Deloitte.

in association with Snyk

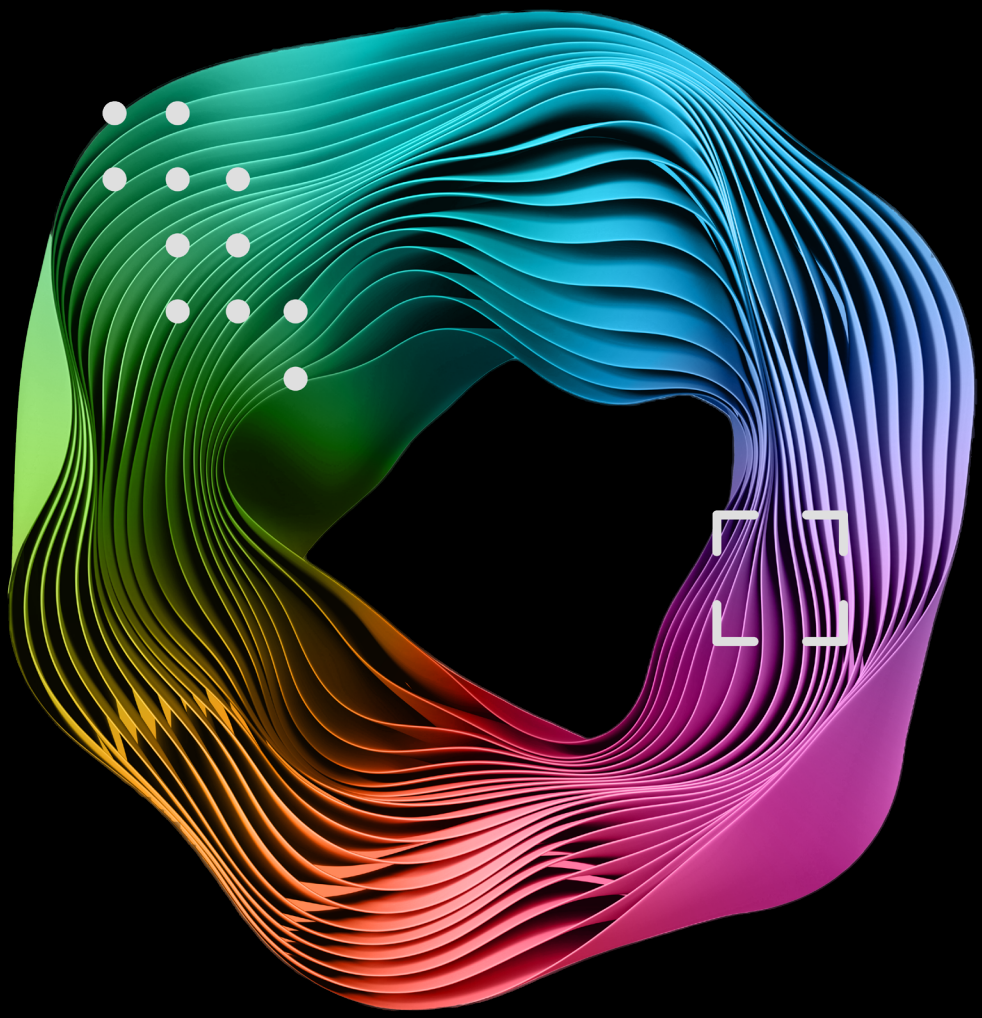

Scaling application security for GenAI with Deloitte and Snyk

Implement GenAI-assisted development tools safely by scaling your application security to meet demand

The introduction of artificial intelligence (AI) is driving a rapid evolution of the technology landscape, causing AI-driven development to outpace typical application security (AppSec) program capabilities.

Many businesses looking to enhance developer capabilities and streamline operations across multiple IT domains turn to Generative AI (GenAI) tools for their ease of use. These tools have become highly accessible to developers, with powerful pre-trained models and open-source versions available to the general public.

However, these new technologies can introduce low-quality or insecure code at a previously unseen pace of development that can widen the already considerable gap between software development and traditional AppSec program capabilities. This issue highlights the need for a multifaceted, modern AppSec or cybersecurity approach to address these rising challenges.

"Gartner® predicts that by 2025, GenAI will be a workforce partner for 90% of companies worldwide."[1] It's expected to cover tasks such as service desk functionality, application development processes, AppSec, and incident response—but like anything else, this technology has its pitfalls.

