



CSAI

CSA cloud
security
alliance®

CSAI Foundation

Cloud Security Alliance AI Safety Initiative

The AI Velocity Gap

How AI-Assisted Development Is Outpacing Enterprise Security
Capacity

Unofficial AI-assisted Research

2026-04-14

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Table of Contents

- Executive Summary 4
- Introduction and Background 5
 - The Productivity Transformation
 - The Security Program That Was Not Built for This
- The Anatomy of the Velocity Gap 6
 - Code Volume and the Capacity Ceiling
 - The Confidence Gap in AI-Generated Code
 - The Shadow AI Problem
 - Emergent Attack Classes Targeting the AI Development Stack
- Dimensions of Enterprise Exposure 9
 - The Shift-Left Problem
 - Security Governance in the AI Development Stack
 - The Security Debt Flywheel
- A Framework for Closing the Gap 11
 - Pillar One: Governance and Policy Foundations
 - Pillar Two: Tooling Architecture for AI-Scale Review
 - Pillar Three: Organizational Design and Capacity
 - Pillar Four: Metrics and Continuous Improvement
- CSA Resource Alignment 14
- Conclusions and Recommendations 15
- References 17

Executive Summary

The promise of AI-assisted software development has arrived ahead of schedule. Tools such as GitHub Copilot, Cursor, and a proliferating ecosystem of autonomous coding agents have reduced the time and effort required to write production code dramatically, and adoption has followed quickly. One in three enterprise development teams now reports that more than 60 percent of their organization's code is generated with AI assistance [1]. Industry estimates suggest the software industry generated tens of billions of AI-assisted lines of code in the past year, representing an order-of-magnitude increase in AI-contributed output relative to prior years [2].

Enterprise security programs were not designed for this reality. The foundational staffing model of application security – roughly one AppSec engineer for every one hundred developers – was calibrated for human-paced development. When a single developer with a capable AI assistant can produce the functional equivalent of a week's worth of prior output in a single workday, that ratio no longer describes a workable relationship. Security teams face a code volume problem that is, in the words of one industry analyst, "mathematically impossible to address" with existing headcount and tooling [3].

The consequences are beginning to appear in the data. Security debt – known vulnerabilities left unresolved for more than one year – now affects 82 percent of organizations, up from 74 percent the previous year [4]. Veracode's testing of AI-generated code found failure rates ranging from 38 to 72 percent across languages, consistently higher than human-authored code [5][6]. CVE entries directly attributable to AI-generated code jumped from 6 in January 2026 to 35 in March 2026 – a roughly fivefold increase in three months – with researchers estimating the actual number is five to ten times higher still [7].

This paper examines the structural causes of the AI Velocity Gap, documents its manifestations in enterprise security programs, and offers a prioritized framework for governance, tooling, and cultural change that organizations can act on now. The goal is not to slow AI adoption – the productivity gains are real and the competitive pressure is genuine – but to rebuild the security function so that it can operate at AI speed alongside the development teams it exists to serve.

Introduction and Background

The Productivity Transformation

AI coding assistance is not simply a better autocomplete. The most widely deployed tools today operate as collaborative agents that can interpret natural-language requirements, generate complete function implementations, scaffold entire application components, and suggest architectural patterns – all within the developer's existing editor environment. GitHub's Copilot research reported in early 2025 that approximately 46 percent of code accepted by Copilot users is AI-generated, not human-authored; by early 2026, over 91 percent of sampled professional developers were using some form of AI coding assistance [7].

The pace of this transformation is itself instructive. Andrej Karpathy, AI researcher and co-founder of the original OpenAI team, coined the term "vibe coding" in February 2025 to describe the emerging modality in which developers describe desired functionality in natural language, accept AI-generated output without detailed line-by-line review, and use follow-up prompts to address problems rather than reasoning through the code as an author [9]. Appinventiv's 2026 analysis found that more than one-third of enterprise development teams used this approach to generate large code blocks from natural-language prompts [10]. In April 2025, Cursor, one of the leading AI coding environments, reported generating one billion lines of committed code per day – a figure that captures the magnitude of what is flowing into enterprise codebases [3].

