sonatype

The Risks & Rewards of Generative AI in Software Development

Since its launch late last year, ChatGPT has dominated the conversation in the tech industry and has been roundly debated.

While there is a range of opinion about the potential of generative AI tools among the developer community, there is a general consensus that it will have a huge impact on the industry, akin to cloud adoption. Software engineers are using the technology to research libraries and frameworks and write new code, while application security professionals are using it to test and analyze code and identify security issues.

To get a better sense of how generative AI is influencing and impacting the work of software engineers and the software development life cycle, Sonatype surveyed 400 Developer (DevOps) and 400 Application Security (SecOps) leaders in the United States. The findings revealed that an overwhelming majority are using generative AI today. Each metric suggests an incredible (even historic) rate of adoption and organizational effort to establish processes, though some report feeling pressured to incorporate the technology despite concerns about security risks. Respondents also raised concerns about workers potentially being replaced by AI.

Here, we look at how perception of generative AI varies depending on role, and compared how DevOps leads responded differently than their SecOps counterparts. In most areas, they were aligned in their perspectives, but SecOps respondents appear more bullish than DevOps leads on use of the technology, its benefits to them and impact on the industry. This may be traced to generative AI tools arming security experts with the ability to more easily scale their specialized skills across the organization. DevOps professionals, on the other hand, may be slightly more apprehensive due to the nature of their work – creating code, a public artifact, with unclear copyright.

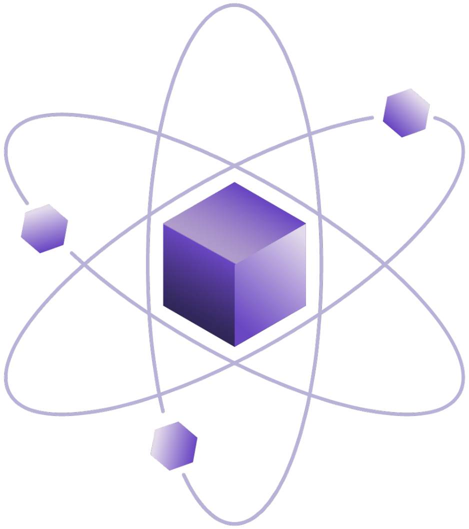

Key Findings

Where do developer and security leads see eye to eye?

74%

Feel pressure to use gen AI despite security risk

Three-quarters of all respondents (74%) said they feel pressure to use generative AI despite the perceived security challenges, with DevOps leads more likely to say that than their counterparts.

67%

Think developers should be paid for code used in AI

Both groups also agreed that creators should own the copyright for AI-generated output in the absence of copyright law (40%) and most said the organization using the code in their software should pay the developers (67%).

97%

Currently use generative AI

An overwhelming majority (97%) currently use generative AI in coding workstreams to some degree.

Among the concerns over generative AI use, security (52%) and job loss (49%) were neck-and-neck at the top of the list for both groups.

The latter references well-documented fears since the introduction of generative AI technology in late 2022, that properly trained AI platforms could ultimately displace technical roles.

Key Findings

Where do developer and security leads differ?

- Nearly half of respondents (49%) think generative AI is overhyped, with 61% of DevOps leads agreeing with that statement.
- Meanwhile, 90% of SecOps leads agreed that generative AI's impact on the industry will be similar to the impact of the cloud.
- More SecOps leads reported time savings from the use of generative AI (57% save at least 6 hours per week, vs. 47% for DevOps). With such clear productivity gains, even cost pressures will not hinder widespread adoption.
- Forty-five percent of SecOps leads have fully implemented it into the software development process, compared to 31% for DevOps.
- SecOps were also more likely to list security issues as the reason their teams don't use the technology (60% vs. 50%).

Dig into our findings to uncover in-depth patterns about generative AI usage. This data covers everything from the technology's impact and more macro concerns, to sentiment around regulation and platform integration.

Report 6 hours per week of time savings with gen AI

47% DevOps

57% SecOps

Fully implemented gen AI into software development

31% DevOps

45% SecOps

Believe security issues are the reason not to use gen AI

50% DevOps

60% SecOps

CHAPTER 1

Sentiment Among Developer, Security Leads

Developers take a more cynical view of generative AI than security leads in general, but there is consensus among the two groups that it's an important disruptive technology and they are pushed to use it despite security concerns.

Sentiments Differ Between DevOps and SecOps

Overall, DevOps leads take a more critical view of Generative AI than SecOps leads do.

Sixty-one percent of developers said the technology was overhyped, compared to 37% among security leads (Figure 1.1). Meanwhile, the majority of respondents (89%) agreed that the technology's impact on the industry will be similar to the impact of the cloud, which began soaring in popularity in the 2000s thanks to the advent of Amazon Web Services (AWS). In a matter of years, on-demand compute and storage (across public, private and hybrid setups) became the norm – and it remains a foundational part of modern computing. Generative AI, then, arguably presents a similar upside.

Figure 1.1 – Is Generative AI Overhyped?

