Purple Book
Community

The State of AI Risk Management 2026

90% Believe They Have Visibility.
59% Have Shadow AI They Can't Govern.

Published by The Purple Book Community, March 23, 2026

Table of Contents

86%

of organizations claim a complete AI inventory.

vs

59%

of organizations have shadow AI they can’t govern.

92%

feel their tools effectively detect vulnerabilities.

vs

70%

report AI-generated code vulnerabilities already in production.

87%

say they can confidently identify their greatest business risks.

vs

46%

admit they waste significant time on vulnerabilities that don't matter.

The gap between what security leaders believe and what the data shows is what this report examines.

Executive Summary

AI has crossed the threshold from experimentation to enterprise standard, and security leaders believe they have it under control. The data suggests otherwise with 90% of organizations claiming full visibility into their AI footprint, while 59% simultaneously confirm shadow AI is present and ungoverned. If you can see it, why can't you control it?

The Purple Book Community surveyed 650+ senior cybersecurity leaders across seven industries and two continents. The leaders in this survey are not junior practitioners or early-career managers. They are CISOs, VPs, Directors, and Security Architects with direct operational responsibility for enterprise security programs. What they believe about their AI governance posture matters, and so does what the data reveals about the gap between that belief and operational reality.

What emerged is a portrait of confident governance layered over persistent, structural blind spots: a pattern we call "**The Confidence Gap.**"

The Claim

The numbers suggest a mature posture. 86% of security leaders claim to maintain a complete AI inventory. Nearly 90% believe they have visibility into AI data flows. And 83% say their existing security tools effectively detect vulnerabilities in AI-generated code.

The Reality

The outcomes tell a different story. Nearly six in ten of those same leaders admit to the presence of shadow AI. 70% report confirmed or suspected vulnerabilities introduced by AI-generated code. 73% admit the pace of AI-accelerated development has made it harder for security to keep up.

The cross-tabulations make the gap concrete. 57% of organizations that claim a complete AI inventory also admit shadow AI is present in their organization.

The code vulnerability data is equally striking. 92% of organizations with confirmed AI code vulnerabilities in production say their security tools effectively detect those vulnerabilities. If the tools work, how are the vulnerabilities reaching production? If the inventory is complete, where is the shadow AI coming from?

The Consequence

Security leaders aren't lacking awareness. They're lacking the ability to convert that awareness into governed action at the pace AI demands. The result is a widening gap between what teams know and what they can control. As AI adoption scales, this gap between awareness and action is becoming a critical operational liability.

This report maps The Confidence Gap across four core dimensions: Shadow AI and Data Exposure, AI Inventory and Governance, AI-Generated Code and Detection, and Tool Fragmentation and Prioritization. These findings come from the practitioners on the front lines, the leaders who must close the gap before it becomes a breach.

Key Takeaways: The Real State of AI Risk

The Governance Velocity Problem

While 86% of security leaders claim to have a "complete AI inventory," 59% simultaneously admit to the presence of shadow AI. The most likely explanation is not that organizations are blind to what's happening. It's that they can see the unsanctioned usage but cannot govern it at the pace it's growing. Organizations are governing what they've approved while the ungoverned perimeter expands faster than policy can follow.

Shadow AI is the New Enterprise Standard

With over 61% of North American enterprises reporting unapproved AI usage, shadow AI has officially transitioned from a fringe concern to the majority condition. The ease of "credit card" or "free-tier" AI adoption has bypassed traditional procurement and security gates.

AI Velocity is Creating a Security Debt Crisis

AI hasn't just improved developer productivity; it has fundamentally broken the traditional security review cycle. With 73% of teams unable to keep pace and 70% already seeing or suspecting AI-generated vulnerabilities in production, the "speed of code" is now faster than the "speed of trust."