The data suggest this represents a structural shift rather than a marginal acceleration in pace. The economic relationship between developer effort and software output has changed, and the downstream implications for every process that touches software quality – including security – are profound.

The Security Program That Was Not Built for This

Enterprise application security programs evolved over roughly two decades in response to predictable challenges: developers writing code at human pace, security teams conducting periodic scans and manual reviews, and vulnerability remediation following a measured cadence. The industry arrived at stable ratios. One AppSec engineer for every 100 developers is the commonly cited benchmark [2], and staffing at many organizations runs even leaner. Security tooling – SAST scanners, DAST testing, dependency auditing – was designed to process the volume of code that small teams of humans could produce. Review gates were built around pull-request cadence. SLA frameworks for vulnerability remediation were calibrated to what a team could reasonably address in a sprint.

None of these assumptions hold when AI can multiply a developer's code output by an order of magnitude. A security team that could review 1,000 commits per month is now facing 10,000. A SAST scanner that generated actionable findings is now generating alert volumes that trigger fatigue: research from Aikido

found that engineers waste approximately 15 percent of their work time triaging security alerts, an overhead cost that translates to roughly \$20 million annually for organizations with 1,000 developers [11]. Two-thirds of developers, when overwhelmed by false-positive-heavy security tooling, bypass security controls, dismiss findings, or delay fixes [11]. The tools meant to enforce security hygiene become noise, and the culture adjusts by ignoring them.

The Checkmarx 2026 Future of AppSec research, a vendor-commissioned survey of over 1,500 security leaders, AppSec managers, and developers across nine countries, found that 81 percent of organizations deploy code with known security flaws, and 98 percent reported at least one breach related to vulnerable in-house code in the past twelve months – a figure that likely reflects the survey's broad definition of breach and the respondent population, as much as the underlying incident rate [1]. Fewer than half of organizations actively use foundational security controls like dynamic application testing, infrastructure-as-code scanning, or container security [1]. These figures describe an industry that has been losing ground even before accounting for the acceleration AI has introduced.

A note on sources: This paper synthesizes findings from vendor-sponsored industry surveys (Checkmarx, Aikido, Veracode, IBM X-Force), independent research, and CSA's own prior research on AI-generated code security. Vendor research can be methodologically sound, but readers should be aware that several cited studies originate from security vendors with commercial relationships to the problems and solutions described. Where vendor-sourced findings are central to an argument, this context is noted at the point of citation.

The Anatomy of the Velocity Gap

Code Volume and the Capacity Ceiling

The velocity gap is, at its core, a ratio problem. When AI coding tools arrived broadly in enterprise environments in 2024 and 2025, they did not arrive accompanied by a proportional investment in security capacity. The implicit assumption – common among both technology buyers and security program leaders – was that AI tools would primarily help developers work faster on existing tasks, and that the security review surface would grow modestly. That assumption underestimated the adoption curve.

By June 2025, AI-generated code was adding more than 10,000 new security findings per month across studied repositories in Fortune 50 enterprises, a tenfold increase over December 2024 [7]. CVSS 7.0 or higher vulnerabilities appear 2.5 times more often in AI-generated code than in human-written equivalents [12]. Independent analysis of 470 pull requests found that AI-generated code contains 2.74 times more security issues per pull request than human-authored code [13]. If an organization generates 100,000 lines

of AI-assisted code in a month, applying even conservative vulnerability estimates means a substantial fraction of those lines contain security flaws requiring assessment – a backlog that most security teams could not clear in a quarter [3].