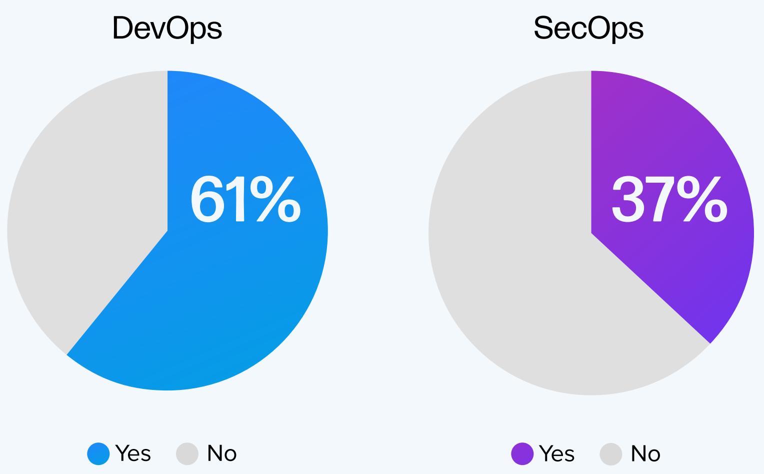

Opinions aside, organizations are pushing development teams to use new generative AI tools to take advantage of the productivity and efficiency increases they can provide. In fact, virtually all DevOps and SecOps leaders (97%) are using generative AI in some capacity — but not necessarily by choice. A majority of both groups (approximately 75%) said they felt pressure to use generative AI despite security concerns. (Fig. 1.2) This shows how organizations are racing to leverage the technology for increased productivity and, ultimately, competitive advantage.

This notion is clear among developers who play a business-critical role in driving innovation and creating real value for organizations. Developers appreciate efficiency because the faster they work, the faster they can add that value, and generative AI poses steep potential time savings.

Figure 1.2 – There Is Pressure To Use Generative AI Despite Perceived Security Challenges

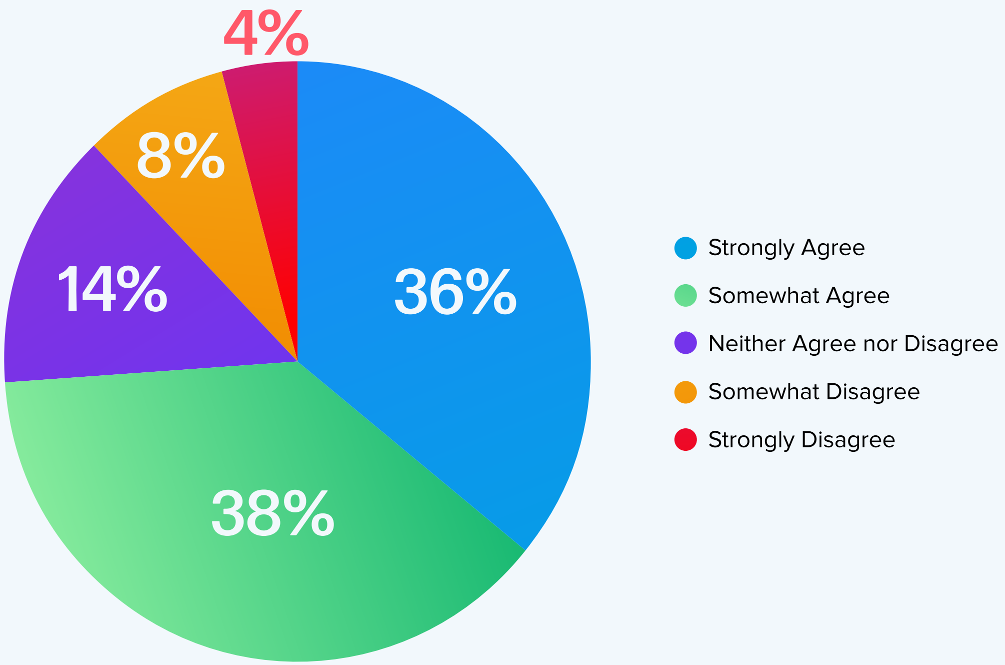

CHAPTER 2

Most Commonly Used Generative AI Tools

Engineers and security teams are frequently consulting these tools, whether for testing and analyzing, or identifying vulnerabilities. And they're doing so with precision and solution integration in mind.

