THE LINUX FOUNDATION | Research

OpenSSF
OPEN SOURCE SECURITY FOUNDATION

State of AI Impact on Software Security and Cybersecurity 2026

Practitioner Perspectives on Risk, Readiness, and Opportunity

Sam Boysel, The Linux Foundation

Foreword by
Laura Guazzelli, *OpenSSF Security Architect, AI / ML / LLM / Agentic Systems*

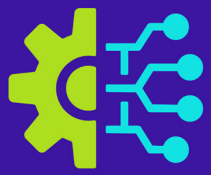
November 2026

activestate

LF AI & DATA

State of AI Impact on Software Security and Cybersecurity 2026

Expansion of AI tool use is both the **primary risk driver** and **a key defense mechanism** for software security.



Experts report that supply chain attacks, prompt injection, vulnerability chaining, automated exploit creation, and maintenance burdens are top software security risks exacerbated by AI.





82% of security practitioners agree that **overreliance on AI tools weakens developers' security skills.**



AI coding tools lead **3 in 5 teams** to **ship code faster than security reviews can keep up.** These tools **introduce transitive dependencies** that are **reviewed only half the time.**

Preparedness lags concern levels for virtually **every AI-related software security threat.** Largest gaps: **automated exploit generation, dependency hallucination,** and **increased maintenance.**



Proactive cybersecurity coincides with preparedness: Teams with dedicated roles, policy, and properly implemented tooling report being more prepared for AI-related security threats.



Ownership and policy are underdeveloped: Only **30%** of teams have a **dedicated role responsible for AI-related threats.** Only **1 in 3** have a formal AI cybersecurity policy.



AI-assisted software security tools boost security development speed by **17%**, but **false positives** from vulnerability scanning **increase workloads by 30%.**





Open source foundations are **trusted sources** of AI cybersecurity guidance (**64%** of respondents). Most practitioners (**83%**) stay informed via convenient channels: **blogs**, **newsletters**, or **podcasts.**



Top barriers to implementing better software security practices: **rapid pace of change, limited resources**, and **lack of clear practical guidance** for AI-specific threats

Top investment priorities for software security: **improved open source tooling, supply chain robustness,** and access to **effective local LLM/SLMs.**



AI-assisted software security tools meaningfully **improve overall security posture** for **70%** of practitioners. **30%** of vulnerabilities discovered in the past 12 months **leveraged AI.**



Copyright © 2026 **The Linux Foundation** | November 2026. This report is licensed under the **Creative Commons Attribution-NoDerivatives 4.0 International Public License**
Please cite this research report when using the infographics. See citation guidance on the report's back page.

Contents

Foreword 4

Executive summary 5

Introduction 6

How AI is changing the cybersecurity landscape 7

Status of software security and cybersecurity in the AI era 11

Impact of AI on development and security workflows 18

Barriers and investment priorities 22

Opportunities for AI-assisted security tools 24

References 27

Conclusion 28

Methodology 30

About the author 35

Acknowledgments 35

Foreword

AI is changing the daily work of software developers and security practitioners. AI is involved in every stage of the Software Development Lifecycle (SDLC); it helps teams write code, generate tests, automate tasks, and find vulnerabilities faster. The speed and efficiency come at a cost, as it creates more work for the people who review that code and maintain the software on which we all depend.

OpenSSF sponsored this research to ground that discussion in practitioners' experience. *State of AI Impact on Software Security and Cybersecurity 2026* looks at where teams see risk, how prepared they feel, and where AI is already helping.

The most pressing finding is that 58% of teams ship code faster than security review can keep up with. AI coding tools add to the load. 60% of respondents say these tools increase their external dependencies, and teams review transitive dependencies only about half the time. Faster development needs to be matched by more review capacity and quicker remediation.

The research also shows the value of clear ownership. Only 30% of teams have a dedicated AI cybersecurity role. Teams with a dedicated role or a formal AI cybersecurity policy report better preparedness for every threat measured. That doesn't prove causation, but it tells leaders where to start: decide who owns AI cybersecurity and what's expected of those individuals.

AI is also part of the answer. Practitioners credit AI-assisted tools with finding about 30% of the vulnerabilities they identified last year. The same tools add work: false positives increase workload by 30%, against a 17% speedup. Getting real value from these tools means budgeting time to validate findings and keeping reviewers' skills sharp. 82% of respondents agree that overreliance on AI weakens developers' security skills.

Many projects can't hire dedicated specialists. Shared open source tooling and practical guidance can close some of that gap. Investing in the communities that maintain our software is investing in our own security. While open source organizations serve as trusted authorities for guidance on AI cybersecurity (cited by 64% of participants), practitioners overwhelmingly rely on routine informational channels—such as blogs, newsletters, and podcasts—to stay informed (83%).

I encourage readers to take one question from this report back to their teams: **where is software development outpacing review, and who is responsible for closing that gap?** Thank you to the practitioners who shared their experiences, the Linux Foundation Research team, and our fellow sponsors, ActiveState and LF AI & Data.

Laura Guazzelli, *Security Architect - AI/ML/LLM/Agentic Systems, OpenSSF*

Executive summary

The rapid integration of AI into software development is reshaping software security and cybersecurity in general, creating a complex landscape defined by heightened risks, evolving team readiness, and powerful new defensive tools. While AI adoption introduces significant potential for new vulnerabilities, it also brings capabilities required to counterbalance rising threat levels. AI even holds out the promise of developing software that is exceedingly hard to successfully attack. A survey of 308 cybersecurity practitioners affiliated with teams in both organizations and open source software (OSS) projects reveals key insights for the way the rise of AI technologies influence software security and cybersecurity.

Expansion of AI tools stands as the primary driver of new software security risks, risking inclusion of insecure code and dramatically expanding attack surfaces. AI coding tools have led 60% of organizations to introduce more external dependencies and teams review transitive dependencies only half the time. 82% of practitioners warn that overreliance on AI weakens developers' core security skills. Experts cite supply chain attacks (93%), prompt injection (90%), and heavy maintenance burdens (87%) as top concerns. Across every evaluated threat, practitioner concern outpaces preparedness, with the widest readiness gaps appearing in automated exploit generation and increased maintenance overhead. Practitioners cite rapid change and limited guidance as top barriers to improving security posture for AI threats.

Teams are actively working to bridge this preparedness gap, though adoption of formal safeguards remains incomplete. Dedicated AI software security and cybersecurity roles, clear policies, and proper tooling directly correlate with higher threat readiness; however, only 30% of teams currently maintain dedicated software security roles for AI threats, and just one-third have implemented formal AI policies. As a result, 58% of teams now ship code faster than security reviews can evaluate, forcing most to triage alerts rather than fix them immediately. Respondents prioritize improved open source tooling, stronger AI supply chain security, effective local LLMs and SLMs, and better training and education on software security best practices in the age of AI.

Despite these pressures, AI-powered security tools present a crucial opportunity to neutralize emerging threats. AI detection tools accounted for 30% of identified vulnerabilities over the past year, with 69% of practitioners reporting meaningful improvements to their overall security posture and 72% viewing AI as a necessary complement to existing defense practices. While these tools accelerate security tasks by 17%, high false-positive rates can increase workloads by 30%.

To navigate this landscape, organizations must move beyond reactive triage and establish a balanced defense. Dedicated AI software security roles, formal AI usage and security policies, and properly implemented AI-assisted software security tooling are characteristic of well prepared teams. While AI tools offer essential speed and detection capabilities, they cannot replace fundamental security hygiene. Long-term resilience requires leadership to prioritize dedicated governance over AI software security concerns, strengthen supply chain safeguards, and deploy AI selectively to augment, rather than substitute, developer expertise.

Introduction

AI is becoming part of ordinary software development and cybersecurity work, creating new tensions between these two rapidly evolving practices. Developers use AI tools to generate code and accelerate delivery. This creates new challenges and threat vectors from a security perspective: AI tools can introduce insecure or malicious code, expand dependency graphs, accelerate attacks, and create new demands for review and maintenance. However, software developers and security practitioners can also use AI to identify vulnerabilities, support remediation, and analyze increasingly complex systems. AI-impacted secure software development is therefore not only a question of responding to new risks and adopting new tools; it is a question of how organizations and open source projects adapt their responsibilities, policies, workflows, and skills.

This report examines the current state of AI's impact on the software security and cybersecurity[1] landscape through a survey of 308 qualified security practitioners working across organizations and open source software (OSS) projects[2]. We asked respondents how AI is changing security risk, which threats concern them, how prepared they feel, and what incidents they have encountered. We also examined the roles and policies teams use to manage AI cybersecurity concerns, the effects of AI on development and alert-handling workflows, the barriers that limit adoption of better practices, and the opportunities and limitations of AI-assisted security tools.

The report begins by identifying the forces reshaping the cybersecurity and software landscape, including increased AI tool use, vulnerability chaining, dependency growth, and concerns about the long-term effect of overreliance on AI. It then assesses the current status of AI software security and cybersecurity in organizations and OSS projects by comparing concern, preparedness, and reported experience across major threat categories. The next section considers the organizational foundations of readiness, including security ownership and formal policy, as well as the effect of AI on shipping speed, alert handling, and security debt. The final sections address the practical barriers and investment priorities identified by practitioners, followed by an assessment of AI-assisted security tools. That assessment considers where these tools are used, whether they improve security posture, the productivity gains they provide, and the false-positive workload that can reduce those gains.

Taken together, the findings describe an ecosystem under pressure but with concrete opportunities to improve readiness through clearer ownership, stronger practices, targeted investment, and responsible use of AI. Throughout the report, we identify action items and policy recommendations for teams to confidently adapt to the new AI threat landscape.

1 The focus of this report is on the way in which AI technology creates new security challenges for software developers and cybersecurity practitioners who are tasked with ensuring the code they both ship and use is secure. For the sake of brevity, we will use terms like "AI software security" and "AI cybersecurity" to refer to this focus throughout this report. The survey instrument used the somewhat nebulous term "AI security" to refer to the same concept, reflected in some questions and response options.

2 Respondents were asked if their responses were based on their experience working in organizations or for OSS projects. Unless otherwise noted, the responses in this study reflect the pooled average of these two perspectives. 73% of respondents answer from the perspective of their work in organizations (Figure 23).

How AI is changing the cybersecurity landscape

Cybersecurity challenges evolve as the use of AI becomes ubiquitous in software development and deployment. A rapid expansion of AI-generated code and AI systems has increased attack surfaces, strained conventional security risks, and introduced new threat vectors. What factors are driving changes in AI software security risk? How does AI exacerbate traditional security risks? Does increased AI reliance lead to even more fundamental challenges to cybersecurity in the long run? In this section, we use survey evidence to explore how and why AI is changing the security landscape. By understanding the mechanisms driving shifts in risk profile, organizations and practitioners can better anticipate and mitigate emerging threats.

Expansion of AI tool use is the primary risk driver

Above all, **the explosion of AI tool use and AI-generated code is the most commonly cited factor driving changes in AI software security risk** (72% of survey respondents) (Figure 1) (top 5 shown). Another important risk factor are **expanded attack surfaces** that arise from this same surge in AI adoption (51%).

These factors are often related: AI coding tools create a deluge of new code that itself can be vulnerable, but those additions often expand the dependency graph, thereby increasing the potential for vulnerabilities to be introduced and exploited through those new dependencies. A tendency to ship code faster than security review can keep up further compounds this risk. Maintainers for top OSS projects are often flooded with AI-generated contributions and vulnerability reports, straining the limited bandwidth they have for thorough review and timely remediation.