Cybersecurity leaders should start preparing now for this new world of unknowns, as it will be critical to understand the effects of increasing GenAI usage on AppSec programs and how it can be safely configured to match the scale of development enabled by GenAI tools.

1: Gartner Press Release, Gartner Says AI Ambition and AI-Ready Scenarios Must Be a Top Priority for CIOs for Next 12-24 Months, November 6, 2023 GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

Dimensions of cybersecurity within GenAI

Deloitte's AI use vision and strategy framework summarizes a multifaceted approach to securing AI across multiple domains, some of which overlap with AppSec strategy domains. In this paper, we will summarize how an organization can effectively transform existing AppSec strategies to accommodate AI.

Defining an AI use strategy includes determining the integrity of AI systems against malicious actors and the system's capability to defend from cyberthreats. The AI strategy should promote the safe and secure launch of AI and GenAI-enabled products, including development tools and development team enablers.

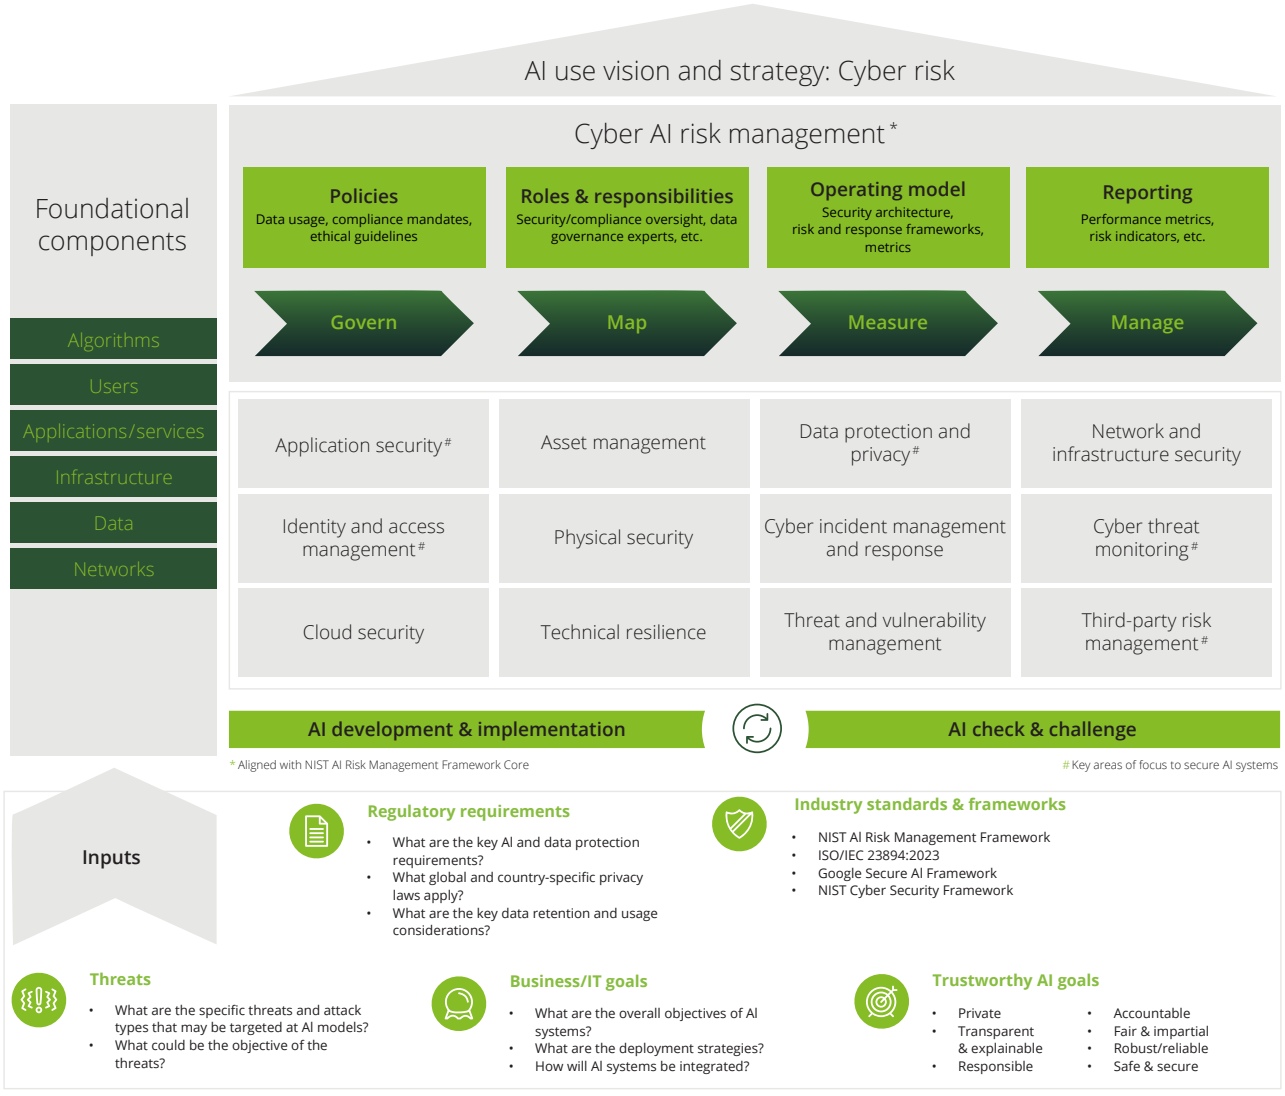

Deloitte's framework for AI use includes inputs from multiple industry-accepted sources, and considerations unique to each organization. Across AI-specific foundational components and general security domains, this exemplifies a thorough approach to identifying and implementing an AI risk strategy.

"With the increasing use of GenAI applications by development teams, organizations should adopt a broad security strategy. By providing the right tools, security leadership can enable developers to leverage automation and capabilities that give feedback on weaknesses in their AI-generated code quickly and set up a scalable AppSec program that integrates security compliance into their workflow without slowing down development operations."

Faris Naffaa
Deloitte Advisory Senior Manager

Cybersecurity for GenAI-enabled development