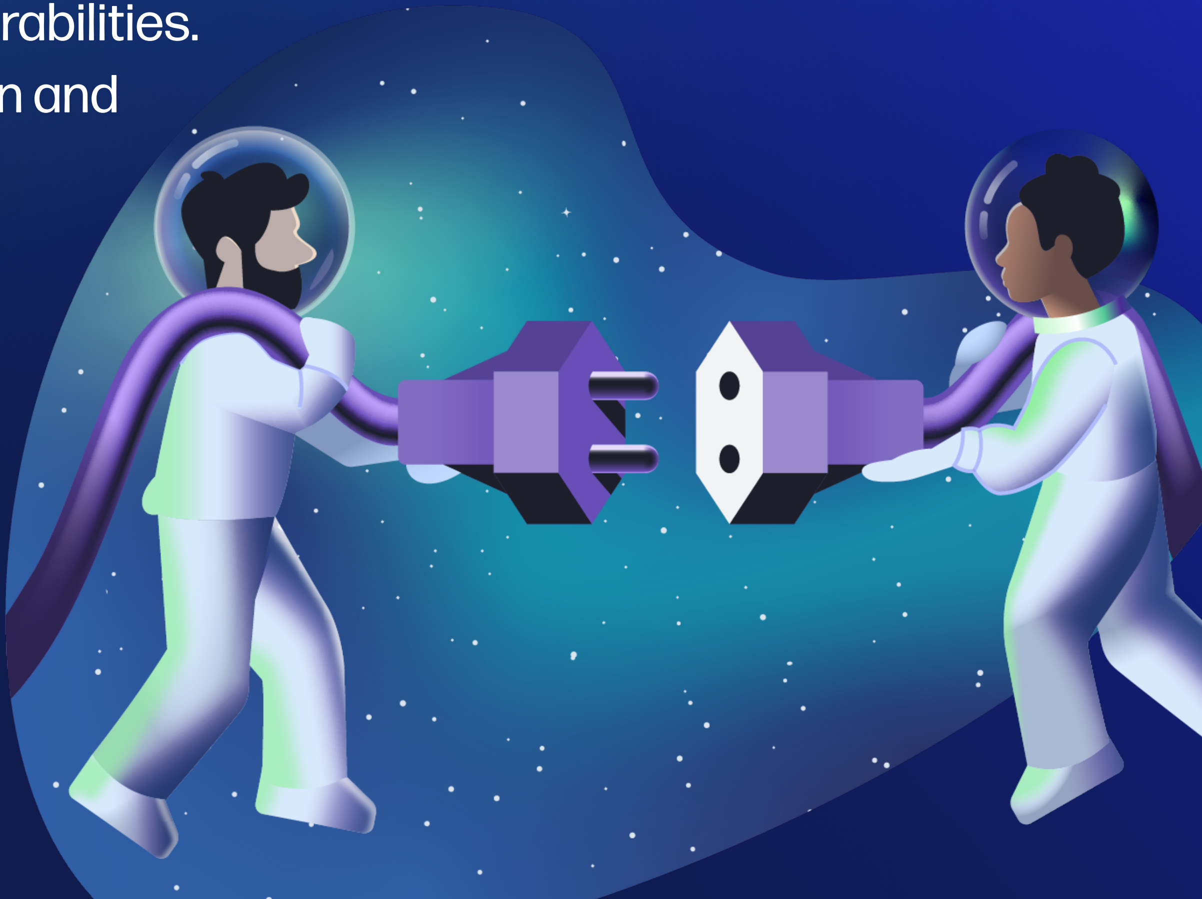

Here's What Tools Developers, Security Leads are Using

Nearly all respondents (97%) are currently using generative AI in their workstreams to some degree, with 84% using it at least a few times per week and 41% using it daily.

What's more, on average, respondents are using at least two tools, with 86% of those polled preferring ChatGPT, and 70% relying on GitHub Copilot (followed by SCM Integrations, IDE Plugins and sourcegraph Cody). (Fig. 2.1)

Figure 2.1 - Generative AI Tools Used By Teams

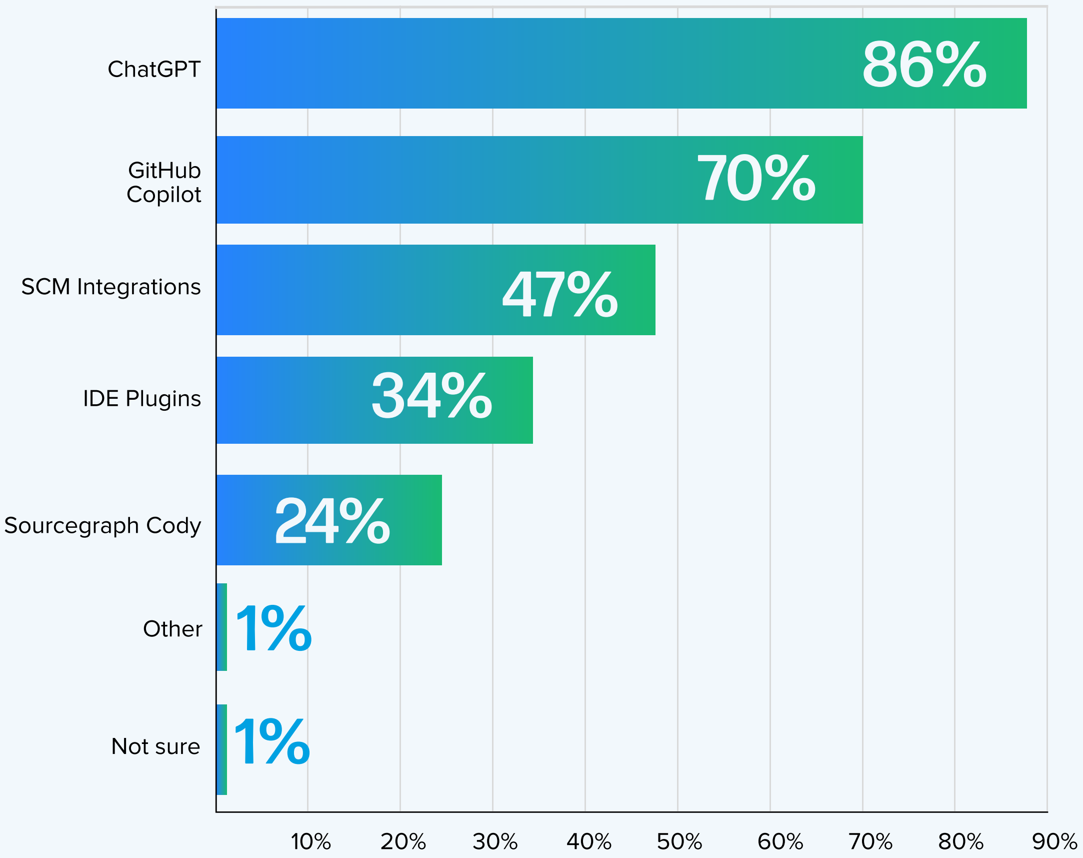

SecOps leads are ahead in fully implementing generative AI into their software development processes, with 45% saying they are “there now” compared to 31% for DevOps. Twenty-three percent of all respondents also indicated that they would have it adopted within a year. Only 2% said they had no plans to adopt the technology.