FIGURE 1

Main factors contributing to changes in AI software security and cybersecurity risk over the past 12 months



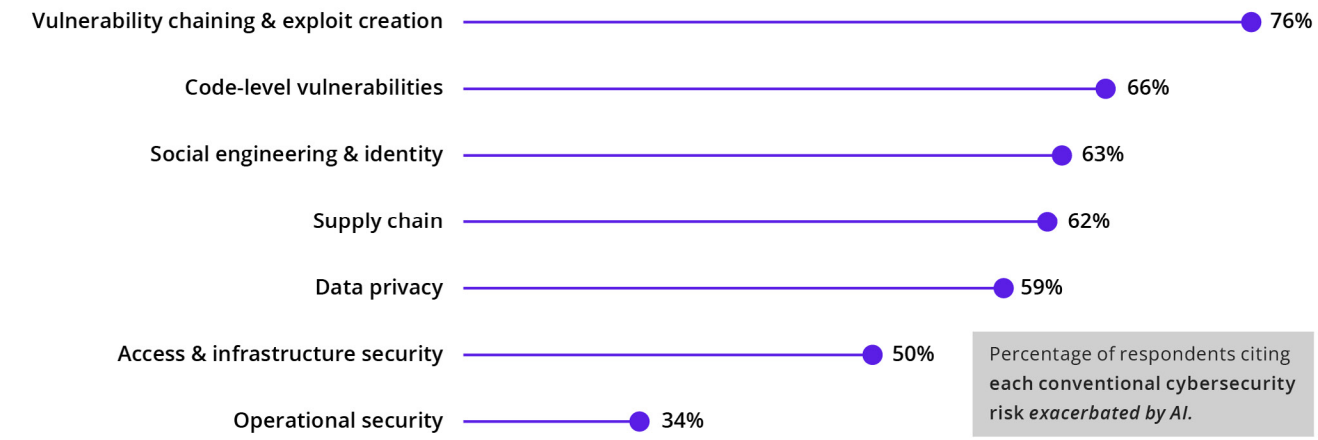
SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q20. NUMBER OF RESPONDENTS = 194. TOTAL RESPONSES = 591. TOP FIVE MOST FREQUENTLY CITED FACTORS DISPLAYED.

AI increases risk of vulnerability chaining and exploit creation

It's important to first establish which conventional security risks AI exacerbates. Of particular concern is the effect of AI-powered vulnerability chaining and exploit creation, cited by 76% of AI cybersecurity experts (Figure 2[3]). For example, AI can assist attackers in identifying sequences of defects that, when combined, enable more sophisticated and damaging exploits than would be possible by exploiting individual defects in isolation. Adding to this concern is that what once required significant manual effort and expertise can now be automated and accelerated with AI, lowering the barrier for attackers to execute complex attacks.

FIGURE 2

Traditional security risks most exacerbated by AI



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q19. NUMBER OF RESPONDENTS = 115. RESPONSES FROM PRACTITIONERS WITH ADVANCED OR PROFICIENT CYBERSECURITY EXPERTISE.

3 Experts are defined as respondents who identified their cybersecurity expertise level as "proficient" or "advanced".

AI-suggested external dependencies increase attack surface

Software supply chains are a cornerstone of modern software development but inherently balance development convenience or efficiency with additional risk. AI coding tools tend to be quite verbose, directly increasing attack surfaces by introducing more lines of code that often rely on third party dependencies. These dependencies are just the tip of the iceberg, as they can bring in further indirect or transitive dependencies, each adding potential vulnerabilities and increasing the overall attack surface.

60% of organizations and projects say that AI coding tools increase the number of external dependencies in their codebases versus only 16% who say they decrease them (Figure 3). This large net increase in dependencies highlights the growing attack surface introduced by AI coding tools and underscores the importance of robust dependency management practices.

If AI coding is increasing transitive dependencies on average, are teams reviewing them thoroughly? The study finds that **projects and organizations review transitive dependencies only about half of the time (53%)** (Figure 3). This lack of scrutiny opens codebases up to additional linked threats: supply chain attacks, slopsquatting, and other forms of dependency-related vulnerabilities. Consequently, organizations and projects should start by controlling where dependencies come from, sourcing components and their transitive dependencies from a vetted, policy-governed repository with verified provenance, and back that up with regular audits, automated dependency scanning, and strict review processes for new dependencies.

FIGURE 3

AI coding tools generally increase external dependencies — which often go unreviewed

60% of security practitioners who find that **AI coding tools *increase*** the number of **external dependencies** in their codebase

16% of security practitioners who find that **AI coding tools *decrease*** the number of **external dependencies** in their codebase

53% average **likelihood of reviewing external indirect (transitive) dependencies**

SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q28 (NUMBER OF RESPONDENTS = 176). Q29 (NUMBER OF RESPONDENTS = 153).

Action item

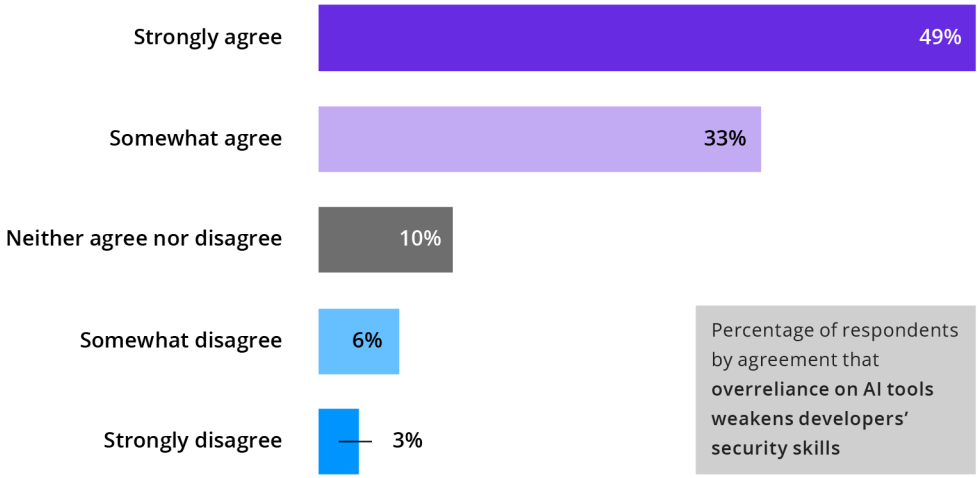
- Start with a trusted source. Pull dependencies from a vetted, policy-governed repository so developers and AI coding tools get an approved path by default.
- **Dependency management** should be prioritized, with **regular audits, automated dependency scanning, and strict review processes for new dependencies** to mitigate the risks introduced by AI-suggested external dependencies.

Overreliance on AI tools weakens developers' security skills

Overreliance on AI tools can create more systematic security risks in the long run. **82% of respondents agree that overreliance on AI tools weakens software developers' security skills**, potentially leading to a decline in the overall security posture of organizations and projects over time (Figure 4).

FIGURE 4

82% of respondents agree that overreliance on AI tools weakens developers' security skills



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026.
Q23. NUMBER OF RESPONDENTS = 196.

Skill erosion jeopardizes the long-term security capabilities of development teams, making it crucial for organizations to balance the use of AI tools with ongoing training and skill development. This is especially concerning for OSS where components are often widely reused and integrated across multiple projects, amplifying the potential impact of weakened security skills.

Action items

- Organizations and projects ought to place policies in place to **reinforce the notion that AI tools are *complements*, not substitutes, to standard cybersecurity best practices.**
- Experienced developers should be encouraged to **mentor new team members** and help **facilitate knowledge sharing**.
- New developers have a **responsibility to invest in their skills first, tools second**.
- Organizations should **foster a culture of security awareness** to mitigate the risks associated with overreliance on AI tools.
- Using systems like **artifact repositories and expert vetted supply chains can create a security safety net** to shore up gaps.

Status of software security and cybersecurity in the AI era

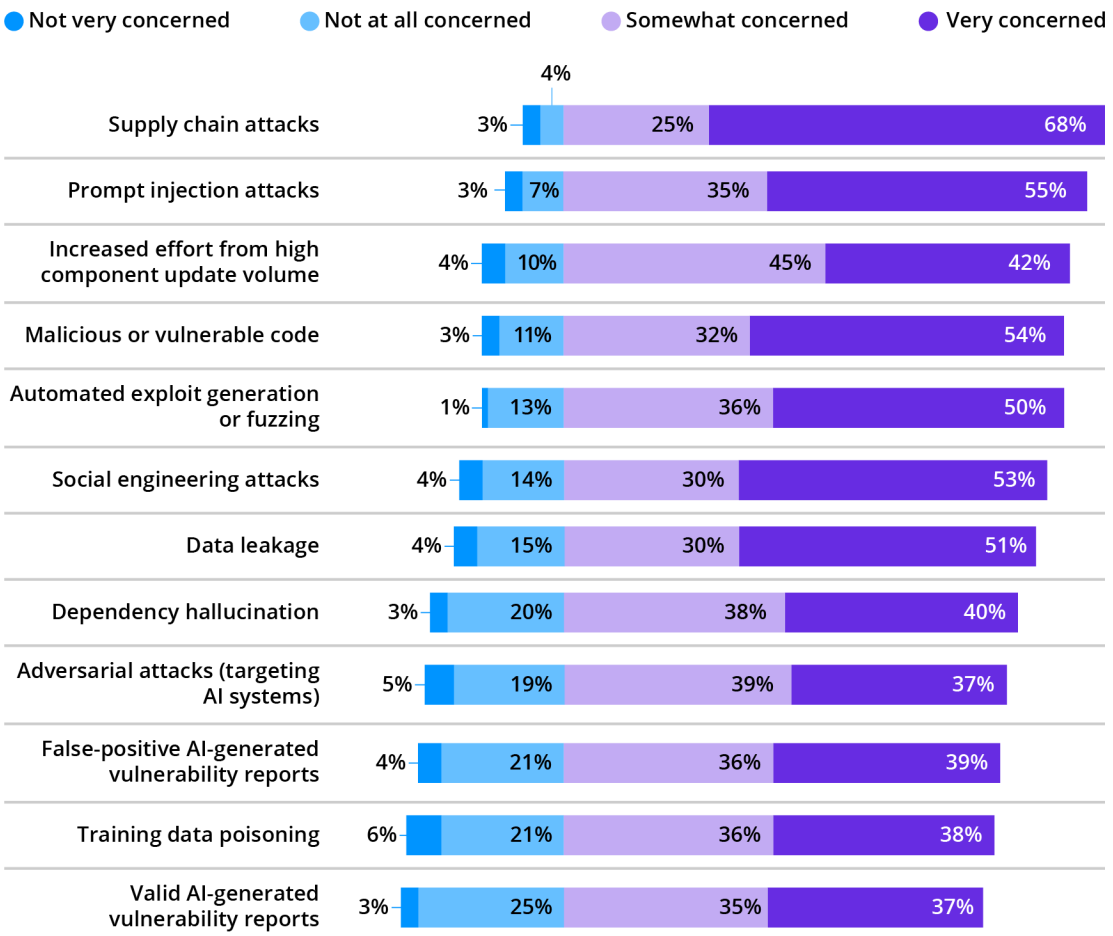
We next explore the current state of AI software security and cybersecurity in OSS projects and organizations. Organizations and practitioners vary in terms of threat perception, preparedness, and personal experiences. We explore how our pool of respondents vary along these dimensions for a common set of emerging security threats driven by AI. What AI threats are top of mind for practitioners? How prepared are they to address them? And how frequently have they encountered them?"

Experts are most concerned about supply chain attacks, prompt injection, and maintenance effort

The rapid pace of change in the AI landscape makes it especially challenging for teams to stay aware and ahead of emerging threats. What AI threats are cybersecurity experts most concerned about? **Supply chain attacks top the list of AI software security threats** (93%, of which 68% are very concerned), followed closely by prompt injection attacks (90%), increased maintenance effort due to high component update volume (87%), and malicious or vulnerable code (86%) (Figure 5).

FIGURE 5

Most concerning threats cited by cybersecurity experts



Percentage of **cybersecurity experts** by level of concern for each **security threat influenced by AI**

SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q12. NUMBER OF RESPONDENTS = 116. SUBSET OF RESPONDENTS WHO INDICATE "PROFICIENT" OR "ADVANCED" AI CYBERSECURITY EXPERTISE.