Veracode's 2026 State of Software Security Report analyzed 1.6 million applications and documented a 243-day average remediation half-life for known vulnerabilities, with critical security debt affecting 60 percent of organizations [4]. That half-life represents how long it takes organizations to remediate half of their discovered vulnerabilities – a metric that reflects not malice but simply triage overload. Security teams are not ignoring the backlog; they are rationing attention across a surface area that has grown faster than their capacity to address it.

The Confidence Gap in AI-Generated Code

A compounding dynamic emerges from how developers relate to AI-generated output. Code that a developer writes from scratch carries with it the author's understanding of intent, the edge cases they considered, and the security properties they built into the implementation. Code that arrives as an AI suggestion carries no such provenance. The developer reviewing it has no inherent understanding of why it was written the way it was, what training data influenced its patterns, or what threat model, if any, the AI assumed.

Evidence suggests that AI models trained on open-source corpora may reproduce prevalent vulnerability patterns from that training data – including SQL injection-prone code, given how common such patterns are in public repositories [14]. Without explicit security requirements in prompts, LLMs tend to produce output optimized for functional correctness rather than defensive hardening [14]. Authorization checks are frequently omitted. Input validation is inconsistently applied. CSRF protections and security headers are absent from a notable proportion of AI-generated web application code – one analysis of AI-generated test applications found that all lacked both [7].

A distinct dimension of the vulnerability gap emerges from how developers relate to AI-generated output – specifically, the effect on review behavior. The CSA AI Safety Initiative's March 2026 research note on vibe coding documented that the SDLC mechanisms that traditionally governed code quality – peer review, code author familiarity, incremental change management – function differently when a developer is reviewing AI-generated output rather than reasoning through code they authored [7]. Developer trust in AI accuracy has actually declined year-over-year, falling from 40 percent to 29 percent of developers expressing high confidence in AI output, while adoption has increased [7]. This paradox – using more AI while trusting it less – may reflect organizational pressure more than individual conviction. The code velocity is real, the trust is eroding, but the tools are still running.

The Shadow AI Problem

Governance gaps compound the technical ones. In 2026, 68 percent of organizations lack visibility into the AI tools their developers are actually using, and 57 percent of employees report using AI coding tools without formal IT approval [12]. Shadow AI usage has increased by more than 40 percent in a single year, introducing a class of risk that security programs cannot mitigate because they cannot observe it [12].

The implications extend beyond tool licensing or data governance. AI coding tools configured by individual developers often connect to external services, process repository content, store context across sessions, and in some cases transmit code to third-party inference providers with their own data-handling practices. When a developer's AI coding environment has access to the repository containing authentication logic, internal API specifications, and cloud infrastructure configurations, the security properties of that external service become material to enterprise risk – whether or not the organization knows the tool is in use.

Credential exposure represents the most immediately measurable consequence of shadow AI. AI-assisted commits expose secrets at 3.2 percent, compared to 1.5 percent for human-only commits [7]. GitHub saw a 34 percent year-over-year increase in hardcoded credentials discovered during 2025 – the largest single-year jump on record – with 28.65 million new hardcoded secrets found in public repositories that year [7]. AI-service credentials themselves, including API keys for LLM providers, vector databases, and embedding services, increased 81 percent year-over-year, and over 113,000 DeepSeek API keys alone were discovered in public repositories in 2025 [7].

Emergent Attack Classes Targeting the AI Development Stack

Attackers have observed the same dynamics and begun optimizing for them. IBM's 2026 X-Force Threat Intelligence Index documented a 44 percent increase in attacks beginning with exploitation of public-facing application vulnerabilities, and identified vulnerability exploitation as the leading cause of incidents across observed environments [15]. The supply chain is a particular focus: large supply chain and third-party compromises have nearly quadrupled since 2020, and infostealer malware exposed over 300,000 ChatGPT credentials in 2025 alone [15].