These teams are using generative AI tools for a variety of tasks, though SecOps places higher in each task category except for one: researching frameworks and libraries. For instance, when it comes to testing and analyzing, 82% of SecOps respondents said they use it (vs. 79%). However, for researching libraries and frameworks, DevOps leads the way at 63% (vs. 54%).

As far as using generative AI to identify security vulnerabilities, SecOps leads are more likely to do so than their counterparts (72% vs. 62%). (Fig. 2.2) Only 7% of all respondents said they don’t use it for identifying vulnerabilities nor do they plan to.

Figure 2.2 – Have You Used Generative AI To Detect Potential Vulnerabilities Before They Become Critical Security Risks?

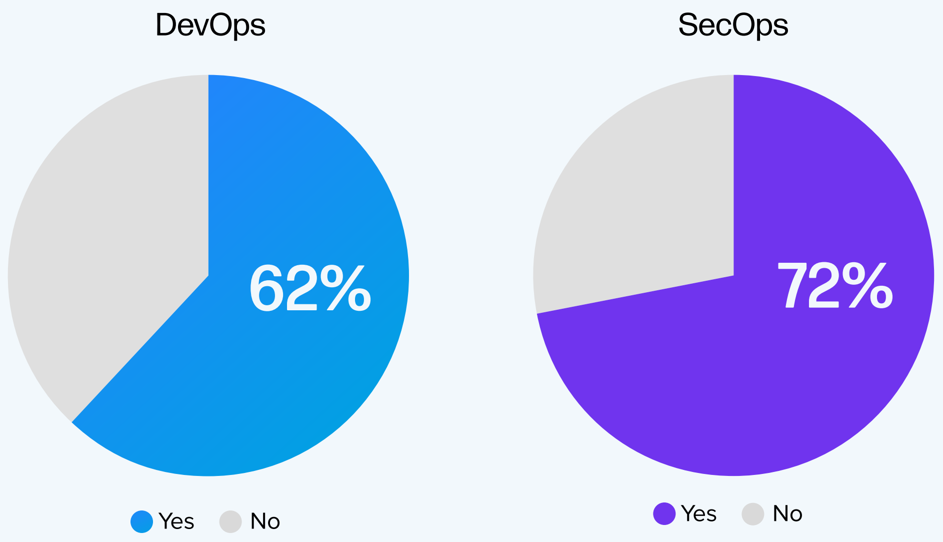

Additionally, most respondents are using behavioral AI and want workflow integration with generative AI. Nearly three-quarters of all respondents (73%) said they currently use a DevSecOps platform built on behavioral AI to protect their software development lifecycle. (Fig. 2.3) And at least 80% currently use a similar platform with AI to automate tasks in said lifecycle. (Fig. 2.4)

Meanwhile, 98% of those surveyed said they would like to see DevSecOps platforms integrate with generative AI, and more than one-third said it would be a big improvement for the software development process.

The biggest benefit from an integrated platform would be getting code to production faster, followed by easier remediation of bugs and errors, greater team collaboration/communication, and finally, making the product more secure.

Figure 2.3

73%

Use a DevSecOps platform
built on behavioral AI to protect the software
development lifecycle (All respondents)

Figure 2.4

80%

Use a DevSecOps platform
with AI to automate tasks in the software
development lifecycle (All respondents)

CHAPTER 3

Impact and Challenges

Generative AI frees up nearly one business day a week for respondents, but it's also posing "code sprawl" problems.

Saving time, but dealing with data and code sprawl

Developer and security leads reported a number of advantages to using generative AI, including time savings, but they also said dealing with the increase in data was challenging.

Asked what the most positive impacts of this technology have been, DevOps leads said faster software development (16%) and more secure software (15%). SecOps leads said increased productivity (21%) and faster issue identification/resolution (16%). (Fig. 3.1)

Overall, a majority of respondents report reclaiming approximately one business day per week, thanks to generative AI, a clear early demonstration of its utility. SecOps leads saw greater time savings than their DevOps counterparts, with 57% saying it saves them six-plus hours per week, compared to 47% of DevOps respondents saying that.

When asked about the top challenges in their use of generative AI, more than half of respondents from both groups (55%) listed data sprawl collected from unfiltered prompts.

Additional challenges included: lack of insight into how code was built (50%) and code sprawl from producing too much code (50%).