While GenAI has not yet reached the peak of its potential, many businesses are already adopting it. As noted in "CISO's Guide: Using AI for Cyber Defense",[2] AI is rapidly evolving due to its "capacity to automate labor-intensive, standardized, or error-prone activities to reduce manual work performed by cyber teams." These helpful abilities of GenAI tools are changing how modern cybersecurity teams operate. To adapt to this change, many tech teams are leveraging GenAI tools to expand application development. However, these rapid changes present new challenges to AppSec and DevSecOps teams. "Gartner® predicts that by 2025, GenAI will cause a spike in the cybersecurity resources required to secure it, causing more than a 15% incremental spend on application and data security."[3]

AI-generated code is inherently insecure, and the quality and security of GenAI code hinges primarily on the security of the large language model (LLM) behind the AI coding assistant, and the data used to train it. The usage of AI coding tools can also carry risks in other ways, as recent incidents reveal that popular GenAI tools could leak secrets based on the provided prompts and the response data used.

Another challenge is the lack of organizational context in open-source GenAI tools. LLMs may lack the contextual knowledge of an organization's policies and leading practices, and they may not account for regulatory and compliance regulations such as payment card information, personally identifiable information, Health Insurance Portability and Accountability Act, and Sarbanes-Oxley based on the data that the organization handles.

There is also the risk of simply growing too fast. The integration of GenAI tools can massively accelerate the pace of code development and exacerbate existing code security issues. Without security scanning automation and appropriate remediation guidance, a surge in security risks could easily overburden an unprepared organization's AppSec team and their vulnerability management process.

Challenges in the marketplace with the rise of GenAI-assisted development highlight the need for a broad, multilayered approach encompassing people, processes, technology, and governance. This methodology involves having multiple safeguards in place to help prevent critical failures and provides a multifaceted security strategy focused on AI-enhanced processes. This approach is crucial, especially to garner stakeholder understanding and leadership buy-in.