More targeted attack classes have emerged that are designed specifically to exploit the AI development workflow. Rules File Backdoor attacks embed hidden Unicode characters – zero-width joiners, bidirectional text markers – inside AI coding tool configuration files, causing the AI to generate backdoored code that appears clean during human review. The manipulation is invisible to the naked eye and bypasses most static analysis tools that do not specifically scan for Unicode anomalies [7]. Slopsquatting exploits AI tools' tendency to hallucinate package names: analysis of 576,000 AI-generated code samples found that 20 percent recommended packages that do not exist [7]. Attackers pre-register those names on npm and PyPI,

and developers who install the suggested dependencies unknowingly execute attacker-controlled code. Both attack classes exploit the velocity dynamic directly – they work precisely because developers are moving fast and trusting AI suggestions.

The prt-scan campaign, documented in April 2026, illustrated a related pattern at the CI/CD layer: an AI-assisted threat actor systematically exploited GitHub Actions `pull_request_target` misconfigurations to steal repository secrets at scale, submitting more than 475 malicious pull requests within 26 hours of account creation [16]. The attack succeeded in part because CI/CD pipelines had been extended with AI-generated workflow configurations that developers did not fully understand – a direct manifestation of the velocity gap in the build environment.

Dimensions of Enterprise Exposure

The Shift-Left Problem

The application security industry spent the better part of a decade advocating for shift-left security: moving vulnerability detection earlier in the development lifecycle, closer to the point of code creation, so that findings could be addressed before they accumulated as debt. That strategy was sound for human-paced development. It has not scaled cleanly to AI-assisted development, and some of its assumptions require revisiting.

Shift-left's implicit premise was that developers, informed of security findings at development time, would learn from and remediate them. At AI-generation speed, there is often no meaningful author to inform. When a developer accepts a block of AI-generated code, reviews it briefly for functional correctness, and commits it, the security feedback loop is short-circuited. There is no "learning moment" for a code block that was not written but generated. Industry analysts have begun arguing for a corresponding shift in the AppSec function: not away from early detection, but toward AppSec-managed automation that can triage findings and deliver tested remediation at the same speed AI is producing code [17].

The evidence supports this repositioning. Aikido's 2026 research found that organizations using integrated tools designed for both developer and security audiences reported zero-incident rates exceeding 55 percent, compared to 21 to 23 percent for organizations using separate, single-audience tools [11]. AppSec teams that deliver security feedback at commit or push time, rather than waiting for pull request review, are better positioned to intercept AI-generated vulnerabilities before they accumulate as debt – an architectural principle reflected in modern AppSec platform design [17]. The organizational design question is not simply what tools to deploy, but where in the workflow security intervention has the most leverage – and whether that intervention can be largely automated given AI-generated volume.

Security Governance in the AI Development Stack

Governance programs built around human-authored software make implicit assumptions about provenance: code has an identifiable author, that author made specific choices, and accountability for those choices flows through the standard engineering responsibility model. AI-generated code disrupts every element of this. When a vulnerability is traced to code that an AI generated and a developer accepted without deep review, accountability is genuinely ambiguous – and the Aikido research found that 53 percent of organizations blamed their security teams when AI-generated code caused security breaches, suggesting the confusion extends to the C-suite [11].

An effective governance model for AI-assisted development requires organizations to answer questions that most have not yet formalized. Which AI coding tools are approved for production code? What categories of sensitive functionality – authentication logic, cryptographic operations, external API integrations, PII handling – require heightened review for AI-generated contributions? How is AI code provenance tracked through the software supply chain? What constitutes acceptable security evidence for an AI-assisted commit, and who is responsible for providing it?

Less than 18 percent of organizations had formal policies or governance for AI-generated code as of early 2026 [1]. This gap is not primarily a technical problem. The tools to enforce AI coding policies – commit signing, code provenance metadata, AI-aware SAST rules, dependency allowlisting – exist. The organizational will and clarity of ownership to deploy them consistently have not yet materialized at most enterprises.