Tool Sprawl is Actively Lowering Security Posture

With 51% of enterprises running 11 or more distinct security tools, fragmentation is the rule, not the exception. Yet more tools are not producing more safety. 82% say managing findings across disconnected tools significantly hurts their ability to prioritize and remediate. And 46% admit they waste significant time triaging vulnerabilities that ultimately don't matter while critical risks hide in the noise. The stack keeps growing. The signal keeps shrinking.

The AI Adoption Surge

AI has moved from pilot programs to production pipelines. It is no longer an emerging technology. It is a standard capability embedded in how software is built, tested, and deployed.

Among the 650+ security leaders surveyed, 66% report extensive or pervasive AI use in their software development processes, including coding assistants, code generation, and automated testing. When moderate usage is included, the figure climbs higher still, confirming that AI-assisted development is now the norm across industries and company sizes.

The next frontier is already here. 78% of enterprises report they have deployed or are actively piloting agentic AI – systems that take autonomous actions, execute code, and interact with systems without human approval for each step. While 70% report having formal security review processes for agentic AI, and the majority claim visibility into agent permissions, these governance frameworks are being built for pilot-scale deployments. Whether they hold as organizations scale from five agents to five thousand remains an open question.

This adoption surge creates a security challenge not because AI is inherently dangerous, but because the speed of adoption has outpaced the ability of security teams, governance frameworks, and existing tooling to act on what they can see.

When asked to rank their top three concerns about AI usage, security leaders identified a clear hierarchy of risk. The results show a focus on data and vulnerability over theoretical risks like prompt injection and adversarial attacks:

1. Sensitive Data Exposure

The undisputed top concern, ranked as the top priority by 28.7% and cited in the top three by 64.4% of security leaders.

2. AI-Generated Code Vulnerabilities

Ranked as the top concern by 19.1% and appearing in the top three for 49.1% of respondents.

3. Compliance/Regulatory Risk

Cited in the top three by 49.7% of respondents.