Action items

- **Focus on the supply chain first.** The software supply chain is a focal point for many cybersecurity risks. Teams should prioritize identifying what they use and depend on in production, where it comes from, and how well it is secured.
- **Eliminate unnecessary dependencies.** Some external components are not runtime critical, are also provided by necessary dependencies, or can easily be replaced with in-house functionality. One less dependency means one less factor beyond the purview of the security team.
- **Counter prompt injection attacks.** Developers must routinely consult documentation and other materials, but attackers may attempt to insert malicious instructions in them for an AI to implement. Use AI tools that filter, extract, and take steps to counter such attacks.
- **Review generated code.** AI can generate code, but it is often vulnerable. Use both AI and human review to counter the insertion of malicious or vulnerable code.

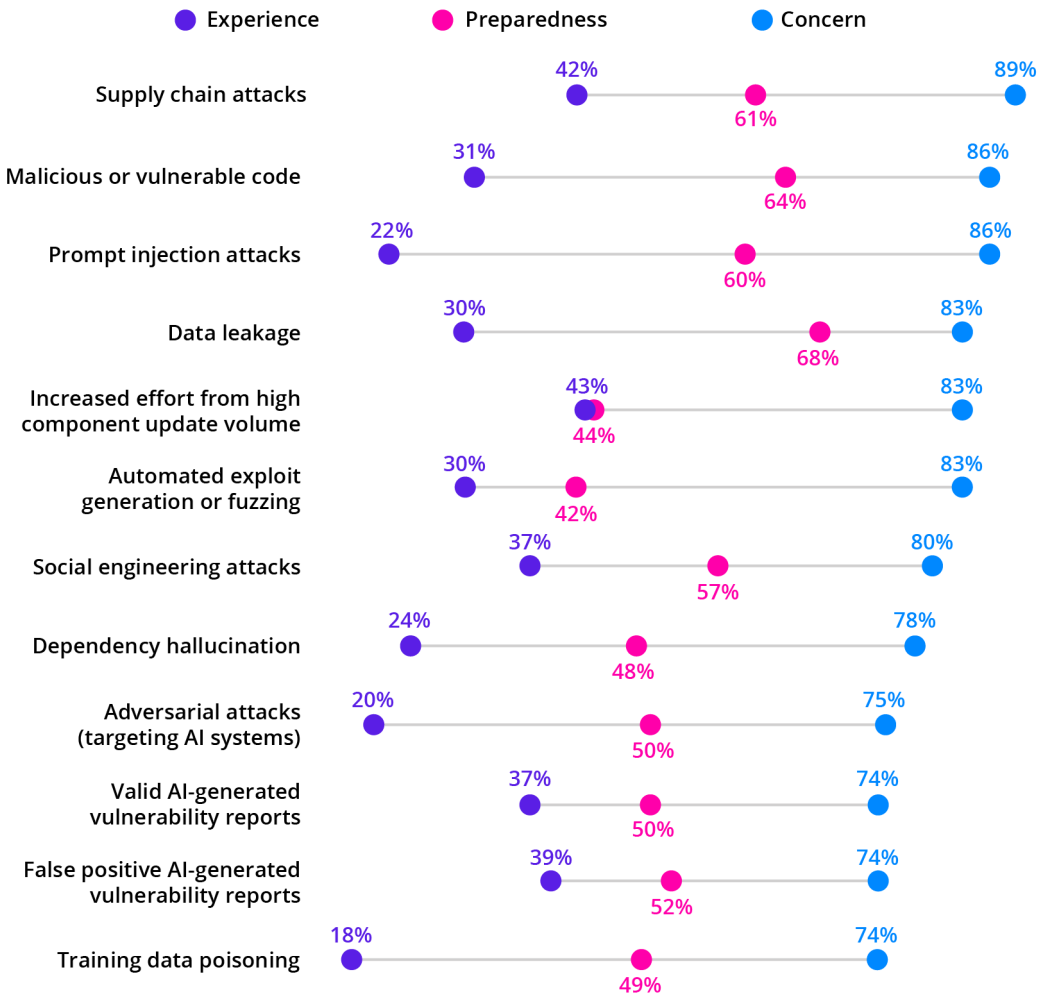
Preparedness lags concern for most AI threats

Practitioners are actively dealing with emerging threats and naturally vary in terms of which threats have the greatest impact within their organization or OSS project. In Figure 6, we summarize the top threats in AI software security by level of concern, preparedness, and actual incident experience.

Across every threat evaluated in the study, concern consistently outpaces preparedness, which in turn outpaces actual threat incidence, revealing critical readiness gaps across virtually every major AI software security risk. While concern is high across the board, the severity of these gaps varies significantly depending on the nature of the threat. The widest gap appears in automated exploit generation and fuzzing, which shows a 41 percentage point difference between concern and readiness. Similarly, while supply chain attacks remain a top overall concern, practitioners feel the operational strain most acutely around component update volume. High update burdens affect 43% of respondents, creating a 39 percentage point gap between concern (83%) and preparedness (44%). In contrast, threats such as data leakage feature much smaller readiness gaps, largely because mature technical solutions and clear operational guidelines are already available.

FIGURE 6

PRACTITIONER CONCERN, PREPAREDNESS, AND ACTUAL EXPERIENCE FOR DIFFERENT AI THREATS TO SOFTWARE SECURITY



% of respondents by **concern, preparedness, and personal experience for each security threat** influenced by AI in their project or organization

SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q12, Q13, Q14. NUMBER OF RESPONDENTS = 207.

Action items

- Development teams are feeling **underprepared for most AI threats to software security.**
- Attention should **focus on the largest preparation gaps**, such as managing the **increased maintenance labor** required by high component update volumes and **automated exploit generation or fuzzing.**
- There isn't a purely technical solution for overworked maintainers but the risks are severe. Teams should **consider holistic solutions that involve improved processes, properly implemented tooling, and increased resource allocation.**

Proactive cybersecurity coincides with preparedness: dedicated roles, policy, and tooling

What do teams well prepared for AI software security threats do differently? One approach is to adopt proactive practices that anticipate and mitigate potential risks before they materialize, such as having dedicated roles security related to AI, formal policies, and properly implemented AI-powered cybersecurity tooling. Our study finds that proactive practices are associated with higher levels of preparedness across the board (Figure 7).

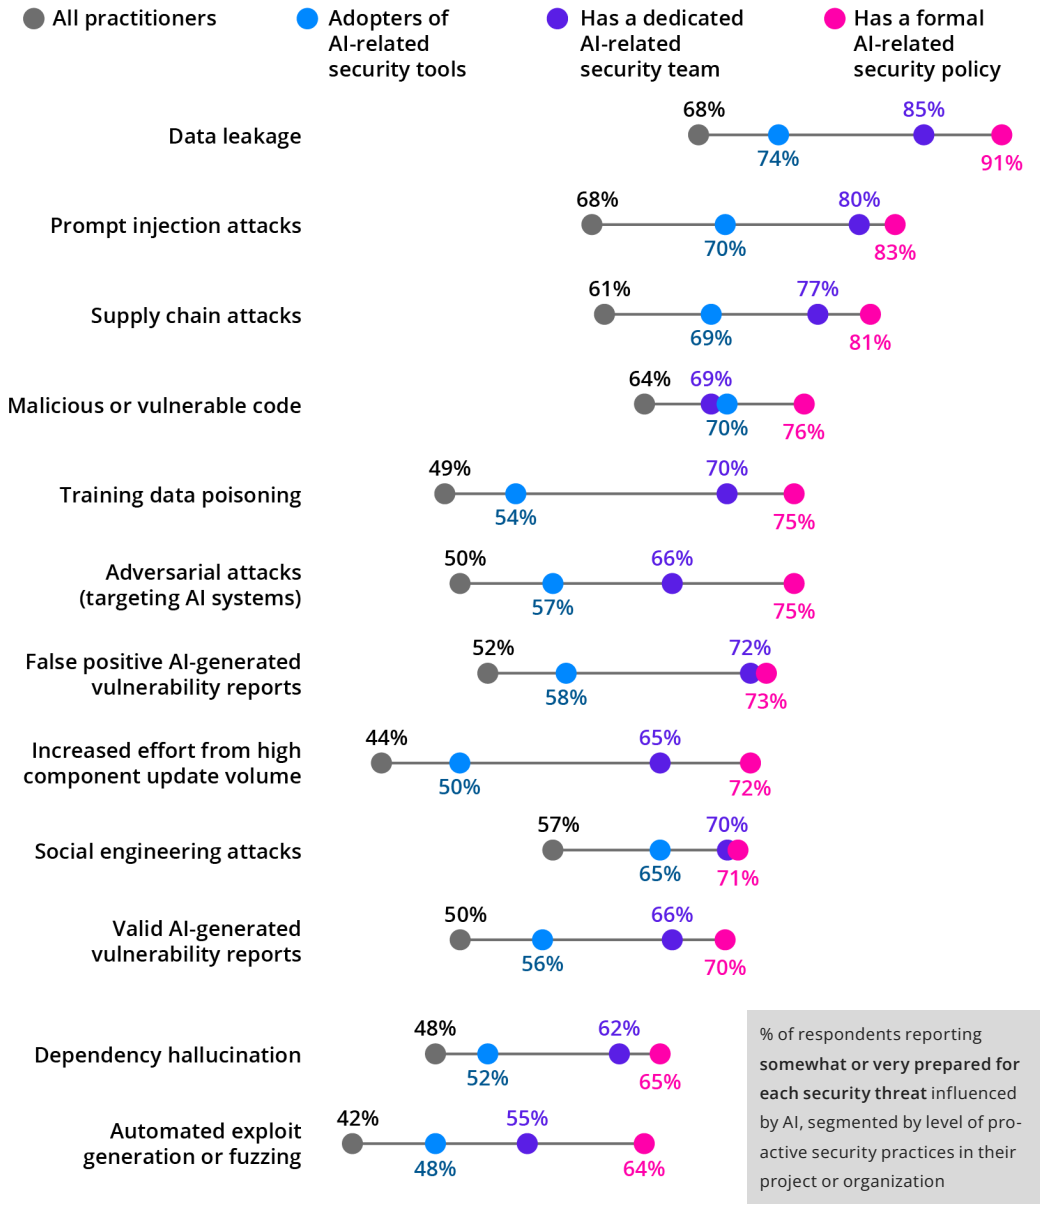
Proactive cybersecurity has a dramatic impact on an organization's overall preparedness. For example, 75% of organizations or projects that have a dedicated cybersecurity role for AI-related threats report being well prepared for adversarial attacks targeting AI systems compared with the ecosystem average of 50%.

Action items

- **Formal security policies addressing AI and well implemented tooling are low hanging fruit**. Proactive security practices have clear value. But not all organizations or projects have the resources for dedicated AI software security specialist roles. Prioritize lower cost measures to better match well prepared teams.

FIGURE 7

PROACTIVE SOFTWARE SECURITY COINCIDES WITH INCREASED PREPAREDNESS FOR EMERGING AI THREATS



% of respondents reporting **somewhat or very prepared for each security threat** influenced by AI, segmented by level of pro-active security practices in their project or organization

SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q13 SEGMENTED BY Q11, Q15, Q17. NUMBER OF RESPONDENTS = 202.