Figure 3.1 -Top 3 Positive Impacts From Generative AI (Based On Overall Respondents' Ordered Ranking)

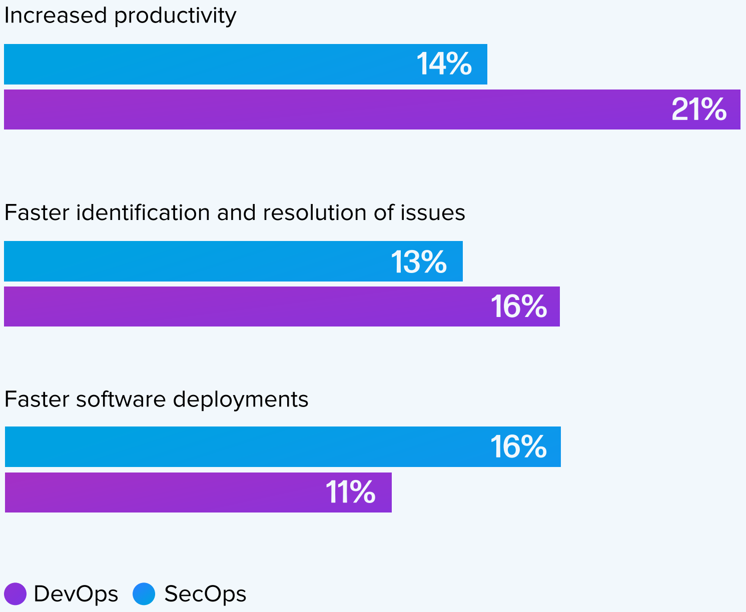

CHAPTER 4

Security Concerns

There is strong consensus among software engineers about security issues related to generative AI use, with DevOps leads being more pessimistic in their belief that the tools will lead to more security vulnerabilities in code. Both feel pressure to use it regardless of the concerns.

Pressured Despite Security Worries

Both DevOps and SecOps leads agree that there is pressure to use generative AI despite security concerns, and DevOps leads feel more strongly than their counterparts that the technology will lead to more vulnerabilities.

Seventy-seven percent of DevOps leads said they felt pressure to use the tools, while 71% of SecOps did. (Fig. 4.1) More than 3 in 4 DevOps leads said generative AI will lead to more pervasive security vulnerabilities (Fig. 4.2), and also result in more vulnerabilities in open source code specifically (Fig. 4.3).

Surprisingly, SecOps leads were less concerned here (58%) – perhaps as historically under-resourced security practitioners eye generative tools to more efficiently scale their efforts. Meanwhile, 55% of DevOps respondents said the technology would make threat detection more complex (vs. 44% of SecOps).

HOW MUCH DO YOU AGREE WITH THE FOLLOWING STATEMENTS?

Figure 4.1 - There Is A Pressure To Use Generative AI Despite Perceived Security Challenges

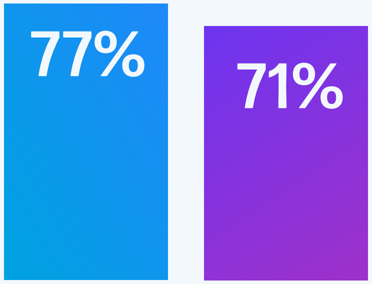

Figure 4.2 - Generative AI In Software Development Will Lead To More Pervasive Security Vulnerabilities

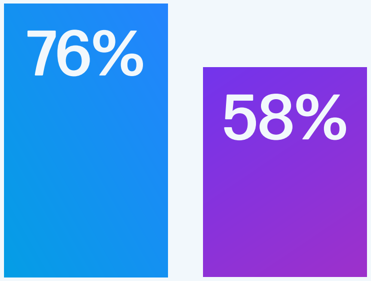

Figure 4.3 - Generative AI In Software Development Will Lead To More Vulnerabilities In Open Source Code

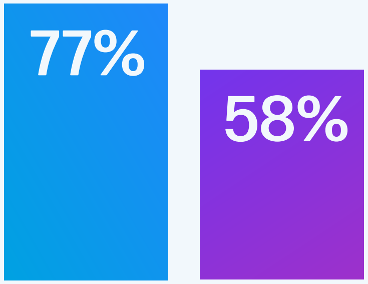

Security risks and the prospect of job loss ranked neck-and-neck when it comes to concerns with generative AI. Asked about what troubles DevOps leads most, their top worries were security and resilience risks, and that it will require special code governance (19% for each). Eighteen percent fear it may lead to layoffs. (Fig. 4.4)

Meanwhile, top concerns for SecOps leads were layoffs (1 in 5 cited this), followed by security risks (just two points lower). The latter is perhaps due to generative tools' ability to test for weaknesses at scale, a functionality that could put software quality assurance in the crosshairs. (Fig. 4.5)