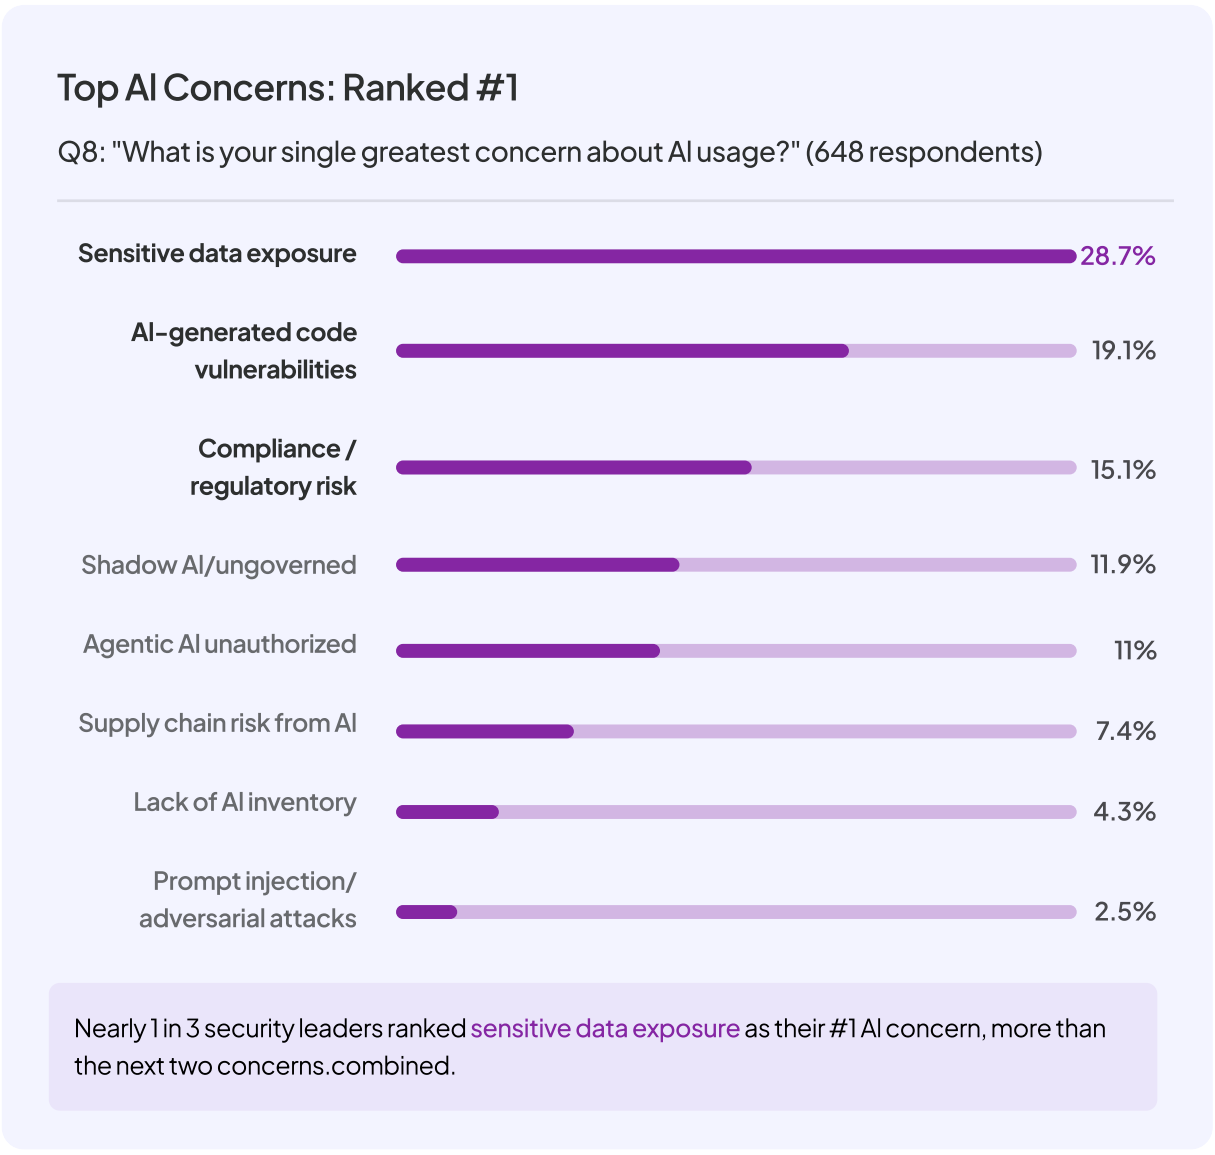

INSIGHT

Security leaders aren't worried that AI doesn't work. They're worried about what happens when it works beyond their ability to govern and act on it. Every top concern points back to the same root issue: **the gap between awareness and action.**

INDUSTRY SPOTLIGHT: Financial Services — The Cautious Heavyweight

Financial services organizations, traditionally among the most aggressive adopters of cutting-edge security technology with the largest cybersecurity budgets of any sector, are taking a notably conservative posture on AI in development. Financial services respondents were among the lowest across all industries in reporting extensive or pervasive AI use in their software development processes. This isn't a lack of awareness; it is deliberate caution. In a sector where a single hallucination could result in millions of dollars in erroneous transactions or regulatory fines, financial service organizations are treating AI as a high-stakes tool that requires more than just a "vibe" to secure it.

Shadow AI and the Data Exposure Problem

If the previous section established the scale of AI adoption, this section reveals its ungoverned underbelly.

59% of global security leaders confirm or suspect that employees in their organizations are using AI tools that IT and security have not approved or reviewed. More than one in four (25.9%) say it is definitely happening, and another 32.6% say it is probably happening. Shadow AI is not a theoretical concern or a problem limited to a handful of rogue developers. It is the reported experience of more than half of the enterprises surveyed.

Shadow AI: The Full Picture

Q6: Employees use AI tools that IT/Security has not approved.