30% of teams have a dedicated role responsible for AI software security threats

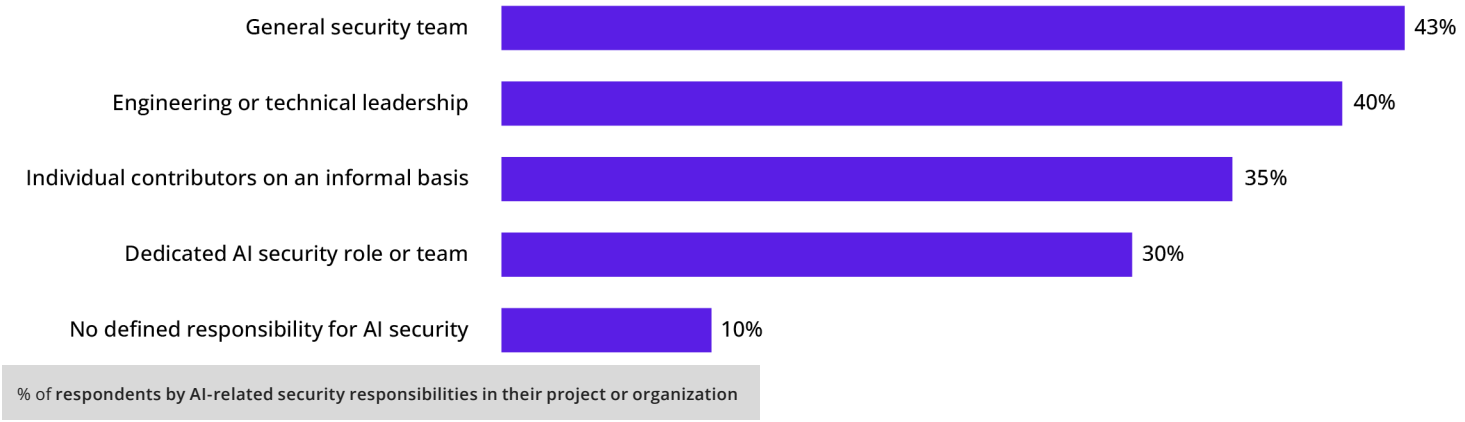
AI software security responsibilities are often distributed across existing teams rather than centralized (Figure 8). The most frequent pattern is that organizations have a dedicated security team but no dedicated role responsible for AI threats (43% of respondents). This suggests that while organizations are aware of the need for someone who knows how to apply AI to support software security, they may not yet have the resources or expertise to dedicate a specific role or team to this area. However, 30% of teams overall do have a dedicated security role or group responsible for AI threats. This rate naturally differs by the type of team. Organizations (35%) are more likely than OSS projects (16%) to have the resources for an AI cybersecurity specialist role or team.

Action item

- Organizations and OSS projects should **consider establishing a dedicated role or team** to improve security posture in the face of new AI threats. This may involve hiring or training personnel with expertise in hardened software security for AI threats, as well as providing them with the necessary resources and support to effectively manage these new risks.

FIGURE 8

WHO CURRENTLY HAS RESPONSIBILITY FOR AI SECURITY IN YOUR PROJECT OR ORGANIZATION? (SELECT ALL THAT APPLY)



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q15. NUMBER OF RESPONDENTS = 206. TOTAL RESPONSES = 366.

Only 1 in 3 organizations have a formal security-policy that addresses AI

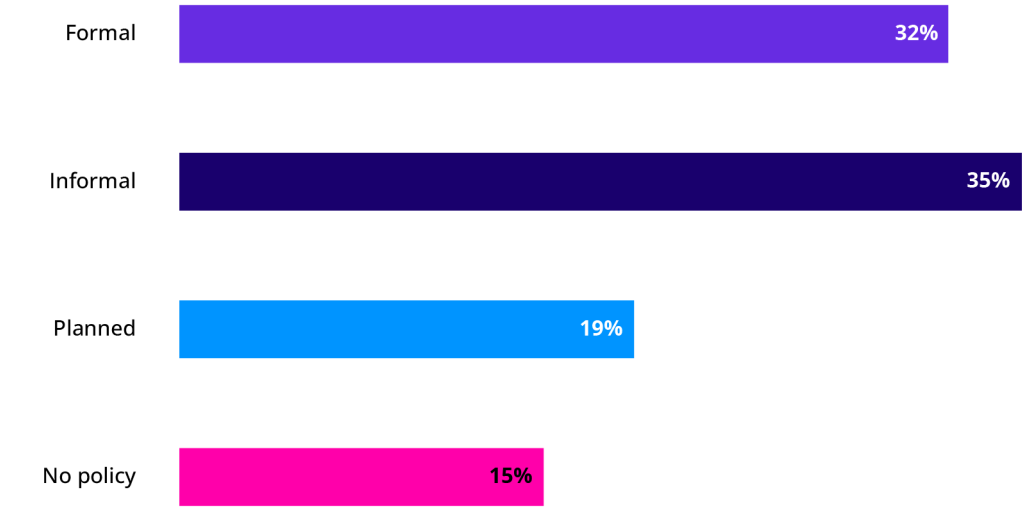
Nearly one third of organizations have a formal security policy that covers AI tool usage and threats (Figure 9). 35% of organizations have an informal security policy regarding AI while 15% do not address AI in their policy at all. Encouragingly, 19% of teams are planning to implement a formal software security policy accounting for AI in the near future, narrowing the formal policy gap.

Action item

- Organizations and projects should consider **formalizing their approach to AI insecurity policies** to ensure that they are comprehensive, clear, and enforceable.
- This may involve defining **specific guidelines for AI tool usage**, establishing **protocols for AI-generated code review**, and outlining **procedures for responding to AI-related security incidents**.

FIGURE 9

MOST ORGANIZATIONS AND OSS PROJECTS ADDRESS AI WITH SECURITY POLICY BUT THE LEVEL OF FORMALIZATION IS MIXED



% of respondents by **AI-related security policy development status**

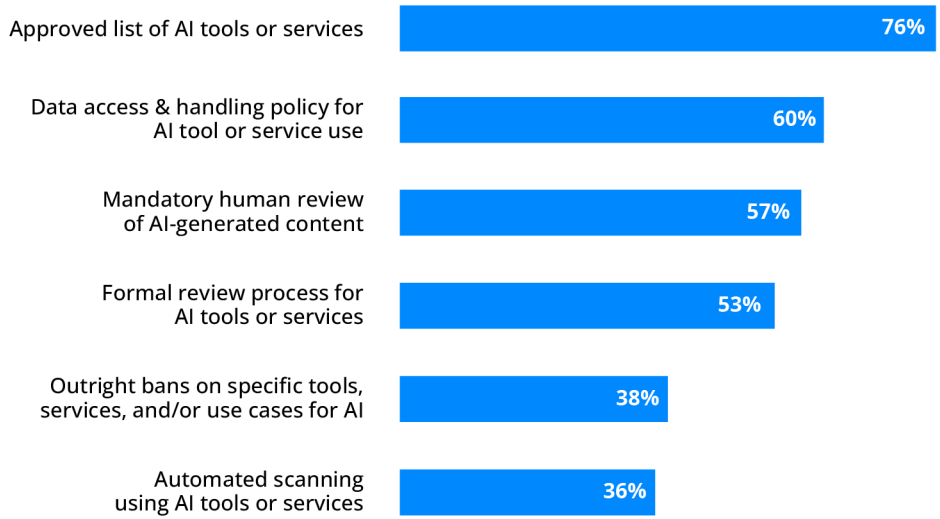
SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026.
Q17. NUMBER OF RESPONDENTS = 190.



A first order priority in the security policy is a list of approved AI tools or services (76%) (Figure 10). This is an excellent place to start but there several important elements are less common across security policies. **Only 57% of policies mandate human review of AI-generated code**. Since the human-in-the-loop paradigm reinforces the complementary nature of AI tools and leverages the best of both worlds, mandatory human review should be a feature of more policies. **Only 36% of policies cover automated scanning using AI tools or services**. Automated scanning is an important tool for teams dealing with large codebases and high throughput development. Policy coverage should match any expansion in automated vulnerability detection uptake.

FIGURE 10

WHICH OF THE FOLLOWING ARE ELEMENTS OF YOUR PROJECT'S OR ORGANIZATION'S AI POLICY? (SELECT ALL THAT APPLY)



% of respondents by AI-related security policy elements in their project or organization

SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026.
Q18. NUMBER OF RESPONDENTS = 156. TOTAL RESPONSES = 516.

Action item

- AI-generated code for mission critical production systems should be prioritized for human review. Skipping methodical review by an expert creates security debt that can range from mildly costly to catastrophic depending on the business importance application. Start a conversation about balancing review responsibility with production tempo.
- Require automated scanning for vulnerability detection for all code, including AI-generated code. Vulnerability detection is an increasingly important tool for security practitioners.

Impact of AI on development and security workflows

How are AI tools impacting development and security workflows within organizations? This section explores the effects of AI on code shipping speed, security alert handling, and overall development practices. In particular, our study finds concerning patterns that ultimately lead to the accumulation of security debt over time.

3 in 5 teams ship code faster than security review can keep up

A majority (58%) of organizations are shipping code faster than security review can keep up, which can lead to a backlog of vulnerabilities that are not being addressed in a timely manner (Figure 11). This suggests a **failure to integrate adequate security review that matches production tempo.**

Action item

Properly implemented AI-powered software security tools can be a helpful part of maximizing review throughput. **Optimize the security review process, increase security resources** (if possible), and **implement automated security checks** to keep up with faster development cycles.

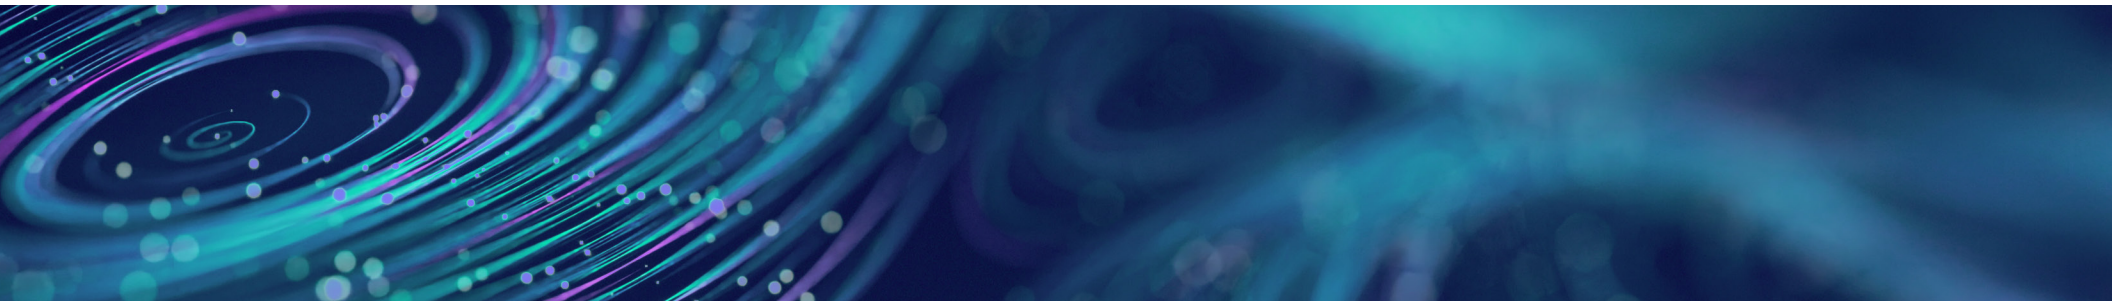
FIGURE 11

Security review lags behind production timelines for a majority of teams

61% of projects or organizations with an ***informal*** AI software security or cybersecurity policy agree they are **shipping code faster than security review can keep up**

52% of projects or organizations with a ***formal*** AI software security or cybersecurity policy agree they are **shipping code faster than security review can keep up**

SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q25 X Q17. NUMBER OF RESPONDENTS = 173.