#1 CONCERN ABOUT USING GENERATIVE AI

Figure 4.4 - DevOps

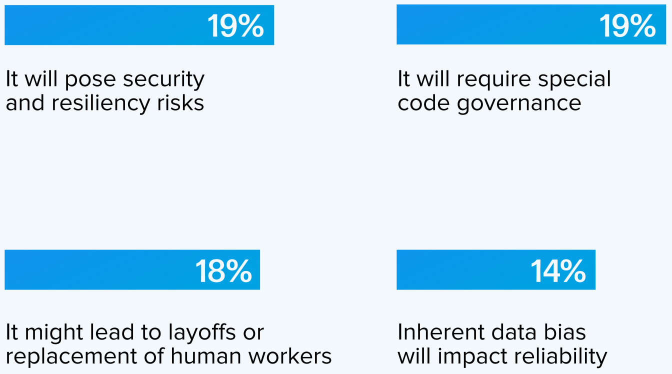

Figure 4.5 - SecOps

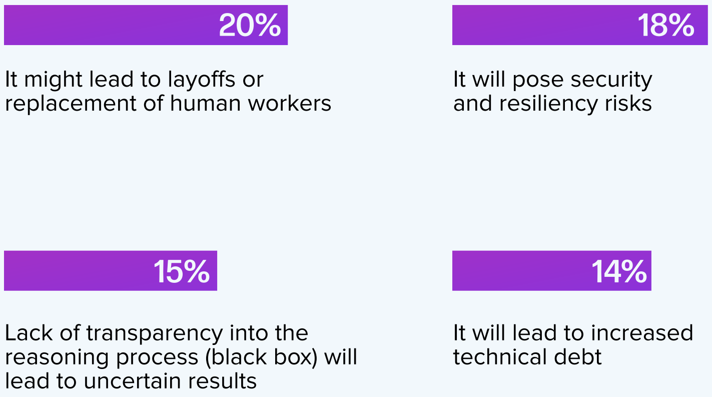

The lack of regulation in the generative AI space was no doubt a cause for concern. Half of those surveyed agreed that it poses security risks to consumers' private information and leaves their company liable, and also that it poses security risks to their company.

Forty-one percent of respondents said lack of regulation could deter developers from contributing to open source projects. This means a near majority fear that AI-platform ambiguity could be harmful to these critical libraries, and hints at a need for more comprehensive guardrails. (Fig. 4.6)

A mere 3% said they were not using generative AI and they cited security risks and unreliability of the tools.

Figure 4.6 - Concerns About Absence Of Regulations Related To Generative AI And Open Source

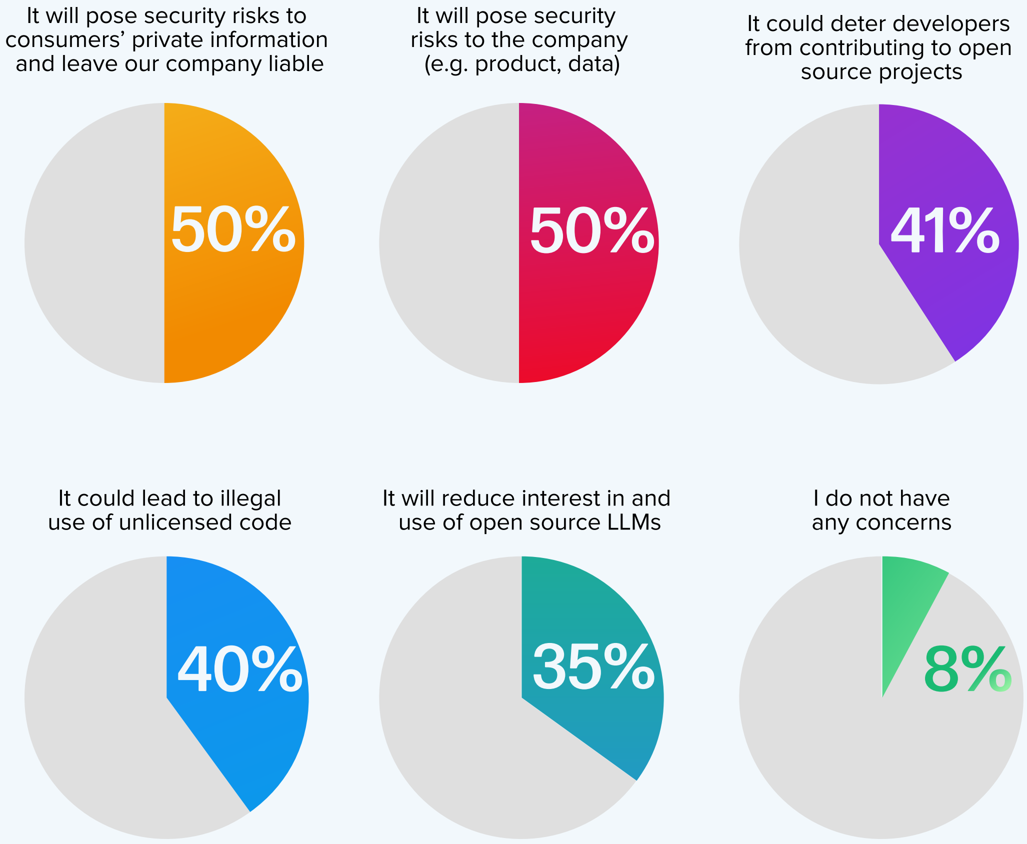

Chapter 5

Using Generative AI Responsibly

Organizations are addressing concerns with new generative AI policies and awaiting regulation – at a time when various stakeholders, the federal government included, are mulling possible policies. What's more, a surprising number of respondents do readily believe that the government has some role to play in reining in powerful AI functionality.