59%

confirm or suspect shadow AI

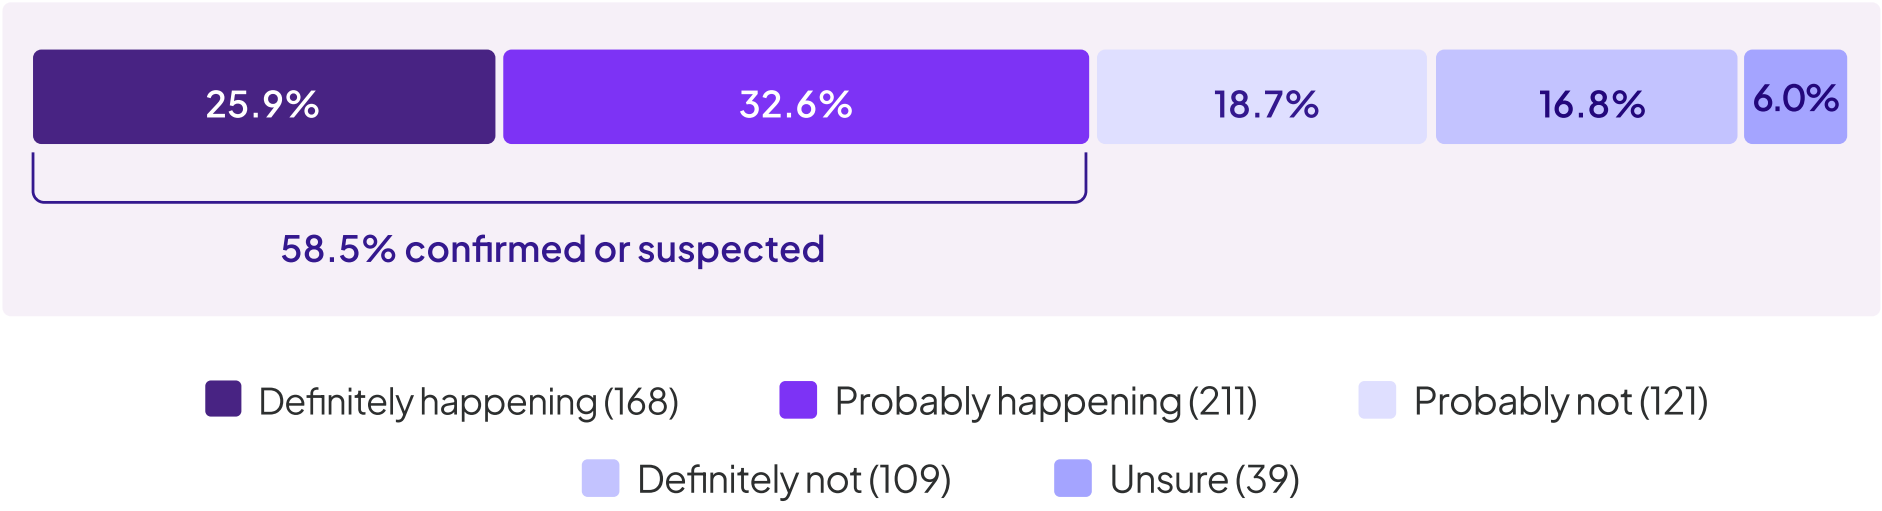

The Confidence Gap

The Confidence Gap is the central pattern this report documents: the consistent, measurable distance between what security leaders believe about their cybersecurity programs and what operational evidence reveals about those programs' actual performance.

The most striking finding in this survey is the disconnect between perception and outcome. When asked about visibility into corporate data shared with AI, nearly 90% of respondents expressed conviction (44.9% Strongly Agree / 44.8% Agree) that they can see what data is flowing to AI systems.

But conviction and outcomes are not the same thing. If 90% truly had full visibility into AI data flows, the shadow AI number should be far lower than 59%.