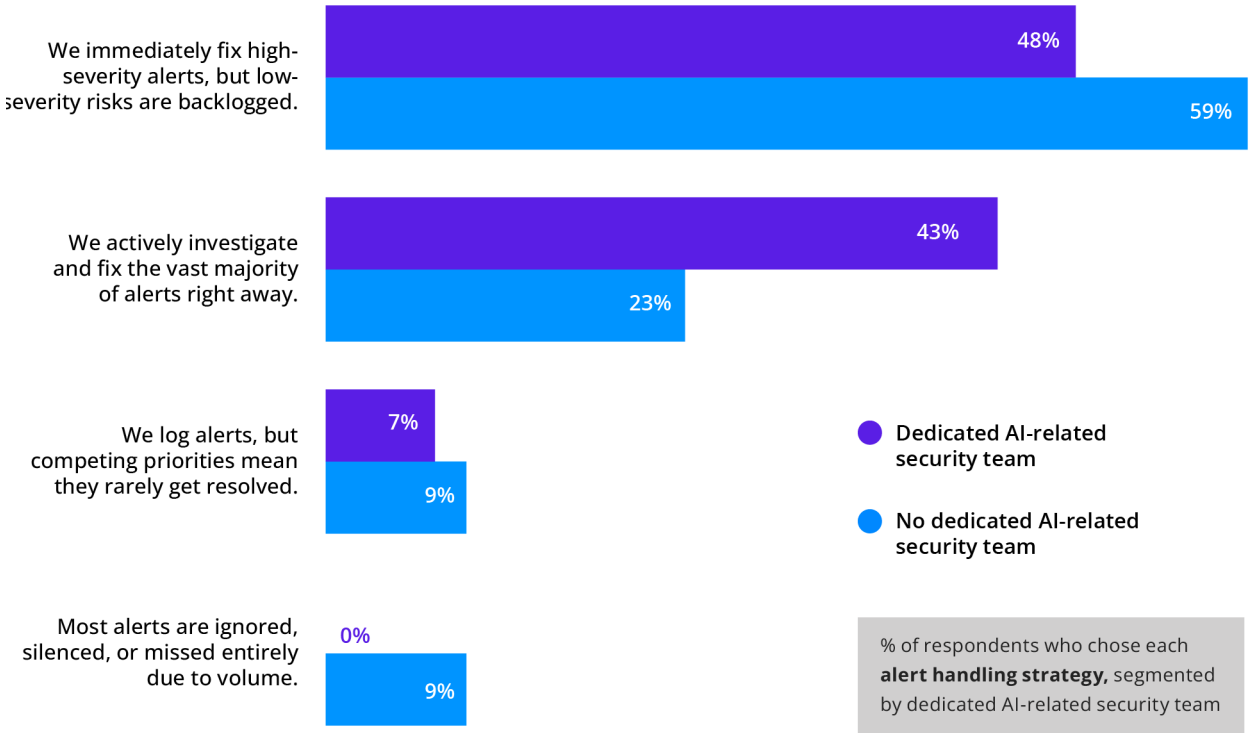


Most teams triage alerts rather than fixing them immediately

Security response motions vary depending on organization resources and structure. Organizations with dedicated security teams with AI threat expertise and responsibilities tend to handle alerts more proactively, while those without such teams may experience delays or backlogs. **Only 23% of organizations or projects without a dedicated security team or role for AI threats actively investigate and fix the vast majority of new security alerts they receive immediately** (Figure 12). When a **dedicated security team for AI threats is present,** this share **increases to 43%**. Even when the backlogged vulnerability alerts are low severity, limited availability of security resources has the potential to create a growing low-severity backlog and even an accumulation of security debt.

FIGURE 12

DEDICATED TEAMS FOR AI THREATS HELP PREVENT SECURITY DEBT BEFORE IT CAN ACCUMULATE

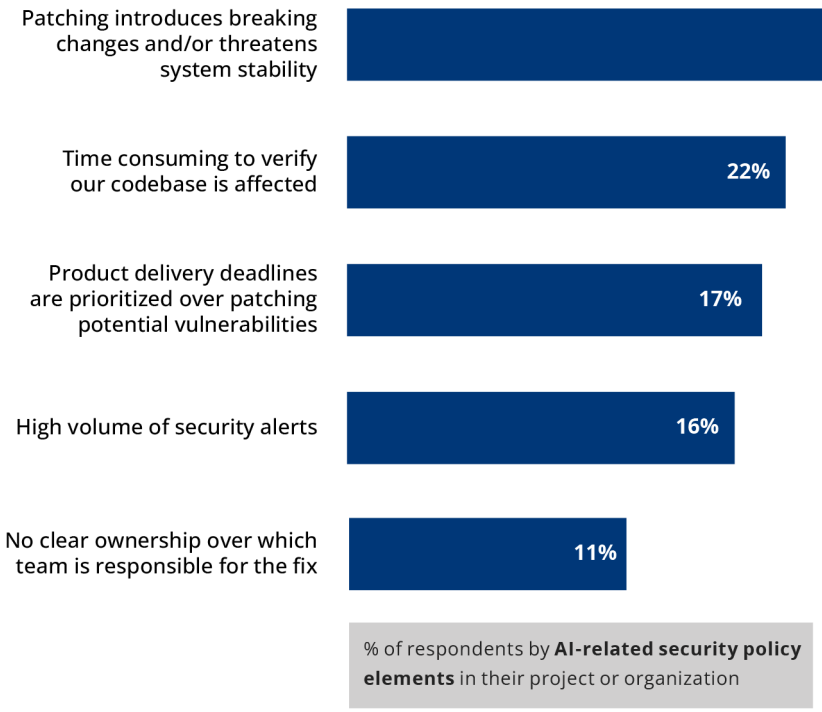


SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q26 SEGMENTED BY Q15.

Security resource constraints are just one factor driving the triage-first approach to vulnerability management. When vulnerabilities are deprioritized, the main reasons cited are potential breaking changes (28%) and time constraints (22%) (Figure 13). **Production stability concerns often play a significant role in these decisions**, as organizations aim to minimize disruptions to their live systems.

FIGURE 13

PRODUCTION PRIORITIES AND TIME CONSTRAINTS ARE COMMON REASONS FOR DEPRIORITIZED VULNERABILITY FIXES



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q27. NUMBER OF RESPONDENTS = 181.

Action item

- Organizations and projects should consider implementing a more structured approach to vulnerability management. While triage is an efficient and necessary process, this creates scope for a growing low-severity backlog.
- Establish clear criteria for severity assessment, prioritize vulnerabilities based on potential impact, and ensure prompt remediation for critical vulnerabilities.
- Eliminate vulnerability potential closest to its source. Whenever possible, focus on upstream causes rather than downstream symptoms. Sourcing dependencies from a vetted, policy-governed repository keeps unvetted and known-vulnerable components out before they ever need triage.
- Organizations should harden the systems they develop to reduce the number of vulnerabilities that need triaging and prioritizing in the first place. Fixing entire classes of problems, at the source, eliminates the need to triage those problems later.

Practitioners look to blogs, newsletters, and open source foundations to stay informed on security developments related to AI

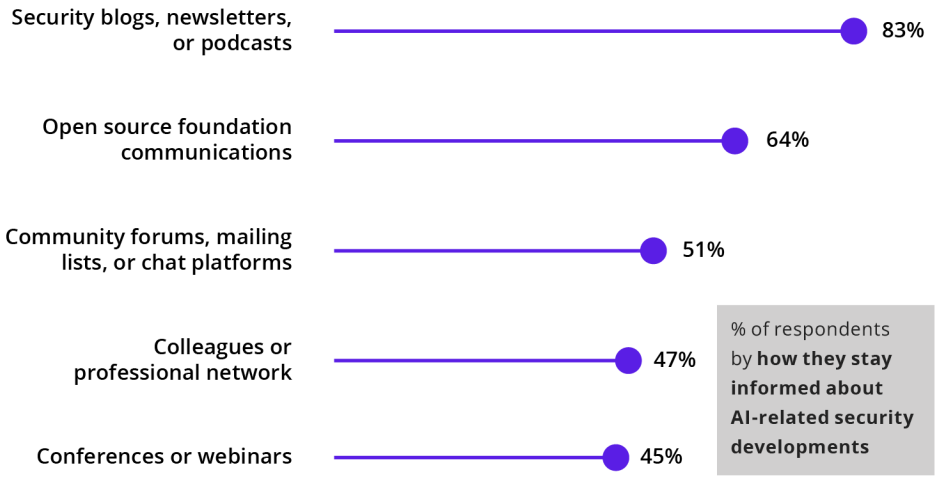
The rapid pace of AI development makes it challenging for practitioners to stay informed on emerging threats and cybersecurity best practices. To **keep up with security developments and news related to AI threats**, practitioners most often **rely on blogs, newsletters, and podcasts (83%) and communications by open source foundations (64%)** (Figure 14). Practitioners clearly value and trust the guidance of open source foundations, which are well positioned to provide authoritative guidance on security best practices for the AI era. At the same time, the catch-all nature of popular media meets practitioner audiences where they are. management practices, including regular audits, automated dependency scanning, and strict review processes for new dependencies.

Action item

Open source foundations should **continue to provide guidance on security best practices for AI-related concerns** and coordinate efforts to better reach busy practitioners. This may involve **expanding into more popular media formats** like podcasts.

FIGURE 14

4 out of 5 practitioners get security guidance on AI from blogs, newsletters, and podcasts



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q16. NUMBER OF RESPONDENTS = 209. TOTAL RESPONSES = 843.

Barriers and investment priorities

As the community navigates rising AI usage and its effects on the security landscape, it's helpful to directly ask practitioners what they consider the key obstacles to implementing better software security practices and what areas warrant additional investment to make operations more robust to AI-driven vulnerabilities.

Figure 15 reports the most cited barriers and investment priorities for improving software security and cybersecurity with respect to AI-induced risks. The rapid pace of AI change is the most commonly cited barrier to adopting better security practices for AI usage (61%), ahead of lack of budget or resources (54%) and lack of clear guidance to deal with AI-specific threats (49%). While the rapid pace of change is beyond the practitioner's control, several cited barriers to improve AI cybersecurity are candidates for community attention. Clear and practical guidance coming from trusted voices in the community would go a long way in increasing awareness and security readiness for AI.

Practitioners also clearly indicate what current gaps in security that warrant additional investment. The top priorities are improved tooling (63%), AI supply chain security[4] (56%), access to effective local LLMs (49%), and better education on AI-related security threats and remediation (47%) (Figure 16). Much like the most commonly cited barriers, respondents to the study highlight key areas where community investments can improve overall security posture across organizations and OSS projects. A key interest from the community is a quality open source scanner specifically tailored to vulnerability detection in AI-generated code. These tools would create cascading spillover benefits across the community. Key community leaders in the security and open source spaces would be wise to spearhead these efforts.

4 In particular, respondents are interested in AI system provenance, which include security vetting for pre-trained models and datasets.

FIGURE 15

TOP BARRIERS TO IMPROVING SOFTWARE SECURITY FOR AI THREATS

RANK	BARRIER	% OF RESPONDENTS
1	Rapid pace of AI change	61%
2	Limited budget or dedicated resources	54%
3	Lack of clear, practical guidance for AI-specific threats	49%
4	Lack of awareness or education of AI systems risks	43%
5	Competing priorities with other work	43%

SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026.
Q21. NUMBER OF RESPONDENTS = 194. TOTAL RESPONSES = 613.

FIGURE 16

TOP SECURITY INVESTMENT PRIORITIES FOR AI RISKS

PRIORITY	INVESTMENT	% OF RESPONDENTS
1	Tooling (e.g., open source scanners for AI-generated code)	63%
2	AI supply chain provenance and vetting	56%
3	Access to effective local LLM/SLMs	49%
4	Education on threats & remediation	47%
5	Standards and threat modeling frameworks	46%

SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026.
Q22. NUMBER OF RESPONDENTS = 196. TOTAL RESPONSES = 800.



