


7 Steps to Safely Adopt GenAI in Application Security



A Case Study in Disruption

Everyone is Experimenting With AI

Everyone is using or experimenting with AI today. It is a real time case study in disruption. There are plenty of examples of how generative AI (GenAI) is helping people be more productive, from students to writers, marketers, and software development teams.

According to [Gartner](#), approximately 80% of enterprises will have used GenAI application programming interfaces (APIs) or models by 2026. As AI increases productivity for organizations, it is fueling further demand and adoption. As a result AI tools are being quickly adopted; however, in many cases it's unofficial or unsanctioned shadow IT.

GenAI tools are free, or relatively low cost, and AppSec and development teams are already seeing the productivity and efficiency increase. AI can also create or expose security vulnerabilities and become a powerful tool for malicious actors. Instead of slamming the brakes on innovation and restricting the use of GenAI in application development and AppSec, we must understand how to use it with minimal introduction of risk.

At Checkmarx, we speak to security professionals daily, and every time we ask them “what’s on your mind?” they invariably mention AI. We surveyed 900+ AppSec Managers and CISOs to learn more about the benefits and challenges around AI and understand how application security teams view the changes happening in development because of AI. In this report, we present the findings of our survey alongside a list of 7 key steps to start including generative AI in your application security practice.

In this report

- See the results of our research
- Understand how fellow AppSec professionals manage (or don't) GenAI
- Learn how to begin building guardrails around GenAI



Key Findings



26% have purchased GenAI tools or plan to purchase within the next two years.



70% say different groups within the organization are allowed to purchase different AI tools on a case-by-case basis.



Only **12%** say augmenting their software development workflow to leverage GenAI is a top priority.




60% are worried about GenAI attacks such as AI hallucinations.



80%

are worried about security threats stemming from developers using AI in code generation.



31%

say AI will help security teams and developers remediate faster.

AI in Software Development

Uncoordinated, Ungoverned, and Uncertain

There is a sense of excitement around the possibilities that AI creates for AppSec and development teams. But at the same time AI is also introducing a sense of uncertainty. Likewise, as AppSec managers work out how best to embrace AI, many also view it as 'yet another technology to govern'. Therefore, how can they use AI positively, without compromising on AppSec principles and governance?

According to a recent [Gartner Peers Insights Poll](#) focused on how software engineering teams are using GenAI, 38% of respondents cited issues creating governance policy and a similar percentage also cited cybersecurity problems. Forty one percent of respondents said the cost of AI tools was a challenge.

Our survey shows a similar picture. Seventy percent of respondents say they allow AI purchases and usage on a case-by-case basis. Only 14% of respondents said that they have a single, official tool that teams are authorized to use for code generation. Overall, there is minimal governance around the use of AI tools for code generation.

Only

14%

of survey respondents

have an official AI tool that teams are authorized to use for code generation.

AI

FIGURE 01

How far along is your organization in the adoption of Generative AI tools?



Despite all the hype around AI organizations report that they are not rushing to officially adopt solutions (Figure 01). 42% of respondents said they have no immediate plans to purchase GenAI and 29% said they were experimenting but had no plans to purchase. In fact, when asked about the level of focus in their organization around augmenting their software development workflow to leverage GenAI, only 12% of respondents said this was a high priority. This means developers are being left to decide when and how to use GenAI. Additionally, we named 11 GenAI tools in the survey, with GitHub Copilot being used by 42% of organizations – the most popular. However, every AI tool listed had between 20 to 40% of respondents saying they knew it was being used in their organization. This implies that there is considerable ad hoc use of a lot of different tools.

In terms of any guidance or governance, 15% of companies have completely banned AI for code generation, which means that 85% are using it in some way. Only one respondent said that 0% of their company's code was AI generated, and only three respondents said that their company is not using AI code generation tools. This indicates a high likelihood of shadow AI tool adoption.

Looking ahead, not a single AppSec professional said that absolutely none of their organization's code will be AI-generated in 2024.