2: CISO's Guide: Using AI for Cyber Defense, The Wall Street Journal, 2024
3: Gartner Press Release, Gartner Forecasts Security and Risk Management Spending in India to Grow 12% in 2024

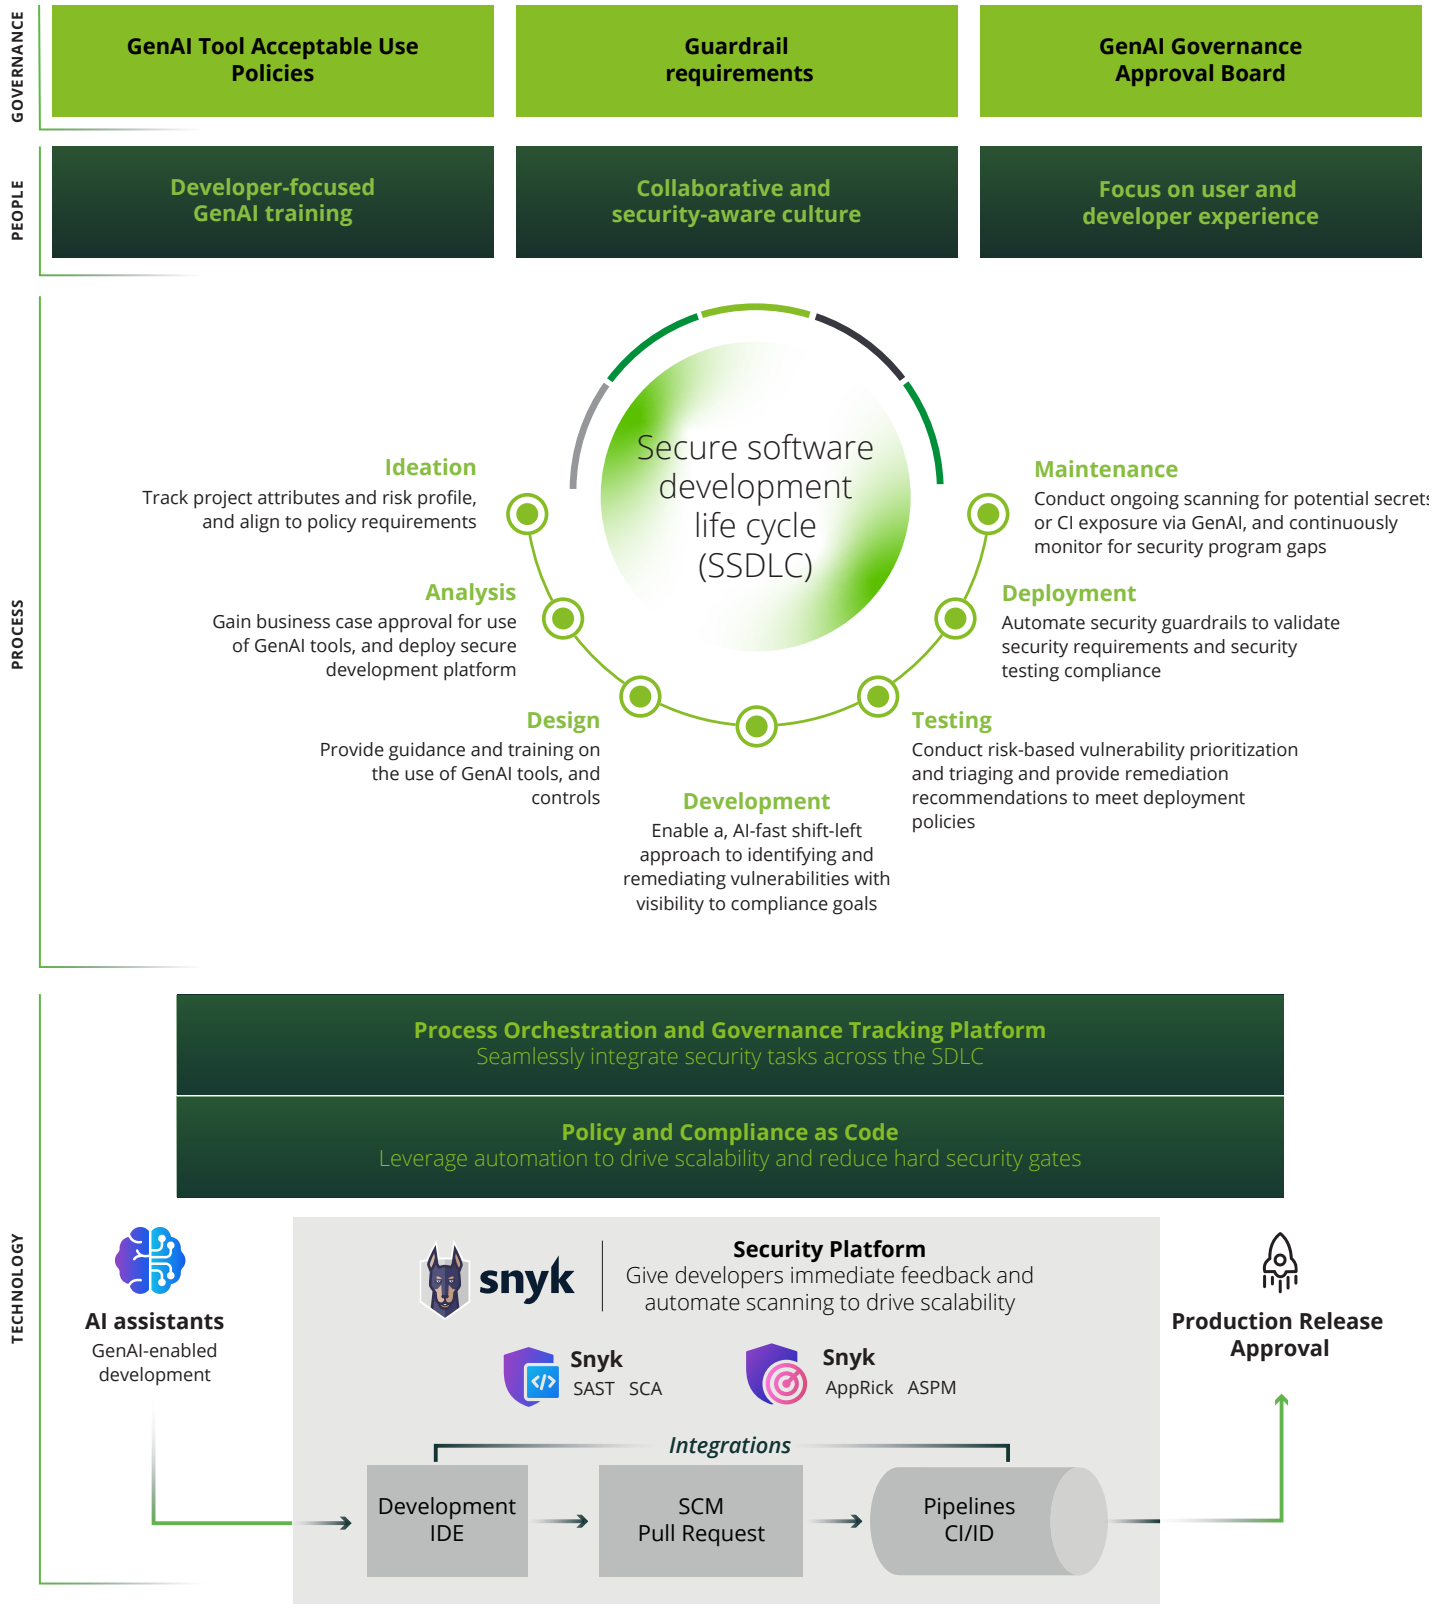

Many organizations are familiar with SSDLC processes and how they may integrate with security capabilities. With the introduction of GenAI tools, it is critical to consider impacts across People, Process, Technology, and Governance to modernize SSDLC programs to accommodate the emergence of new technologies.

Core principles: People and process

People

The "human factor" plays a crucial role in securing AI systems. Training is vital to help team members understand risks, leading practices, and the rationale behind safeguards.

It's critical that organizations retain skilled professionals who can leverage AI but also understand the importance of not compromising governance and review guardrails. Training should also cover guidance on using native integrated development environment (IDE) security tools that provide real-time, proactive safety protocols for AI-generated code. This focus should include guidelines on what data can (or cannot) be entered into GenAI tools to protect intellectual property and remain compliant with data regulations such as the General Data Protection Regulation.

Clear guidance for developers is critical. A report by Snyk, a leader in developer security, found that more than half of AI users commonly encounter vulnerabilities in AI-generated code, with some of these users having concerns about the security implications of using AI code completion tools. Despite this, 80% of developers were reported to bypass AI code security policies. "In our experience, there is a noticeable trend in the industry where developers who use AI coding assistants tend to trust GenAI more than they trust themselves to produce secure code" says Faris Naffaa. Considerations in this domain go hand in hand with critical technology and governance guardrails.