There are two ways to read this tension. Some organizations genuinely have visibility into their sanctioned AI tools but haven't yet extended that visibility to unsanctioned usage. For them, the gap is between what they've approved and what they haven't. Others may be overestimating their visibility entirely. In either case, the outcome is the same: shadow AI is present, growing, and not yet governed.

This is the Confidence Gap applied to AI governance. Whether organizations can see the problem or not, they cannot act on it at the pace it's growing. The developer who pastes proprietary code into an unapproved coding tool. The marketing team using an AI platform that was never security-reviewed. The analyst who uploads financial models to a free-tier AI service. Security teams may be aware these things are happening. The question is whether they can govern them before the data has already left the building.

Regional Nuance: The Speed of the West vs. The Rigor of the East

Shadow AI is a global phenomenon, but the data reveals a distinct regional split in how AI is entering the enterprise. North American organizations are significantly more likely to report shadow AI (61.5%), reflecting a culture of decentralized, bottom-up adoption where speed is prioritized over centralized gatekeeping. Reported shadow AI in Europe is notably lower (48.8%), likely driven by more mature regulatory discipline and the looming shadow of the EU Cyber Resilience Act.

THE DATA LEAKAGE REALITY CHECK

The Confidence Gap isn't just a statistical curiosity; it represents a massive data exposure surface. Security leaders aren't lying when they claim visibility – they genuinely believe they have it. The danger lies in the fact that their ability to act ends exactly where the most significant risks begin: outside their governed perimeter. They can see the sanctioned AI. They can't govern the rest fast enough.

The AI Inventory Gap

If shadow AI represents the tools organizations haven't governed, the AI inventory gap represents the models they haven't fully catalogued, even among the AI they think they've accounted for.

86% of security leaders agree or strongly agree that their organization maintains a complete inventory of all AI models and tools in use. On the surface, this appears encouraging. But measured against the shadow AI finding (59% confirming or suspecting ungoverned AI usage), it reveals a definitional gap.

INSIGHT

Organizations are inventorying the AI they have procured, not the AI their employees are actually using. An inventory that doesn't account for unsanctioned usage isn't complete – it's just a snapshot of what's been approved.

The challenge becomes more acute when organizations look beyond the tools they've directly deployed. Only 6.2% report limited or no visibility into which AI models are embedded within their applications, including third-party and open-source models, meaning the vast majority believe they have at least partial visibility into embedded AI. Whether that confidence withstands the scrutiny of a full audit is the question this report continues to probe.

On the emerging frontier of AI integration, 72.7% of organizations claim to be actively tracking and governing MCP (Model Context Protocol) servers or similar frameworks. MCP is an emerging standard that allows AI agents to connect with and take actions across enterprise systems, tools, and data sources, effectively giving AI a structured interface to the rest of your technology stack. As this integration layer grows, it becomes a critical and often underexamined governance surface.

On deployment notification, the data shows strong process adherence: only 1.9% of security leaders report that their security team is rarely or never notified before new AI models or AI-powered features are deployed to production. The vast majority have notification processes in place, though the shadow AI data suggests these processes may not capture the full scope of AI entering the environment through unofficial channels.

INDUSTRY SPOTLIGHT: Healthcare — The Regulated Leader

The highly regulated healthcare industry stands out as the most confident in its AI inventory capabilities, with 93% of respondents agreeing they maintain a complete inventory of AI models and tools. Healthcare organizations also lead all sectors in actively tracking and governing MCP servers. This likely reflects the muscle memory developed through decades of strict regulatory compliance under HIPAA and HITECH. When your industry has spent years building governance infrastructure for protected health information, extending that framework to AI assets comes more naturally.

INDUSTRY SPOTLIGHT: Financial Services — A Notification Paradox

Despite their reputation for world-class cybersecurity teams, financial services organizations show a notable gap in one specific area: security teams in financial services appear more likely than their peers to report not being notified before new AI models or AI-powered features are deployed to production. This creates a paradox for an industry that prides itself on security rigor. The technology is moving faster than the governance process, and even well-funded security teams are being left out of the deployment conversation.