Action items

- High quality open source vulnerability detection tools for AI-generated code are a clear priority for the community. Open source foundations and organizations should consider investing in the development and maintenance of such tools, as well as providing guidance on how to effectively use them.
- Supply chain security demands focus. The supply chain faces even greater challenges due to widespread AI adoption: rapidly generated code prone to vulnerabilities, increasingly sophisticated malicious actors, and new vectors for exploits or leaks. Efforts like Akrites and OpenSSF's supply chain security initiatives are critical to the community's ability to address these challenges.
- Much like open source security tooling, effective local LLMs and SLMs can dramatically increase average cybersecurity posture for AI. Practitioners need effective AI tools to efficiently counteract AI software security threats at production scale.
- Although not directly addressed by our survey instrument, vetted open source AI models, not just open weight, are of great interest to the community. Users seek models[5] they can use locally with provenance assurances, attributes that cloud providers and even open weight models cannot guarantee.
- Education and training on AI-related security threats and remediation are essential to equip practitioners with the necessary skills and knowledge. Organizations should invest in developing comprehensive training programs, workshops, and resources to ensure their teams stay up-to-date with the latest AI-related security practices.

5 **https://opensource.org/ai/open-source-ai-definition**

Opportunities for AI-assisted security tools

A key finding of this study is that AI tools can also be part of the solution to the new challenges they create. While they introduce new risks, they also offer significant potential to enhance security practices and equip security practitioners with powerful capabilities. We next explore how AI-assisted security tools influence security posture overall, how they are being used to help secure code, and what limitations arise in practice.

7 out of 10 practitioners consider AI-assisted security tools a net positive complement to existing practices

Among cybersecurity practitioners that actively use AI-assisted tools, **69% agree that these tools meaningfully improve overall security posture within their organization or project** (Figure 17). Only 10% believe that security tooling leveraging AI actually degrades cybersecurity. But it should be noted that AI tools are not a replacement for existing security practices and workflows. **72% of respondents agree that AI-assisted security tools complement their existing security practices and workflows rather than replace or disrupt them.**

FIGURE 17

PRACTITIONERS FIND VALUE IN AI-ASSISTED SECURITY TOOLS AS COMPLEMENTS TO EXISTING PRACTICE

69% of practitioners agree that **AI-assisted security tools have meaningfully improved overall security posture** for their project or organization

62% of practitioners agree that **AI-assisted security tools complement existing security practices and workflows**

SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q32 (NUMBER OF RESPONDENTS = 115). Q33 (NUMBER OF RESPONDENTS = 112).

Examples of AI-assisted tools for software security

AI-assisted software security tools streamline developer workflows by automating vulnerability detection, generating code fixes, and securing novel generative AI deployments. In code analysis, AI-enhanced static testing platforms like **Snyk Code** and **GitHub Advanced Security** detect flaws in real time and auto-generate code fixes, while deterministic static analysis platforms like **Semgrep** (via Semgrep Assistant) and **CodeQL** (via GitHub Copilot Autofix) leverage AI to triage findings and draft patch recommendations. To protect third-party components, supply chain tools like **Dependabot** automatically open pull requests to resolve vulnerable open source dependencies. Security teams are also deploying autonomous agent ensembles, often combining agentic frameworks like **AutoGen** and **LangChain** with frontier LLMs, to perform dynamic agentic fuzzing, reason through complex application logic, and autonomously generate proof-of-concept exploits.

Action item

- The majority finds AI-assisted security to be a net positive within their organization or project. Practitioners, especially holdouts, should **strongly consider exploring integrating AI tools into their existing workflows.**
- However, it is important to **carefully evaluate the specific tools and their capabilities to ensure that they align with the organization's security needs and objectives**. Clear, formal policies can go a long way in ensuring these tools are being used properly. Additionally, organizations should provide training and support to their teams to effectively integrate AI-assisted security tools into their existing workflows.

Vulnerability detection is the top security use case for AI

Where exactly can AI-assisted security tools be most useful? There are a number of good applications and important use cases for these tools, but next explore which ones are most commonly used in practice within organizations and OSS projects.

Our study finds that using AI for vulnerability detection has become a well established cybersecurity application. Vulnerability detection is the top recommended security use case for AI tools by cybersecurity experts (74% of experts agree). After AI coding generation and completion tools, AI-assisted vulnerability detection is the most widely deployed application of AI for software security tools across organizations and OSS projects (68% of practitioners use it). Practitioners report that 30% of vulnerabilities were successfully detected using AI-assisted security tools within the last 12 months. Vulnerability detection is a clear standout use case, greatly augmenting practitioners ability to review and identify potential security issues in continuously evolving and expanding codebases. The next step is vulnerability remediation at scale, which allows developers to operationalize detection and efficiently introduce fixes.

FIGURE 18

AI-ASSISTED VULNERABILITY DETECTION IS RECOMMENDED CYBERSECURITY EXPERTS AND WIDELY USED IN PRACTICE

74% of security practitioners that consider **vulnerability detection a good use case for AI**

68% of projects and organizations that **currently use AI for vulnerability detection**

30% of security **vulnerabilities were successfully detected using AI assistance** in the last 12 months

SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q30 (NUMBER OF RESPONDENTS = 109). Q11 (NUMBER OF RESPONDENTS = 207, TOTAL RESPONSES = 747). Q31 (NUMBER OF RESPONDENTS = 146).

Action item

Vulnerability detection using AI tools is a top priority for teams considering how to improve their cybersecurity posture. Not only is it recommended by experts and widely used, but it's proven to be successful in practice.w

AI productivity gains under threat from false positive vulnerability reports

AI tools let practitioners complete security related development work 17% faster on average (Figure 19). AI tools confer a significant productivity gain for security developers, considering the importance of timely vulnerability remediation, the opportunity cost of the practitioner's time, and the increased pace of production overall. However, the increased incidence of false positive vulnerability reports from AI-assisted security tools threatens to erode these productivity gains. On average, practitioners report a 30% increase in workload from false positive vulnerability reports. The cost to generate a vulnerability report plummets with AI, but extracting correct reports, never mind the signal from noise, becomes the new challenge. While it's encouraging to see that AI-powered vulnerability detection is being widely accepted by security experts and practitioners.

FIGURE 19

PRODUCTIVITY GAINS AI-ASSISTED SECURITY TOOLS ARE MIXED

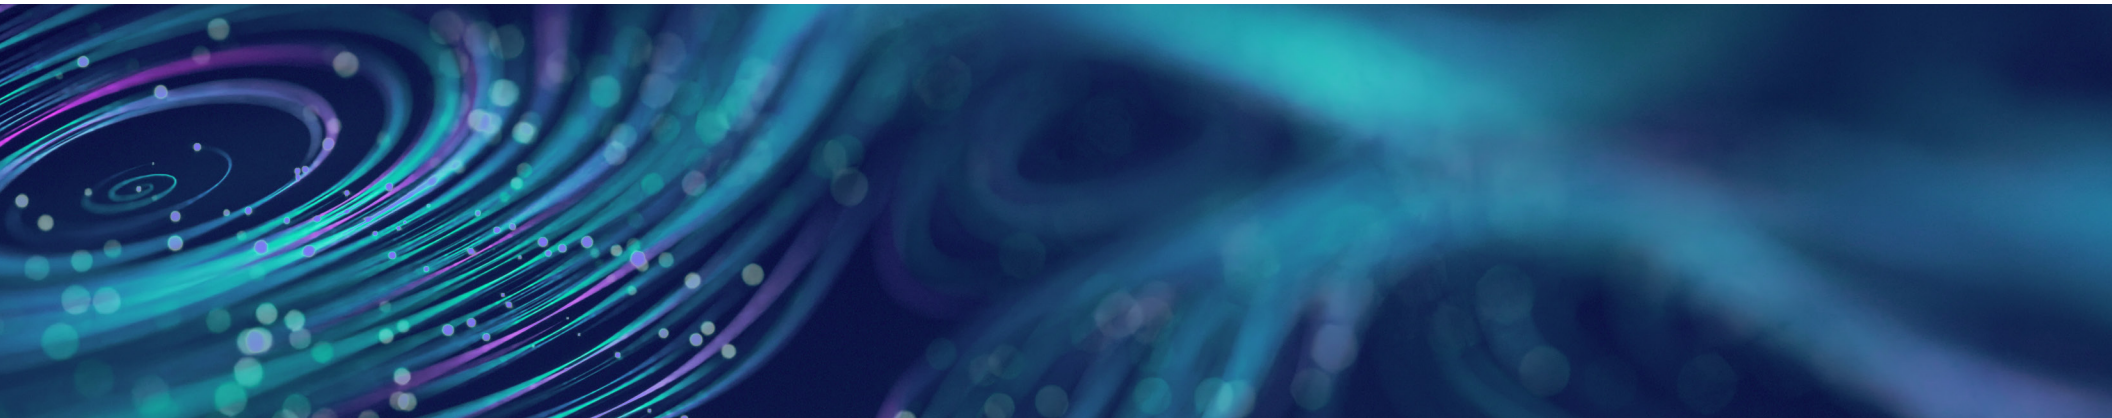
+17% **increase in security development productivity** due to AI tool use

+30% **increase in workload** due to **false positive vulnerability reports** from AI vulnerability detection tools

SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q30 (NUMBER OF RESPONDENTS = 109). Q11 (NUMBER OF RESPONDENTS = 207, TOTAL RESPONSES = 747). Q31 (NUMBER OF RESPONDENTS = 146).

Action item

- Organizations and OSS projects should equip their development and security teams with AI-assisted security tools to help them work more efficiently and effectively.
- However, the productivity gains from AI-assisted security tools should be weighed against the additional workload caused by false positives. While we find evidence that effects might offset on average, a concerning imbalance between underpowered productivity gains and overwhelming false positive security reports is a likely scenario for some organizations and projects.



References

Cybersecurity and Infrastructure Security Agency (CISA) **"Cybersecurity Collaboration Playbook"** Operational guidance published by CISA outlining best practices, protocols, and workflows for joint cyber defense and multi-organization collaboration.

Elastic **"What Is AI Cybersecurity?"** An overview by Elastic detailing the role of artificial intelligence in threat detection, automated response, and modern cybersecurity operations.

European Commission **"EU Action Plan on Cybersecurity and Artificial Intelligence"** An official European Commission policy document outlining a coordinated strategy to promote the safe deployment of AI models, strengthen EU cyber resilience, and scale up AI capabilities for proactive threat detection and defense.

National Institute of Standards and Technology (NIST) **"NIST Special Publication 1353 (Initial Public Draft)"** Official draft guidance from NIST on managing cybersecurity risk, technical standards, and security frameworks.

Open Source Security Foundation **"Security-Focused Guide for AI Code Assistant Instructions"** A developer guide curated by the OpenSSF providing security principles, prompt engineering guidelines, and safety practices for deploying AI-powered coding assistants.





Conclusion

AI is now part of the software development and security processes, and its effects are consequential in both directions. The expansion of AI tools and AI-generated code is increasing code volume, dependency complexity, attack surface, and the speed at which vulnerabilities can be introduced or exploited. At the same time, AI-assisted security tools are already helping practitioners detect vulnerabilities, review code, manage dependencies, and complete security-related development work more quickly.

The central challenge is not whether to use AI, but whether security practices are developing quickly enough alongside its use. Practitioners report greater concern than preparedness for every AI threat to software security measured in this study. Supply chain attacks, prompt injection, and maintenance demands are especially prominent concerns, while elevated component update volume and automated exploit generation or fuzzing show some of the largest preparation gaps. The findings describe a forward-looking risk environment: concern is widespread even where reported incident experience remains lower.

The evidence also points to organizational readiness as a practical lever. Teams with dedicated software security and cybersecurity roles for threats introduced by AI, formal policies, and properly implemented tooling report stronger preparedness across threat categories. This is an association rather than proof that any single practice causes better outcomes, but it is consistent across the study and provides a concrete direction for action. The limited reach of dedicated roles and formal policies, combined with the fact that many teams ship faster than security review can keep up, shows why ownership and process cannot be left implicit.

AI-assisted security tools offer meaningful benefits when integrated with human judgment and established controls. Among active users, most report improved security posture and regard AI as a complement rather than a replacement for existing practices. However, the average productivity gain is accompanied by a larger reported increase in workload from false-positive vulnerability reports. Tool adoption should therefore be evaluated by the quality of the complete workflow, including validation, prioritization, remediation, and feedback, rather than by detection volume or speed alone. The perceived risk that overreliance weakens developers' security skills also reinforces the need to preserve human expertise through review and training.