This is understandable because AI is incredibly easy to adopt. It is intuitive, versatile and useful. Most of all, it has the potential to make development teams faster and more productive. In business, the ability to make things happen quickly often outweighs the risks.

Governance Gaps

Paving the Way for Diverse New Attack Vectors

Looking beyond anecdotal industry conversations and what we read in the media, our survey revealed a range of AI-related security issues confronting organizations.

When we asked respondents what their biggest AppSec-related concerns were when incorporating AI code generation tools into their organization, they were worried about literally everything. In fact, only 8% had no concerns. The highest level of concerns related to:

92%

AI

of respondents

cited AppSec related concerns when incorporating AI code generation tools.

60%

Hallucination Attacks

These are false data points or patterns that AI models might perceive due to adversarial inputs or misinterpretations, which can be exploited by malicious actors.

60%

Prompt Injections

Threat actors can manipulate AI models by introducing or “injecting” specially crafted prompts, tricking the system into undesired behaviors or outputs.

41%

Secrets Leakage

This is not an attack, but an accidental leak. For example, developers input their custom code into a GenAI tool to work on the code; the LLM trains itself in this custom code and regurgitates it later to other developers outside the company. This is an easy way for companies to lose control of their IP.

55%
**Data Privacy /
Intellectual Property**

GenAI tools are as susceptible to data leaks as other technologies, with the potential to leak user data, systems inputs, and more.

43%
**Insecure Coding
From GenAI**

GenAI tools are not trained in secure coding practices. They may utilize insecure libraries, packages, or simply write bad code.

45%
License Risks


The GenAI community is facing a lot of questions about the data that it uses to train its models, and whether they have a right to use that data. Therefore, it is possible that output used by an organization from GenAI is IP that they do not have a license for. This opens the business up to risk.

This is not surprising. Any new technology often lacks proper governance. Likewise, the effectiveness of security measures against AI-based attacks are still uncertain; only one-third of AppSec professionals said that such

measures have been helpful. The rest (64%) admit that they are either unsure how helpful these measures have been, or they have not helped at all.

FIGURE 02

What is your level of concern about security threats stemming from developers using AI code generation tools to write code?



However, when we zoom in on application security and look at what practitioners think about GenAI as a security threat from a coding perspective, we see somewhat of a bell curve (Figure 02).

AppSec practitioners appear to view GenAI coding as just another threat. It's something to worry about and plan for, but nothing to panic over. This is not terribly surprising given other survey data Checkmarx has reported on and considering the state of secure coding practices in the industry at large. In the 2023 Global Pulse on AppSec

Survey, Checkmarx found that only 22% of CISOs believed that their development teams were skilled in coding best practices. In the eyes of some security practitioners, the addition of GenAI isn't that much of a change from the status quo.

AI Creates New Application Security Possibilities

There is significant optimism that AI will benefit security teams. When we asked respondents where they thought AI has the biggest potential to enhance security posture; 31% said it will help security teams or developers to remediate vulnerabilities faster and 32% said AI will make it easier to tune AppSec solutions for different applications.

Security vendors are already building and deploying AI tools to help security teams become more effective. There are several potential roles for AI to play in aiding AppSec teams.

Overcoming developer skills deficits to accelerate remediation:


Most developers don't have much experience in application security. They often do not have the knowledge to identify and quickly remediate a vulnerability. Instead, they end up trying to create a net new solution, which can be difficult and time-consuming to manage. Checkmarx has typically addressed this through Codebashing, our interactive security learning and development program.

The addition of the GenAI-based AI Security Champion feature for automatic remediation to our platform allows developers to quickly interpret, and act on, security scan results. This drastically reduces the time between spotting and addressing vulnerabilities.

Addressing the skills and resource gap in AppSec teams:

Security tools require tuning and specialized skill sets. When implemented properly by a security vendor, GenAI can alleviate the need for security professionals to spend hours mastering intricate query languages. This is certainly a requirement that AppSec teams are interested in, with 17% saying correlating security data to simplify analysis would be a benefit.