Process

Organizations should consider how process and policy changes should be updated to accommodate the increased use of GenAI tools for development without sacrificing security. GenAI tools allow for a rapid pace of development, which may significantly affect the efficacy of existing AppSec processes. In addition to processes and policies governing the use of GenAI, consider the following recommendations to scale key components of your AppSec program to meet increased demand introduced by GenAI-enabled development:

Communication: Establish clear guidance for developers communicating with AppSec teams on the use of and outputs from GenAI tools, as well as acceptable use cases and appropriate inputs to the tools. Development teams should know how and when to engage the AppSec team to review outputs from scanning tools, and the communication cadence between teams should be early and often.

False-positive guidance: As developers adopt GenAI tools to support development needs, they will likely see an increase in results from their security scanning. Developers should be equipped with proper training and guidance to review scan results, identify potential false positives, and engage the AppSec team for remediation or acceptance guidance.

Vulnerability prioritization: With the increased scale of development and scanning enabled by GenAI tools, organizations will need to implement prioritization processes to align remediation efforts with business needs. Prioritization should be conducted based on the vulnerability criticality, the application's risk classification, and the business criticality of the application or service.

Securing GenAI with Snyk: Technology

Gartner® predicts that, "By 2026, 40% of development organizations will use the AI-based auto remediation of insecure code from AST [application security testing] vendors as a default, up from less than 5% in 2023."[4]

Selecting cybersecurity tools that meet the velocity and scale of AI output is an essential decision with significant impact on an organization's security posture. Yet a 2023 report published by Snyk found that fewer than 10% of organizations automate a majority of their security scanning capabilities. When considering the added complexity of GenAI in development, it's critical to integrate AI-ready, scalable security scanning tools into the end-to-end, GenAI-enabled development life cycle.

Purpose-built security tooling

Industry experience indicates that developers have and will continue to adopt GenAI-assisted development tools. These tools are becoming highly accessible in the market. They're easy to use, with minimal blockers to leverage GenAI-written material in existing codebases.

However, there is an assumption that AI coding tools can inherently produce secure code. Cybersecurity leaders should consider integrating GenAI tools with security solutions that have established credentials in the cybersecurity industry. These solutions should be purpose-built to secure code with these considerations:

4: Gartner, Predicts 2024: AI & Cybersecurity — Turning Disruption Into an Opportunity, Jeremy D'Hoinne, Avivah Litan, et al., 4 December 2023

Shift-left: Security tools need to be able to catch vulnerabilities early in the development process, ideally while code is being developed. They should provide visibility into applications security posture, baked into the IDE itself. These tools should apply impartial security principles to maintain a high level of protection.

Speed: To facilitate developer adoption, security tools should not slow down or substantially change developer workflows. Scans should run as early and as efficiently as possible, ideally in a native IDE. Excessive scanning time results in frustration and the idea that security becomes a blocker, which can add up to hundreds of thousands of developer hours wasted each month.

Security expertise: Because GenAI lacks cognition, process design and curation by human subject-matter experts is vital for maintaining security standards. And though more general-purpose AI models excel at breadth of topics, they lack the depth and accuracy of more narrowly focused AI models that are trained on specific data and fine-tuned for just one purpose. This is why built-for-purpose security tools with narrowly focused AI models generally perform better at expert tasks. Hybrid, multi-model AI models also generally perform better. For example, when coupled with an LLM, a symbolic AI model's structured rules and knowledge base can help it to greatly reduce the likelihood of hallucinations and errors from the LLM.

Integration: Selecting scanning tools that integrate efficiently into existing workflows can reduce stakeholder resistance to additional security requirements. It can also provide an automated layer of security to GenAI-developed code. Consider implementing static application security testing and software composition analysis tools that integrate into existing processes and work environments to maximize developer adoption.

Insight into application risk: A sophisticated AppSec tool may enhance program capabilities by combining application context—including runtime intelligence, observability, and cloud security data sources—to better identify, prioritize, and remediate top risks to the business. When results from AppSec tests are processed against security intelligence from public sources, proprietary research, machine learning, and human-in-the-loop AI, organizations can gain valuable insight into how to prioritize business risks most effectively.

AppSec program considerations for securing GenAI: Governance

More than 50% of organizations surveyed by Snyk implement advanced risk tools, policies, and procedures but fail to ensure that their teams actually adhere to them. This ultimately results in AI-introduced risks. Trustworthy and reliable AI guardrails enable organizations to control risks while reaping the benefits of GenAI development tools.