The path forward is disciplined adoption supported by investment. Organizations and OSS projects can reduce exposure by clarifying responsibility, formalizing expectations, strengthening dependency and supply chain practices, and using AI tools where they provide measurable value. Open source foundations and other ecosystem institutions have an important role in providing practical guidance, accessible tooling, education, and support for teams that cannot sustain specialized security capacity on their own.



Policy recommendations

1. **Leverage AI tools for security tasks.** Use AI-assisted tools to enhance vulnerability detection, code review, dependency management, and other security-related activities. Ensure that AI outputs are reviewed and validated by human experts, when it's important that the outputs correct, to maintain accuracy and accountability.

2. **Assign clear security responsibilities for threats introduced by AI.** Organizations should designate an accountable owner or team for AI-related security threats, while OSS projects should define AI threat responsibility within their maintainer group, security policy, or governance model. The structure can scale to match the project's size and resources, but responsibility should never be left undefined.

3. **Formalize security policies regarding AI.** Policies should identify approved AI tools and services, and approved sources for open source components, set expectations for human review of AI-generated code, require appropriate automated scanning, address model and dataset provenance, and define procedures for responding to AI-related incidents.

4. **Strengthen dependency and supply chain controls.** Source direct and transitive dependencies from vetted, policy-governed repositories, require review of new dependencies, maintain inventories and provenance records, scan generated code and dependencies, and establish a timely process for evaluating updates and remediation. These controls are especially important as AI coding tools increase the volume and complexity of dependency graphs.

5. **Leverage AI tools to expand security bandwidth rather than replace the practitioner.** Like any other tool, treat AI-assisted security as a complement to established practices. Preserve reviewer accountability, provide continuing security training, and monitor whether tool use is reducing opportunities for developers and maintainers to exercise and develop security skills.

6. **Manage false positive vulnerability reports as an operational risk.** Evaluate tools using precision, remediation outcomes, time saved, and workload added. Establish triage and feedback loops so that false positive reports are validated, prioritized, and used to improve configurations and tools rather than becoming unmanaged security debt.

7. **Invest in accessible ecosystem capabilities.** Prioritize high-quality open source security tooling, robust AI supply chain infrastructure, effective local LLMs and SLMs, practical threat guidance, and education. Shared resources can help smaller organizations and OSS projects address risks that would otherwise require dedicated specialists.

Methodology

Data collection method	Web survey
Survey target population	Information technology professionals, software developers, and open source software contributors who are involved with cybersecurity
Survey distribution	Linux Foundation newsletter subscribers, Linux Foundation social media channels, research partner email/social channels (OpenSSF, ActiveState, LF AI & Data, IoTM2M)
Survey instrument development period	April - May 2026
Survey fielding period	June - July 2026
Analytical sample size	308 qualified respondents
Margin of error (95% confidence level)	+/- 7%
Data access	**https://zenodo.org/communities/linux-foundation-research**
Declaration of AI in the research process	LF Research uses generative AI and AI-assisted technologies in varying degrees for research-related tasks including aggregating secondary sources, synthesizing data, supporting report drafting, copyediting, and translation. All AI-generated content is reviewed and revised as needed.

Survey demographics

Respondent Demographics

The survey captures a global perspective, drawing from a geographically diverse population of respondents. The largest single concentration of participants resides in the United States (27%), followed closely by representation across the European Union (18%) and India (16%). The remaining 39% of respondents are distributed across various other regions worldwide, offering a well-rounded international viewpoint (Figure 29).

This global sample is matched by a rich mix of technical expertise and professional backgrounds. Software developers and engineers constitute the largest professional segment at 31%, with DevOps and platform engineers following at 30%. Additionally, OSS contributors and engineering or technical leaders each represent 28% of the sample (Figure 20). In terms of employment status, the vast majority of participants are actively engaged in the workforce. Fully 67% of respondents are employed full-time, while an additional 12% operate as self-employed professionals (Figure 21). Together, these figures reflect a highly technical audience with direct hands-on and managerial involvement in technology ecosystems.

Building on this solid professional base, respondents demonstrate significant grounding in cybersecurity, largely developed through practical and academic avenues. A striking 74% report that their cybersecurity knowledge stems directly from organizational work, complemented by 57% who cite formal training, education, or academic study. Furthermore, 29% draw security expertise from their contributions to open-source software projects (Figure 22). When looking specifically at the cybersecurity subsample, nearly three-quarters (73%) frame their answers through the lens of organizational employment, whereas 20% respond within the context of an OSS project (Figure 23). Deep domain knowledge is further reinforced by the sample's extensive industry tenure. An impressive 41% of respondents bring more than 20 years of IT experience to the table, while another 29% boast between 11 and 20 years in the field (Figure 24). Consequently, the overwhelming majority of participants possess a high baseline of experience, making them uniquely qualified to evaluate emerging security challenges.

Respondents exhibit a robust understanding of modern AI risks. Combined, 98% assess their familiarity with AI threats as notable: 42% describe themselves as familiar, 39% as proficient, and 17% as advanced, leaving a mere 2% with no familiarity at all (Figure 25). Reflecting the real-world priority of these issues, participants report spending an average of 27% of their working time dealing with tasks related to AI threats (Figure 26). Overall, these metrics establish a respondent base that is not only deeply experienced, but actively engaged in addressing contemporary security challenges.

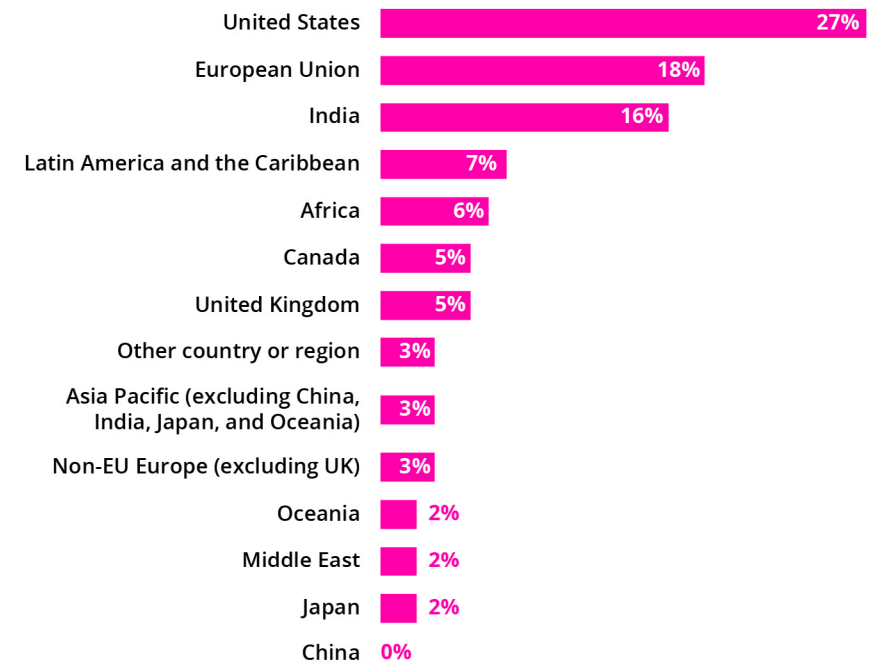
Organization Demographics

The organizations that survey respondents work for span a wide array of economic sectors. The sample is naturally heavily represented by the cross-industry information technology sector, which accounts for 37% of organizations. Financial services (11%) and telecommunications or internet services (9%) round out the next largest enterprise segments (Figure 27). This distribution underscores a strong representation from core digital infrastructure and high-security sectors.

The sample covers organizations of all scales, from small startups to global enterprises. Interestingly, the distribution of organization sizes (measured by number of employees) is approximately bimodal: 22% of respondents work at small organizations with 1 to 49 employees, while another 22% belong to massive enterprises employing 20,000 or more people (Figure 28). Ultimately, this balanced representation ensures that the survey insights reflect perspectives from both agile, small-scale environments and large enterprise infrastructures.

FIGURE 20

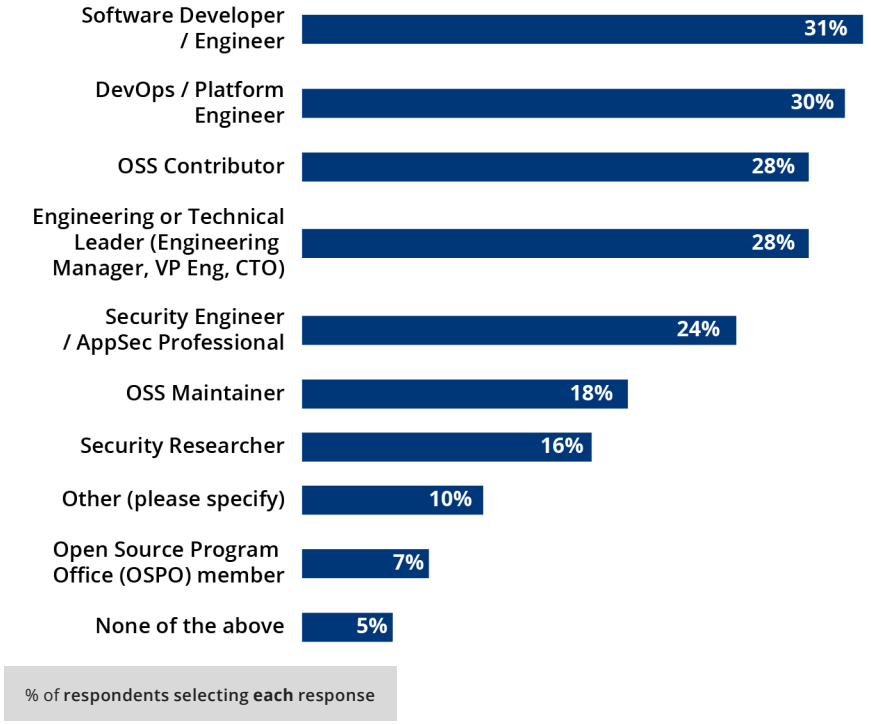
IN WHICH COUNTRY OR REGION ARE YOU CURRENTLY LOCATED? (SELECT ONE)



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026.
Q8. NUMBER OF RESPONDENTS = 266.

FIGURE 21

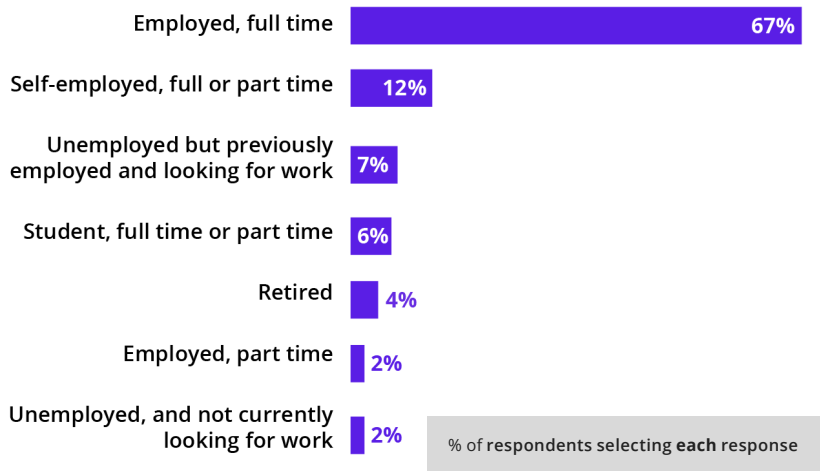
WHICH OF THE FOLLOWING DESCRIBE YOUR ROLE IN THE OSS ECOSYSTEM? (SELECT ALL THAT APPLY)



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026.
Q1. NUMBER OF RESPONDENTS = 308. TOTAL RESPONSES = 609.