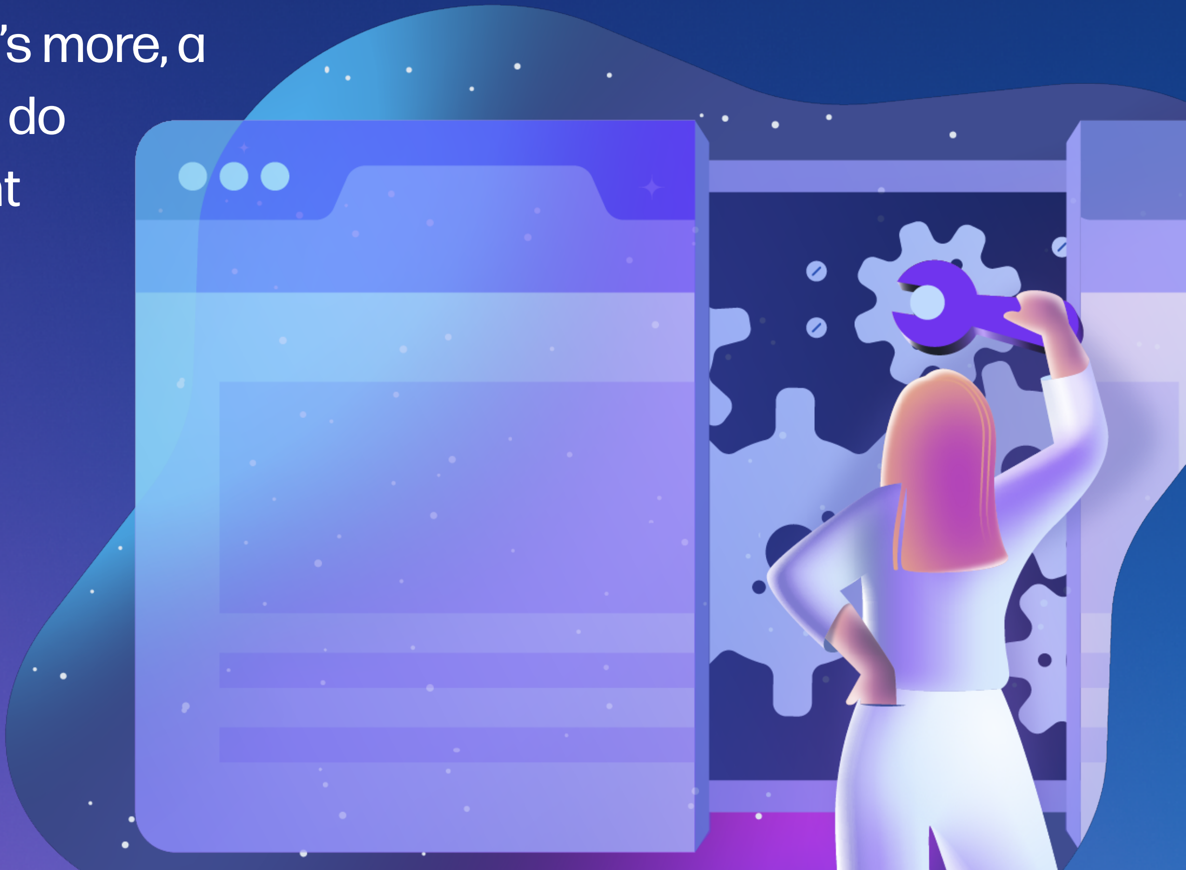

Looking Toward Policies, Regulation and Law

Seventy-one percent of those surveyed said their organization has policies in place for generative AI use, while 20% said their organizations were in the process of developing them.

Asked who they believe is responsible for regulating the use of generative AI, 15% of DevOps leads said the government, compared to SecOps leads (6%). Nearly one-quarter of DevOps respondents said companies should regulate – vs. 14% for SecOps leads. Finally, 59% of DevOps leads and 78% of SecOps leads said both the government and individual companies should be in charge of regulation. A scant 2% said the technology should not be regulated at all. (Fig. 5.1) The consensus here is that technical teams largely want more policies and procedures in place as generative AI use continues to climb.

Figure 5.1 - Who Is Responsible For Regulating Use Of Generative AI?