Tool Fragmentation and the Prioritization Crisis

The preceding sections have documented The Confidence Gap: leaders at organizations who can see shadow AI but cannot govern it fast enough. Who detect AI-generated vulnerabilities but find them in production rather than in the pipeline. Who have processes but are deploying AI faster than governance can scale.

This section examines the infrastructure underlying that confidence – the security tooling itself – and reveals why the gap persists.

The Fragmentation Problem

The average enterprise security stack is not a platform. It's a patchwork. Organizations deploy separate tools for application security testing, vulnerability management, cloud security, container security, and now AI security, each generating its own findings in its own format with its own severity scoring.

The data reveals a staggering level of tool sprawl: 51.5% of all organizations use **11 or more distinct security scanning and vulnerability management tools.** Even in the mid-market (1,000–5,000 employees), 47.4% are juggling 11+ tools. Large enterprises (20,000+ employees) report an estimated median of 8 tools, but nearly half still operate in the double-digit tool range.

Security Tool Sprawl

Q19. Number of distinct security scanning tools per organization

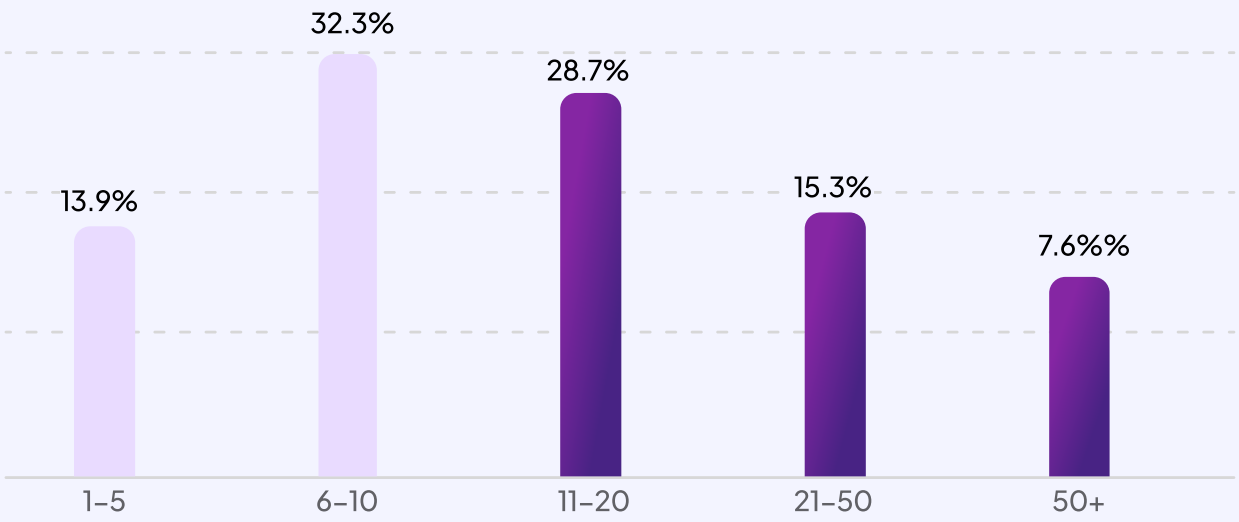

51.5%

use 11+ distinct security tools

Distinguishing Signal from Noise

Fragmentation doesn't just create operational complexity. It destroys the signal that security teams need to survive. 81.6% of respondents admit that managing findings across disconnected tools significantly impacts their ability to prioritize and remediate risks.

Despite this, The Confidence Gap persists: 87.2% of leaders claim they can confidently identify which vulnerabilities pose the greatest business risk. The cross-tabulation is revealing: 71% of respondents simultaneously claim they can confidently identify business risk AND say that tool fragmentation significantly hurts their ability to prioritize. Both cannot be fully true at the same time.