AppSec teams are often outnumbered by developers 150:1. Getting AI integrations from vendors right is high on the list for AppSec practitioners. There is also quite a high level of trust, with 64% of respondents saying they place significant trust in AI tools from




Providing actionable remediation within integrated development environments (IDEs), **Checkmarx AI Guided Remediation** helps developers better understand IaC and API misconfigurations without additional resources.

Now organizations can address issues in their IaC templates faster, reduce management overhead, boost developer adoption, and deliver more secure applications faster.

FIGURE 03

Based on your view of AI in security tools today, what would you trust AI to do?



Surprisingly, security teams are also quite interested in how AI can provide additional help, with only 6% of respondents saying that they wouldn't trust AI to be involved in security actions within their vendor tools.

There is also definite interest in having AI support developers. Respondents said they would be most interested in AI support when it comes to outlining and implementing fixes, as well as support with training. However, a significant percentage (47%) indicated that they are interested in allowing AI to make unsupervised changes to code.

This result was surprisingly high, and likely reflects some wishful thinking on the part of respondents: "eventually it would be nice to automate as much of this work as possible."

We're just at the beginning of utilizing AI to help security teams. Only 4% of respondents engaged in early access trials or proof of concept to improve application security, however there is definite interest and anticipation.

Where Do We Go From Here?

7 Steps to Include Generative AI in Application Security

As our survey data shows, generative AI tools are here – whether we like it or not. They are free or relatively low cost, and development teams are already seeing the productivity and efficiency benefits gained from AI automation. However, we know that AI can also create or expose security vulnerabilities and become a powerful tool for malicious actors.

Rather than trying to put the brakes on innovation and restrict AI use in application development and AppSec, CISOs must find ways to take advantage of AI safely and mitigate its risks. With AI still in its infancy, now is the time to ensure the guardrails are in place, because where there is opportunity, there is also threat. Here are 7 best practices to get you started:

70%

of respondents

allow AI tools to be purchased by different teams on a case-by-case basis

1 Implement Effective Governance for GenAI Adoption

When Generative AI arrived, it hit hard and broke things quickly. Many departments in your organization are already adopting AI tools because they offer tremendous productivity benefits. Governance has not had a chance to catch up – but getting effective governance in place is a must for CISOs in 2024. There are parallels here to work previously done for the cloud, where the question moved from “do we put data in it at all?” to “can we keep data in X region to stay compliant?” Effective governance can reap benefits, such as a standardized internal LLM instance that keeps answers in company-context. The key next steps here are first to fully research how and why your company uses GenAI; then to work cross-functionally to identify specific tools appropriate to meet those needs that are generally recommended; and finally, to approve those technologies while creating effective security policies around their use.

2 Assess and Understand the Risks That AI Can Present

You know that AI is being taken up and implemented across your organization regardless of company policy – but do you know how? As part of partnering closely with development, your AppSec team must understand how exactly developers are using GenAI throughout the SDLC. Where and how is GenAI changing the developer workflow? Does any of it involve uploading your company's intellectual property into an LLM model? As a CISO, you and your AppSec team must map out the changes catalyzed by GenAI, and how your AppSec program will seek to mitigate resultant threats while not disrupting relations with increasingly productive developers.

3 Identify Attack Vectors Targeting AI

AI is just as new to malicious actors as it is to you – so the entire threat landscape is in flux. Within the security community, many are still trying to build comprehensive lists of known threats. OWASP has a top 10 list that Checkmarx contributes to and recommends. Begin by analyzing this list and prioritizing those your business is most vulnerable to. In addition to these known risks – we recommend that your AppSec team takes a hard look at how generative AI is being used at your company, and brainstorm potential unknown threats. This is new to everyone, and existing lists cannot be comprehensive yet.