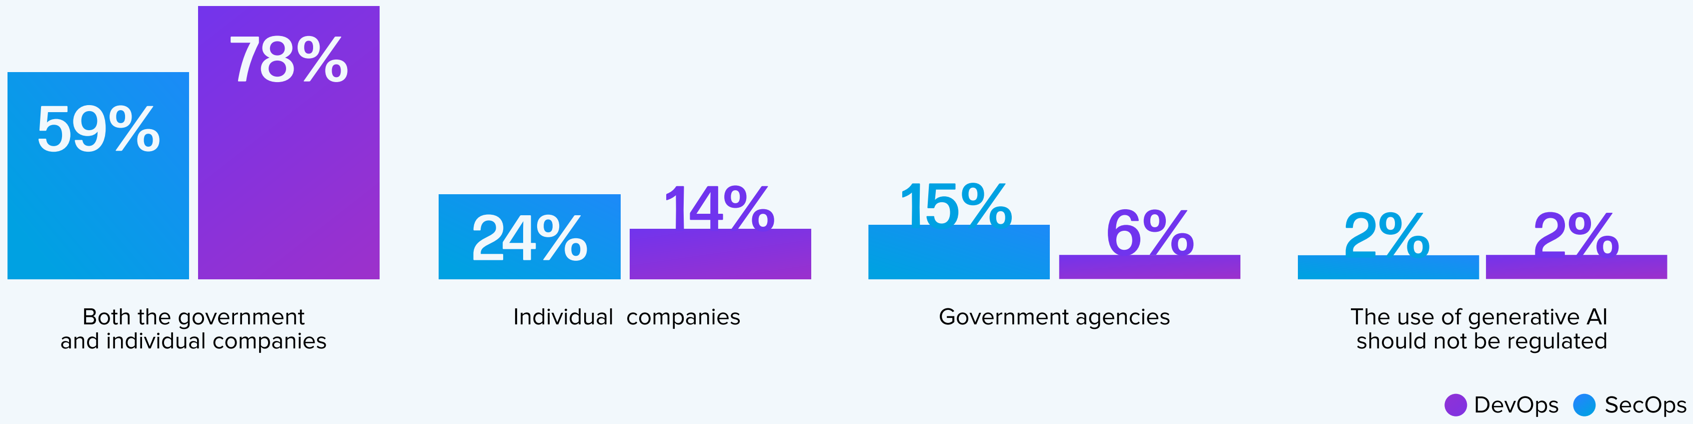

Additionally, copyright ownership for AI-generated content based on open source software remains up for debate, which creates a legal limbo for developers with plagiarism claims against LLMs. Both groups agreed that creators should own the copyright for AI-generated output in the absence of copyright law (40%) and both agreed that developers should be compensated for the code they wrote if it's used in open source artifacts in Large Language Models (LLMs) (90%). (Fig. 5.2)

Asked who should be responsible for compensating developers, a solid majority of both groups (67%) said the organization using the code in their software should pay the developers, while only 33% believe the organization that built and trained the LLM should pay.

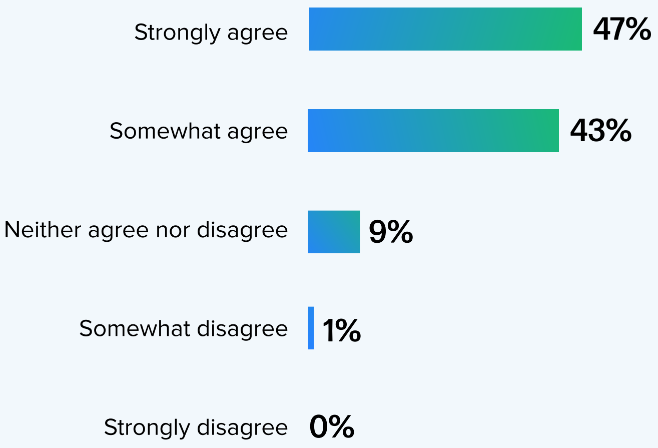

Figure 5.2 – In The Absence Of A Copyright Law For The Use Of Open Source Artifacts In LLMs, I Believe Developers Should Be Compensated For The Code They Wrote.

WHO IS RESPONSIBLE FOR COMPENSATING DEVELOPERS?

63% The organization using the code in their software development

33% The organization that built and trained the LLM

Conclusion

It's still relatively early stages for generative AI, but already it's making a difference for software engineers struggling to develop, test and audit software, and for application security professionals trying to stay on top of vulnerabilities and scale their efforts across their organization. They are seeing faster code and security fixes, more productivity and significant time savings. But they report feeling pressure to use the technology despite security concerns that it will lead to more pervasive vulnerabilities, particularly in open source code.

Along with those security concerns, they're worried about generative AI taking over their jobs and that their work will be used in LLMs. The concern, more specifically, is that code will be plagiarized (without compensation) due to current uncertainty around whether AI-generated content is covered by copyright law.

The industry has a lot to sort out to address the risks this new technology poses. Strong internal governance policies, legal protections and regulation can help mitigate these issues. As far as generative AI replacing developers and application security experts, however, most industry experts agree that the technology will not be able to apply the level of creative and strategic thinking that knowledgeable and trained humans can bring to the software supply chain.

Methodology

Sonatype commissioned research panel provider Sago to conduct a survey of 400 DevOps leaders and 400 SecOps leaders in the United States whose responsibilities involve software development, coding and developer operations or application security, threat intelligence and analysis, and security operations. The web-based survey was fielded July 12-21, 2023. The margin of error is plus or minus 3.5 percentage points.

sonatype

Sonatype is the software supply chain management company. Recognized by globally renowned analysts as a leader in the industry, Sonatype enables organizations to innovate faster in a highly competitive market. We allow engineers to develop software fearlessly and focus on building products that power businesses. Sonatype researchers have analyzed more than 120 million open source components – 40x more than its competitors – and the Sonatype platform has automatically blocked over 145,000 malicious components from entering developers' code. Enabling high-quality, secure software helps organizations meet their business needs and those of their customers and partners. More than 2,000 organizations, including 70% of the Fortune 100 and 15 million software developers, rely on our tools and guidance to be ambitious, move fast and do it securely. To learn more about Sonatype, please visit www.sonatype.com.

Headquarters
8161 Maple Lawn Blvd, Suite 250
Fulton, MD 20759
USA • 1.877.866.2836

European Office
168 Shoreditch High St, 5th Fl
London E1 6JE
United Kingdom

APAC Office
60 Martin Place Level 1
Sydney, NSW 2000
Australia

Sonatype Inc.
www.sonatype.com