Governance

Effective governance is critical in the rapidly changing realm of GenAI. Such governance, however, involves stringent security validation, implementing AI-ready security policies, and installing security tools before deploying GenAI for development. Nonetheless, Snyk's research shows a mere 24.8% of entities employ software composition analysis to verify the security of code suggestions provided by AI tools. Organizations should establish programmatic policies to prevent the hasty, uncritical acceptance of AI-generated code and to make sure that generated code undergoes review.

Internal policies concerning the use of GenAI tools for development should address the following:

Clear definition of acceptable use cases: Policies should clearly define which AI tools are allowed for developers to use. Be sure to define and enforce additional requirements that developers must follow to use GenAI tools for development. Effective application risk management enables security teams to implement and manage acceptable use policies across the organization.

Data selection and usage: Specify the type and quality of data that can be input to GenAI tools to generate a response. It's a leading practice to restrict developers from using inputs that contain secrets or an organization's confidential information.

Data privacy and security: Confirm that any data used in AI models is properly protected, and take appropriate measures to maintain data privacy and security.

Compliance requirements: Confirm that the use of AI tools complies with legal, regulatory, and ethical requirements.

Regular updates: Continuously update policies as new AI tools are developed and as business needs evolve.

4: Gartner, Predicts 2024: AI & Cybersecurity — Turning Disruption Into an Opportunity, Jeremy D'Hoinne, Avivah Litan, et al., 4 December 2023

Moving forward with GenAI

As the influence of GenAI expands, it is critical for organizations to adopt a multifaceted approach to cybersecurity. Increased use of GenAI tools without proper review of acceptable use cases or additional steps to secure the outputs may lead to large gaps in your security. A way to protect yourself is by providing developers with tools that offer quick feedback on vulnerabilities in their code and implementing a scalable AppSec program that enforces security compliance.

Deloitte's AppSec program modernization services and Trustworthy AI™ framework, combined with Snyk's AI-ready AppSec platform, can help organizations mature developer security initiatives and improve visibility into enterprise-wide application risks from code to cloud.

Learn about Secure by Design capabilities from Deloitte here.

Contact Snyk to learn more about securing AI-generated code.

Contact Deloitte to explore opportunities to enhance your AppSec program capabilities and harness the power of GenAI for your development teams.

Authors

Faris Naffaa
Senior Manager
Secure by Design Solution Leader
Deloitte & Touche LLP
fnaffaa@deloitte.com

Ben Desjardins
VP
Product Marketing
Snyk

Snyk Alliance contacts

Vikram Kunchala
Principal
US Cyber
Solutions and Platforms Leader
Deloitte & Touche LLP
vkunchala@deloitte.com

Kelli Wolfe
Snyk Alliance Manager
Deloitte & Touche LLP
kelwolfe@deloitte.com

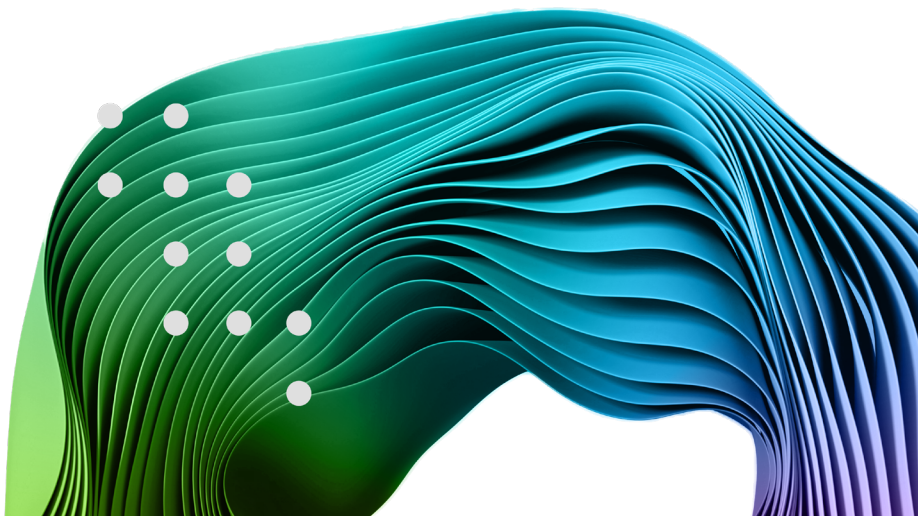

Deloitte.