FIGURE 22

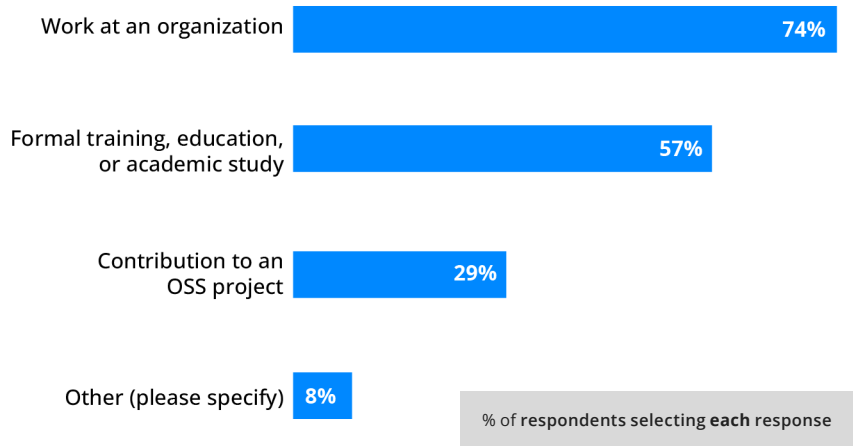
WHAT IS YOUR CURRENT EMPLOYMENT STATUS? (SELECT ONE)



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q3. NUMBER OF RESPONDENTS = 308.

FIGURE 23

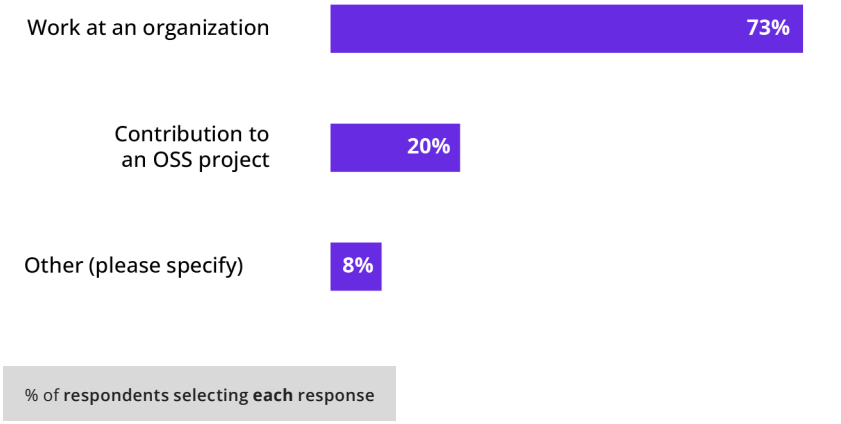
WHAT IS THE BASIS OF YOUR CYBERSECURITY EXPERTISE? (SELECT ALL THAT APPLY)



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q4. NUMBER OF RESPONDENTS = 266. TOTAL RESPONSES = 445.

FIGURE 24

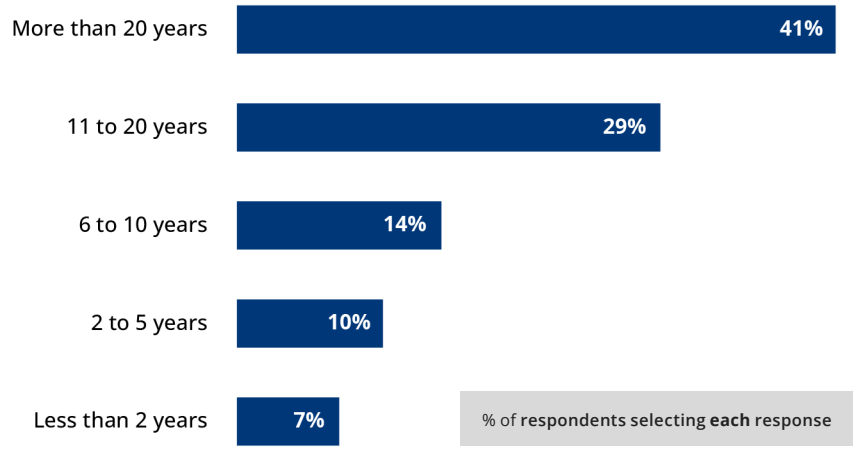
WHICH CYBERSECURITY WORK CONTEXT ARE YOU USING TO ANSWER THIS SURVEY? (SELECT ONE)



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q5. NUMBER OF RESPONDENTS = 266.

FIGURE 25

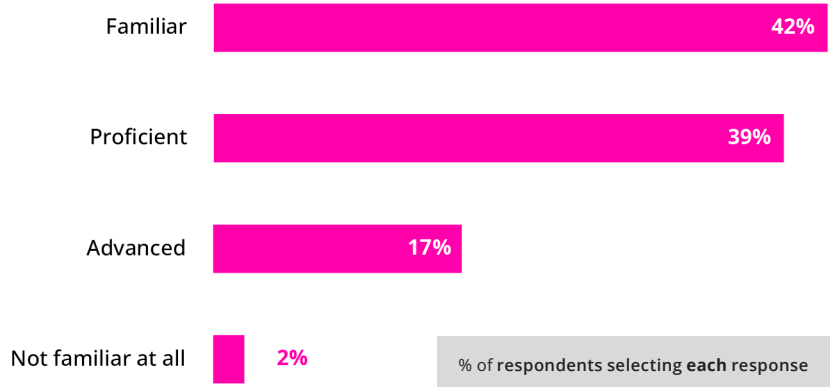
HOW MANY YEARS OF IT EXPERIENCE DO YOU HAVE? (SELECT ONE)



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q6. NUMBER OF RESPONDENTS = 266.

FIGURE 26

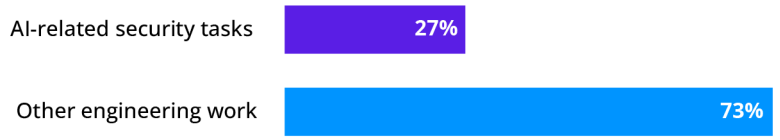
HOW WOULD YOU CHARACTERIZE YOUR LEVEL OF UNDERSTANDING OF AI SECURITY RISKS AND BEST PRACTICES IN THE CONTEXT OF OSS DEVELOPMENT? (SELECT ONE)



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q7. NUMBER OF RESPONDENTS = 266.

FIGURE 27

AVERAGE SHARE OF ENGINEERING TIME SPENT ON AI SECURITY TASKS

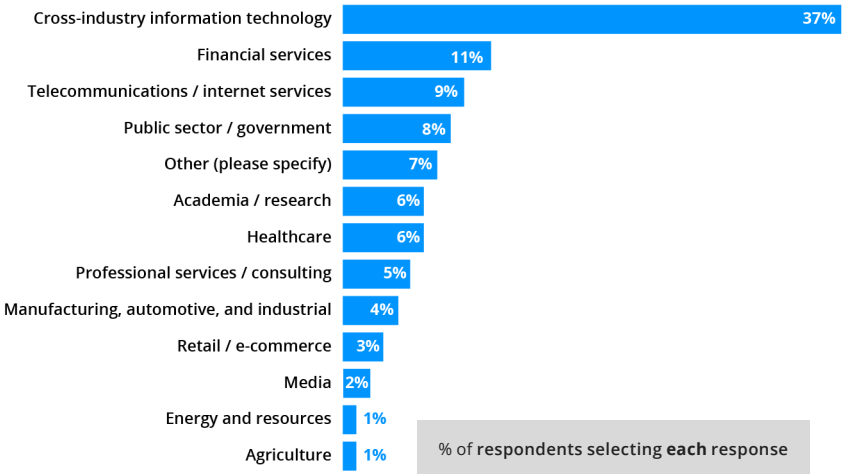


Average **share of engineering time** spent on **dedicated AI-related security tasks** reported by survey respondents.

SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q24. NUMBER OF RESPONDENTS = 174.

FIGURE 28

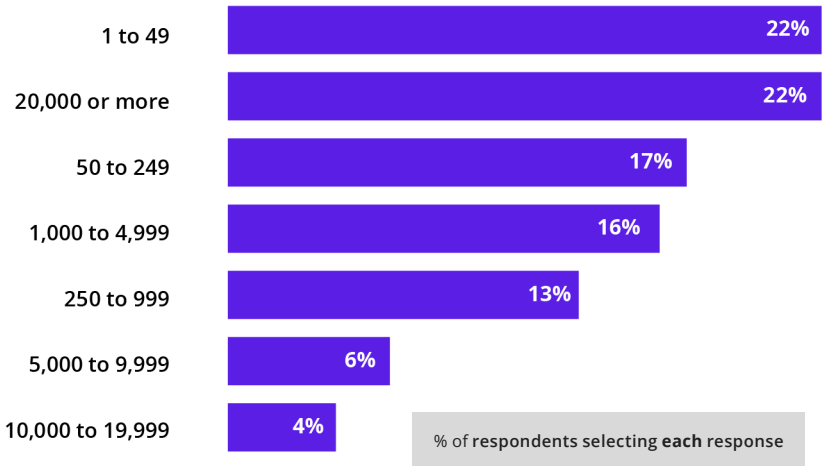
WHICH RESPONSE BEST DESCRIBES YOUR ORGANIZATION'S PRIMARY INDUSTRY? (SELECT ONE)



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q9. NUMBER OF RESPONDENTS = 190.

FIGURE 29

PLEASE ESTIMATE HOW MANY TOTAL EMPLOYEES ARE IN THE ORGANIZATION YOU WORK FOR. (SELECT ONE)



SOURCE: STATE OF AI IMPACT ON SOFTWARE SECURITY AND CYBERSECURITY 2026. Q10. NUMBER OF RESPONDENTS = 186.



About the author

Sam Boysel is a Data Scientist at the Linux Foundation. In this role, he contributes to a variety of empirical initiatives to uncover new insights into open source communities. His work often leverages microeconomic theory to analyze behaviors, incentives, and community dynamics across open source ecosystems. Before joining the Linux Foundation, he was a postdoctoral researcher with the Laboratory for Innovation Science at Harvard. He holds a B.A., M.S., and Ph.D. in Economics.

Acknowledgments

This study has been a collaboration between several different partners and individual contributors. In particular, the author would like to thank OpenSSF, ActiveState, and LF AI & Data for sponsoring this work. A special thanks to Rebecca Banks, Hilary Carter, Laura Guazzelli, Anna Hermansen, Adrienn Lawson, Adrianne Marcum, Christopher "CRob" Robinson, and David A. Wheeler for their support in survey development, review of findings, and report production.

LF Research is grateful to the organizations sponsoring this study:

OpenSSF OPEN SOURCE SECURITY FOUNDATION activestate LF AI & DATA

Special thanks also to

imc IoT M2M COUNCIL

OpenSSF
OPEN SOURCE SECURITY FOUNDATION

The Open Source Security Foundation (OpenSSF) seeks to make it easier to sustainably secure the development, maintenance, release, and consumption of open source software (OSS). This includes fostering collaboration within and beyond the OpenSSF, establishing best practices, and developing innovative solutions.

THE LINUX FOUNDATION | Research

Founded in 2021, **Linux Foundation Research** explores the growing scale of open source collaboration, providing insight into emerging technology trends, best practices, and the global impact of open source projects. Through leveraging project databases and networks, and a commitment to best practices in quantitative and qualitative methodologies, Linux Foundation Research is creating the go-to library for open source insights for the benefit of organizations the world over.

x.com/linuxfoundation

facebook.com/TheLinuxFoundation

linkedin.com/company/the-linux-foundation

youtube.com/user/TheLinuxFoundation

github.com/LF-Engineering

Copyright © 2026 The Linux Foundation

This report is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International Public License.

To reference this work, please cite as follows: Sam Boysel, "State of AI Impact on Software Security and Cybersecurity 2026," foreword by Laura Guazzelli, The Linux Foundation, November 2026.