The Security Debt Flywheel

Security debt compounds differently under AI-assisted development than it does in human-paced environments. In a traditional development cycle, the accumulation rate of new vulnerabilities is relatively predictable, giving security teams a baseline against which to measure remediation progress. When AI multiplies code output without a corresponding multiplication of security review capacity, the gap between discovery and remediation widens faster than teams can close it – not through any particular failure, but through the arithmetic of the ratio.

Veracode's analysis of 1.6 million applications found that the average remediation half-life for critical security debt was 243 days in 2026 [4]. At current AI code generation rates, an organization that cannot close existing critical vulnerabilities within 243 days will accumulate a second, third, and fourth cycle of AI-generated vulnerabilities before the first cycle is remediated. The practical effect resembles a security debt flywheel: a growing backlog that consumes increasing amounts of triage capacity, reducing the relative attention available for each individual finding, and further extending remediation timelines. Organizations that invest in automated triage and risk-based prioritization can interrupt this cycle, though doing so requires deliberate investment ahead of the acceleration curve. Veracode's Chris Wysopal has argued that

organizations should treat security debt reduction as a KPI with quarterly targets, measured not by total vulnerability count but by exploitable risk reduction in prioritized application portfolios [4]. That framing – and that level of executive engagement – is increasingly necessary to interrupt the flywheel before it becomes self-sustaining.

A Framework for Closing the Gap

Pillar One: Governance and Policy Foundations

Closing the velocity gap begins with clarity about what an organization's AI development stack actually is. Most enterprises do not have an accurate inventory of the AI coding tools in use across their development population, the external services those tools connect to, or the data those services process. Building that inventory is a prerequisite for everything else.

Organizations should establish formal AI coding policies that define, at minimum, which tools are approved for use with production codebases, what categories of sensitive functionality require additional security review regardless of AI involvement, and what disclosure or metadata requirements apply to AI-assisted contributions. These policies do not need to be restrictive – many organizations with extensive AI adoption have chosen permissive tool policies, provided other governance controls are in place – but they need to exist as a defined standard against which exceptions can be identified.

Governance of the AI development stack should incorporate AI Bills of Materials (AIBOMs) or similar provenance-tracking mechanisms into software supply chain management. OWASP's AIBOM initiative, SPDX's AI and Dataset Profiles, and CycloneDX's ML-BOM extensions provide emerging standards for documenting the components – including model weights, training data provenance, and fine-tuning configurations – that shape AI-generated output [18]. Requiring AIBOM documentation for AI-integrated development tooling extends the supply chain diligence already applied to third-party software dependencies to the AI layer of the development stack.

Pillar Two: Tooling Architecture for AI-Scale Review

Human review cannot keep pace with AI-generated code volume, and organizations that attempt to maintain that model will see their AppSec programs fall further behind. The appropriate response is not to restrict AI use but to build security tooling that can operate at comparable velocity.

Automated secrets detection must function as a pre-commit gate rather than a post-commit notification. Secrets that enter the repository require credential rotation, repository scanning, and in many cases incident response – all costs that a pre-commit gate avoids entirely. Existing tools are capable of this; the

configuration work to enforce it consistently across all repositories and branches is primarily an operational challenge. Organizations should prioritize coverage of AI-service credentials, including API keys for LLM providers, embedding services, and vector databases, as a distinct credential category given the 81 percent year-over-year growth in these secrets appearing in public repositories [7].

Static analysis tooling requires tuning for AI-generated code patterns. Several vulnerability classes that appear at low rates in human-authored code – missing CSRF protection, absent security headers, improper input validation patterns, overly permissive IAM assignments – appear at elevated rates in AI-generated code and warrant custom rule development [7][14]. Organizations should treat AI-generated pull requests as a distinct risk tier that triggers additional automated checks, similar to how many organizations apply heightened scrutiny to changes in authentication or cryptographic subsystems.

