphic below comes from a poll, done during a webinar with more than 700 IT executives. It confirms the top concerns for these IT executives and aligns with what Gartner analysts hear during client inquiries (see Figure 4).

Figure 4: IT Executive Poll on GenAI Concerns

### IT Executive Poll on GenAI Concerns



Source: Gartner  
800663\_C

Gartner

---

## Recommended by the Authors

Some documents may not be available as part of your current Gartner subscription.

4 Ways Generative AI Will Impact CISOs and Their Teams

Innovation Guide for Generative AI in Trust, Risk and Security Management

---

© 2024 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner research may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.



# Actionable, objective insight

Position your organization for success. Explore these additional complimentary resources and tools for cybersecurity leaders:



## Report

### Cybersecurity Trends: Optimize for Resilience and Performance

Use this report to equip your cybersecurity function for greater resilience.

[Download Now](#)



## Roadmap

### IT Roadmap for Cybersecurity

Create a resilient, scalable and agile cybersecurity strategy.

[Download Now](#)



## eBook

### Leadership Vision for Security and Risk Management Leaders

Explore the top 3 strategic priorities for security and risk management leaders.

[Download Now](#)



## Webinar

### Strengthen Your Cybersecurity Leadership to Navigate Evolving Security Landscape

Explore this five-part series for insights into the evolving landscape.

[Watch Now](#)

Already a client?

Get access to even more resources in your client portal. [Log In](#)

# Connect With Us

Get actionable, objective insight that drives smarter decisions and stronger performance on your mission-critical priorities. Contact us to become a client:

**U.S.:** 1 866 263 8917

**International:** +44 (0) 03301 628 476

[Become a Client](#)

**Learn more about Gartner for Cybersecurity Leaders**

[gartner.com/en/cybersecurity](https://gartner.com/en/cybersecurity)

**Stay connected to the latest insight**