4 Improve the Developer Experience with Secure GenAI

GenAI isn't only a challenge for CISOs and AppSec teams, it also presents opportunities for increased productivity. As a CISO, ask yourself how you can improve developers' experience with AppSec by creating a secure GenAI experience that they love. Check your vendors' roadmaps for GenAI tools to help developers perform security tasks faster. This can include AI-guided vulnerability remediation, vulnerability scans that happen within the GenAI tools themselves, and more.

5 Boost AppSec Team Productivity With GenAI

GenAI not only provides opportunities for improving developer productivity – there is also an opportunity for GenAI to address the skills and resource gap in AppSec teams. Research tools that make AppSec easier for your teams. Are there tools that help generate queries to cut down on their own coding time? Are there AI solutions to help correlate vulnerability data? Work with your key AppSec vendors to understand what they offer, and their roadmap to making your own team more effective.

6 Assess Where Your Organization Is Today

You likely have multiple teams within your organization interacting with multiple different LLMs. To start, a CISO must find out what tools are in use, by whom, and for what purposes. Until you know how and why employees are using GenAI, you won't be able to write or enforce reasonable and effective security policies that effectively mitigate risk. Do not try banning GenAI – it is proven to be ineffective. This survey found that while 15% of respondents said that their company had completely banned the use of GenAI tools, only 1 respondent out of 900 said that 0% of their company's code was AI-generated. As a CISO, focus on corralling users into safer methods of interacting with GenAI without eroding trust in the policies you want to deploy.

7 Evaluate AppSec Vendors Based on Their AI Roadmap

Everyone is talking about GenAI – but which of your vendors are doing something with it that you find useful? How are vendors planning to use generative AI? Do they have ways to integrate it into their products, services and platforms? The developers whose code you are responsible to secure will use GenAI to vastly increase their output. It is also known that GenAI does not create secure code. As a CISO you and your AppSec team must find vendors who can provide the tools your developers need to securely use GenAI, and tools to increase the efficiency of your AppSec teams in the face of ever-increasing developer output.

GenAI Is Like Every Technological Disruption That Has Come Before.

It makes employees more productive... but brings security risks. That means that there are both challenges and opportunities available. We recommend you take a calm, realistic approach to managing GenAI within your organization. Use the best practices in this document to understand how your company uses GenAI, how to mitigate its risks without alienating your colleagues, and leverage its benefits for your security teams.

If you have questions – Checkmarx is here to help. To read more about Checkmarx's vision for generative AI in Application Security, and read about our newest features, check out our [blog](#).

A purple shield-shaped icon with the letters "AI" in white, positioned at the bottom left of the page. A thin purple line extends from the top of the shield upwards and to the right, ending near the top left corner of the page.

Methodology

To get more insight into current trends in software supply chain security, we commissioned a survey of 900 CISOs and application security professionals to shed some light on their key challenges and priorities.


The survey was conducted online by Global Surveyz Research, an independent survey company. Respondents included a mix of CISOs, Deputy CISOs, VPs, Directors, and application security managers from companies in North America, W. Europe, and APAC with an annual revenue

of \$750M+, across a variety of industries, including: Banking & Finance, Insurance, Software, Technology, Engineering, Manufacturing, Industrial, and the Public Sector. The respondents were recruited through a global B2B research panel and invited via email to complete the survey. Answers to most of the nonnumerical questions were randomized to prevent order bias in the answers.

Checkmarx One

Use AI to empower developers and AppSec teams to make application security easier

[Discover How →](#)



Checkmarx

Checkmarx is the leader in application security and ensures that enterprises worldwide can secure their application development from code to cloud. Our consolidated platform and services address the needs of enterprises by improving security and reducing TCO, while simultaneously building trust between AppSec, developers, and CISOs. At Checkmarx, we believe it's not just about finding risk, but remediating it across the entire application footprint and software supply chain with one seamless process for all relevant stakeholders.

We are honored to serve more than 1,800 customers, which includes 40 percent of all Fortune 100 companies including Siemens, Airbus, Salesforce, Stellantis, Adidas, Walmart and Sanofi.