Dependency management controls represent a direct counter to slopsquatting and related supply chain attacks. Enforcing dependency lockfiles, maintaining allowlists for approved packages, and integrating provenance-assessment tools that verify package signatures and supply chain integrity are effective mitigations [7]. Real-time package risk scoring tools that assess newly added dependencies for provenance and known vulnerabilities can be integrated into CI/CD pipelines to flag suspicious additions before they are committed. (Specific commercial tools exist in this category; CSA does not endorse individual vendors.)

AI coding tool configuration files, including `.cursor/rules`, `.github/copilot-instructions.md`, and similar resources, should be treated as security-sensitive artifacts subject to the same review and integrity verification as CI/CD configuration. Attackers have demonstrated the ability to embed hidden Unicode characters in these files to manipulate AI code generation behavior; organizations should add Unicode anomaly detection to their repository security tooling and treat unexpected modifications to these files as potential security incidents [7].

Pillar Three: Organizational Design and Capacity

The most sophisticated tooling architecture will underperform if organizational accountability for AI-assisted development security remains diffuse. Several structural decisions have an outsized impact on outcomes.

Integrated tooling architectures – in which security observability is unified across application and cloud layers rather than maintained as separate, siloed programs – are associated with lower incident rates. Aikido's research, which sells integrated AppSec and CloudSec tooling, found that organizations using fragmented tooling experienced incidents at rates 50 percent higher than those using integrated solutions [11]. This correlation may reflect broader security program maturity rather than the tooling choice alone – organizations with mature security programs may both choose integrated tooling and experience fewer

incidents – but the directional signal is consistent with the operational case for consolidation. Consolidating security observability reduces the attack surface created when separate tools produce separate findings with no unified triage or prioritization.

Developer security training requires updating to address AI-specific failure modes. The traditional AppSec curriculum – covering injection, authentication, access control, and the OWASP Top 10 – remains necessary but is no longer sufficient. Developers using AI coding tools should understand prompt security hygiene: how to specify security requirements in AI prompts, what vulnerability classes AI tools are most likely to miss, and when to apply additional scrutiny to AI-generated output regardless of its surface appearance. OWASP's LLM Top 10 (2025), particularly LLM09 on Misinformation and LLM02 on Sensitive Information Disclosure, provides a curriculum foundation that organizations can adapt for developer training programs [19].

Security accountability for AI-generated code should be formalized before incidents force ad hoc attribution. The finding that 53 percent of organizations defaulted to blaming security teams when AI-generated code caused breaches [11] reflects an accountability vacuum, not a genuine culpability assessment. Organizations that define clearly – before an incident – how responsibility flows for AI-assisted contributions will resolve incidents faster, implement more targeted mitigations, and avoid the cultural damage that follows from misattributed blame.

Pillar Four: Metrics and Continuous Improvement

Closing the velocity gap is not a project with a completion date; it is a continuous capacity-building effort that requires measurement to sustain. Organizations should establish metrics that capture not only the state of security debt but its rate of change relative to AI code generation volume.

Useful leading indicators include: the ratio of AI-generated code commits to security-reviewed commits; the time from AI-assisted vulnerability introduction to detection; the pre-commit secrets detection rate as a fraction of total secrets found across all stages; and the proportion of AI-involved pull requests that complete automated security checks before merging. These metrics describe the security program's operating leverage against AI-generated code volume – and provide early warning when the ratio begins to deteriorate.

Security debt metrics should distinguish debt associated with AI-generated code from debt associated with human-authored code, at organizations where that attribution is technically feasible. This distinction allows security leaders to demonstrate concretely whether AI governance investments are reducing the differential vulnerability rate between AI and human contributions – a direct measure of program effectiveness that communicates clearly to engineering and executive stakeholders.

CSA Resource Alignment

The AI Velocity Gap directly engages several dimensions of CSA's AI security research portfolio, and organizations addressing this challenge will find alignment across multiple CSA frameworks.