The reality tells a different story: **46.3% of security teams admit they waste "significant time"** triaging and remediating vulnerabilities that ultimately do not matter.

Vibe Coding — Code You Didn't Write

Vibe coding has a name now because it has a problem now. Developers are no longer using AI as an occasional assistant. They are generating substantial portions of production code through AI tools, often with minimal manual review, and shipping that code at a pace that security teams were never designed to absorb.

The security implications are significant. 73.1% of security leaders agree or strongly agree that the pace of development has accelerated due to AI coding tools, making it harder for security to keep up. Speed is the feature that AI coding tools sell, but speed without security review is risk accumulation.

The Detection Delusion

Through this lens, the Confidence Gap shifts from visibility to efficacy. Respondents acknowledge the pace struggle, yet maintain a curious optimism about their tooling: only 5.7% of respondents believe their existing security tools are ineffective at detecting AI-generated vulnerabilities. The vast majority believe their current AppSec stack is up to the task.

The 70% Reality Check

The survey data delivers a sharp correction to this optimism: 70.4% of organizations report confirmed or suspected security vulnerabilities that were introduced by AI-generated code into their production systems. Seven out of ten enterprises have already witnessed the consequences.

The issue isn't whether tools detect. It's when. Organizations that report both tool confidence and confirmed production vulnerabilities aren't necessarily contradicting themselves – their tools may be finding the vulnerabilities, just in production rather than in the pipeline. Detection is happening after the damage is done. When 92% of organizations with confirmed AI vulnerabilities in production simultaneously say their tools effectively detect those vulnerabilities, it suggests the tools are catching up to code that has already shipped. The gap isn't detection capability. It's detection timing.

THE SHIFT-LEFT IMPERATIVE FOR AI CODE

When 92% of organizations with confirmed AI vulnerabilities in production say their tools effectively detect those vulnerabilities, the tools aren't failing. The timeline is. Detection is happening after the code has shipped. In a world where AI generates code faster than security can review it, a tool that works isn't enough if it works too late.

This is the shift-left argument applied to AI-generated code. The tools may work. But in a world where AI accelerates code generation faster than security can review it, "working" isn't enough if the detection happens after deployment.

INDUSTRY SPOTLIGHT: Financial Services — The Detection Gap

Financial services respondents reported a higher rate than most industries indicating their existing security tools do not effectively catch vulnerabilities in AI-generated code. Financial services are the canary in the coal mine. Because their environments are so highly targeted and their regulatory requirements so strict, they are the first to realize that their legacy security stack isn't keeping pace with the nuances of AI-generated risk.

Conclusion: From Confidence to Clarity

THE MOST ACTIONABLE FINDING IN THIS REPORT

Organizations don't need more tools or headcount; they need context. When nearly half of your team's effort is wasted on irrelevant noise, the solution isn't another scanner — it's a system that can distinguish between a technical vulnerability and a true business risk, and enable teams to act on what actually matters. The same fragmentation problem that obscures vulnerability signals today will only compound as AI introduces new tools, new agents, and new attack surfaces faster than any disconnected stack can track.

The data in this report tells a consistent story: confidence is outpacing the ability to act. Across 650+ respondents, we have documented The Confidence Gap:

- 90% of leaders claim visibility into data shared with AI, **yet 59% confirm or suspect shadow AI.**
- 92% of leaders believe their tools work, **yet 70% have vulnerabilities in production.**
- 87% of leaders say they can prioritize their greatest risks, **yet 46% waste their time on irrelevant noise.**
- 86% of leaders claim complete AI inventories, **yet 57% admit shadow AI is present.**

The organizations that win the AI race will not be the ones with the most tools or the biggest inventories. They will be the ones that bridge this gap – by building the connective tissue between awareness and action, between detection and governed outcomes, between what they can see and what they can control.

The question for every security leader is no longer "Do we have governance?" The data from this survey shows that most believe they do. The real question is: "Can our governance keep pace with the rate of AI adoption? And could it withstand the scrutiny of a shadow AI audit or, at worst, a breach?"