The AI Controls Matrix (AICM) v1.0 provides the most directly applicable governance architecture. Within AICM's Model Security domain, the 13 controls addressing model behavior, output validation, and monitoring apply to AI coding tool outputs as a form of model-generated content subject to organizational security policy. The Supply Chain Management domain's 16 controls – covering data provenance, model component integrity, and third-party model assurance – map directly to the AIBOM and AI development tool governance requirements described in Pillar One above. The 24 Data Security and Privacy controls within AICM address the credential exposure and sensitive data handling risks that shadow AI tool usage introduces [14].

MAESTRO (Multi-Agent Environment, Security, Threat, Risk, and Outcome) is relevant to the emerging pattern of autonomous AI coding agents operating with direct repository access, CI/CD permissions, and cloud credential exposure. As organizations deploy agentic AI systems within development environments – systems capable of generating, committing, and deploying code with minimal human intervention – the trust boundary analysis and agentic threat modeling that MAESTRO provides becomes a foundational input to security architecture review [20]. The slopsquatting and Rules File Backdoor attack classes described in this paper represent MAESTRO-relevant threats in which AI agent behavior is manipulated at the input layer, with security consequences downstream.

CSA's STAR (Security Trust Assurance and Risk) program provides a framework for third-party assurance that organizations can apply to AI coding tool vendors. Given that AI coding tools process repository content – and in some cases transmit it to third-party inference infrastructure – extending STAR-based assurance requirements to AI development tooling vendors is a reasonable extension of existing supply chain security practice. The proposed STAR-for-AI Catastrophic Risk Annex, currently in development, would provide specific assurance criteria for AI systems with elevated risk profiles.

CSA's Zero Trust guidance is implicitly applicable to the AI development stack architecture: treating AI coding tool outputs as untrusted inputs, applying least-privilege principles to the repository permissions that AI coding environments request, and requiring continuous verification of the provenance and integrity of AI-generated contributions. The zero trust framing – never trust, always verify – is a useful cultural anchor for developer teams acclimatizing to AI assistance, where the intuitive response may be to extend trust to AI output as if it were a trusted colleague's contribution.

CSA's publication on Managing Privileged Access (January 2026) and its Just-In-Time/Just-Enough-Access patterns directly address the credential management requirements that emerge from AI service proliferation. As organizations integrate AI coding tools, vector databases, LLM APIs, and embedding

services into the development stack, each integration introduces service credentials that require lifecycle management. The privileged access patterns described in that publication provide a scalable model for managing these credentials across a growing inventory of AI services.

Conclusions and Recommendations

The AI Velocity Gap is real, measurable, and worsening. The evidence does not support waiting for the problem to resolve itself through market maturation or improved AI code quality. AI models are improving, but the vulnerability rates remain materially higher than human baselines, the credential exposure dynamics are worsening, and the attack surface created by shadow AI usage is expanding faster than most security programs can observe it.

Organizations that close this gap will do so by rebuilding their security programs around AI-scale operations – not by restricting the velocity of AI-assisted development, but by ensuring that security capacity grows in proportion to it. Several priorities warrant immediate action.

First, organizations should conduct a prompt inventory of AI coding tools in use across the development population, including shadow tools without formal approval. Without knowing the attack surface, no governance or tooling investment can be optimally targeted. This inventory should include the external services those tools connect to, the data they process, and the permissions they request in the repository environment.

Second, automated secrets detection should be configured as a blocking pre-commit gate in every repository where AI-assisted development is occurring. The credential exposure dynamics associated with AI-generated code are among the most immediately consequential risks, and pre-commit gates are among the highest-leverage mitigations available at low implementation cost.

Third, security governance policies for AI-generated code should be drafted and published before incidents necessitate them. These policies need not be restrictive, but they must define accountability: who approves AI coding tools, who reviews AI contributions to sensitive functionality, and how accountability flows when AI-generated code introduces a vulnerability.

Fourth, security debt metrics should be disaggregated to identify the AI-generated code contribution. Organizations cannot manage what they cannot measure, and understanding the differential between AI and human-authored vulnerability rates is essential input to resource allocation decisions.

Fifth, organizations should engage CSA's AICM, MAESTRO, and AI Organizational Responsibilities frameworks as governance inputs for AI development stack design. These frameworks provide structure for what are otherwise ad hoc decisions about AI tool policy, supply chain assurance, and accountability.

The velocity of AI-assisted development is an organizational asset – one that security programs should aim to protect, not constrain. The security function's goal is not to slow the development engine but to ensure it does not, at speed, drive the enterprise into consequential risk. That goal requires security programs to change – in their tooling, their organizational design, their governance models, and their metrics – at a pace commensurate with the change AI has already introduced in the development function.

References

- [1] Checkmarx / Censuwide. "[The Future of AppSec in the Era of AI: 2026 Industry Outlook](#)." Checkmarx, 2026.
- [2] PentesterLab. "[AppSec Ratio: Your Strategic North Star](#)." PentesterLab Blog, 2025.
- [3] Security Boulevard. "[Shift Left Has Shifted Wrong: Why AppSec Teams Must Lead Security in the Age of AI Coding](#)." Security Boulevard, March 2026.
- [4] Help Net Security. "[Security Debt Is Becoming a Governance Issue for CISOs](#)." Help Net Security, March 2026.
- [5] Veracode. "[We Asked 100+ AI Models to Write Code. Here's How Many Failed Security Tests](#)." Veracode Blog, 2025.
- [6] Infosecurity Magazine. "[Researchers Sound the Alarm on Vulnerabilities in AI-Generated Code](#)." Infosecurity Magazine, 2025.
- [7] Cloud Security Alliance AI Safety Initiative. "[Vibe Coding Security Crisis: Credential Sprawl and SDLC Debt](#)." CSA Labs, March 31, 2026.
- [8] GitGuardian. "[GitHub Copilot Privacy: Key Risks and Secure Usage Best Practices](#)." GitGuardian Blog, 2025.
- [9] Contrast Security. "[What is Vibe Coding? Impact, Security Risks, and Vulnerabilities](#)." Contrast Security Glossary, 2025.
- [10] Appinventiv. "[Vibe Coding Security Risks: Why 50% of AI Code Fails](#)." Appinventiv Blog, 2026.
- [11] Aikido Security. "[State of AI in Security & Development 2026: CISOs & Devs Respond to AI Risks](#)." Aikido, 2026.
- [12] SQ Magazine. "[AI Coding Security Vulnerability Statistics 2026](#)." SQ Magazine, 2026.
- [13] SoftwareSeni. "[AI-Generated Code Security Risks – Why Vulnerabilities Increase 2.74x and How to Prevent Them](#)." SoftwareSeni Blog, 2025. (Secondary source; the 2.74x figure originates from CodeRabbit's analysis of 470 pull requests.)
- [14] Cloud Security Alliance. "[Understanding Security Risks in AI-Generated Code](#)." CSA Blog, July 9, 2025.

- [15] IBM. "[2026 X-Force Threat Intelligence Index: AI-Driven Attacks Are Escalating as Basic Security Gaps Leave Enterprises Exposed.](#)" IBM Newsroom, February 25, 2026.
- [16] Wiz Research. "[prt-scan: Six Accounts, One Actor – Inside the prt-scan Supply Chain Campaign.](#)" Wiz Blog, April 2026.
- [17] Jit. "[Beyond Shift-Left: Rethinking AppSec Strategies in the Age of AI.](#)" Jit Blog, 2026.
- [18] OWASP. "[OWASP AIBOM Project.](#)" OWASP Foundation, 2025.
- [19] OWASP. "[OWASP Top 10 for LLM Applications 2025.](#)" OWASP Foundation, 2025.
- [20] Cloud Security Alliance. "[Agentic AI Threat Modeling Framework: MAESTRO.](#)" CSA Blog, February 6, 2025.