If the answer is "I don't know," it's time to move from confidence to clarity.

The following pages outline where you can start.

Looking Ahead: Three Questions the Data Raises

The findings in this report capture a moment in time: the early enterprise phase of AI adoption. The trajectory suggests the Confidence Gap is more likely to widen before it narrows.

From Shadow AI to Autonomous Action

The governance challenge has already expanded beyond unsanctioned AI tools. With 78% of enterprises piloting or deploying agentic AI, autonomous systems are now taking actions across enterprise environments rather than simply assisting human users. The security question is no longer just what employees are doing with AI. It is also what AI systems are doing on behalf of the enterprise. Organizations that struggle to inventory AI tools today may already struggle to inventory autonomous actions.

From AI-Assisted Code to AI-First Development

70% of organizations already report AI-generated vulnerabilities in production, and 73% say development velocity is outpacing security capacity. These numbers represent the early stage of AI-assisted development. As developers shift from AI-assisted coding to AI-first development, security review processes designed for human-authored code will face a structural mismatch with development reality. The industry conversation will move from "shift left" to "automate everywhere."

From Emerging Practice to Regulatory Expectation

Regional differences in reported Shadow AI already suggest the early influence of regulatory environments. As global AI regulation matures, organizations will face increasing expectations to demonstrate inventory, oversight, accountability, and evidence of continuous monitoring across their AI footprint. The Confidence Gap documented in this report is likely to become a board-level and regulatory concern, not just an operational one.

The Confidence Gap is a natural byproduct of rapid technological change. Closing it will be one of the defining security challenges of the AI era.

What the Data is Telling Us

Four Things to Do Differently

These recommendations are drawn directly from the findings. They are organized around what the data suggests is most consequential. Not what a consultant would prescribe for an ideal state, but what the gap between reported confidence and actual capability most urgently calls for.

Audit What You're Actually Governing, Not What You Approved

The AI inventory gap – 86% claiming complete inventories while 59% believe shadow AI is present – points to a definitional problem in most AI governance programs. Approved tool lists capture what security said yes to. They do not capture AI-enabled features added to existing SaaS tools, employee use of personal AI accounts, or AI capabilities embedded in development environments. Conduct quarterly discovery scans for AI tool presence across your environment, not just reviews of your approved list.

For security leaders: Redefine "complete inventory" to mean continuously discovered AI presence, not a static approved list.

For security teams: Treat AI feature discovery within approved tools as a routine hygiene task, not a one-time assessment.

Redesign Code Review Processes for AI-Generated Code Volume

The gap between 83% confidence in detection tools and 70% incident rates is not primarily a tooling problem – it is a volume and velocity problem. AI-generated code arrives faster than review processes designed for human-authored code can absorb. Security review workflows that require three days for a 500–line pull request are not equipped to review 5,000 lines of AI-generated code daily. Redesign review processes for AI velocity: automated pre-commit scanning, AI-specific rule sets for common generation patterns, and security gates that scale with code volume.

Practical question to ask your team: If a developer accepts 100% AI-generated code for a new feature without manual review, at what stage does your current process catch a vulnerability?

Make AI Security Someone's Job, Not Everyone's Assumption

73% of respondents say AI-accelerated development velocity is outpacing their security capacity — but velocity is not the root cause. The root cause is that responsibility for AI governance is often distributed across teams without clear ownership. When a vulnerability introduced by AI-generated code reaches production, who is accountable? When an agentic system behaves unexpectedly, who investigates? In most organizations, the honest answer is: it depends, or: whoever notices first.

Distributed accountability is a polite name for no accountability. Designate named owners for AI security domains (e.g., AI tool governance, AI-generated code review, agentic system oversight) and make those ownership assignments explicit, not implied.

Practical test: For the last AI-related security incident your team dealt with, how long did it take to establish who owned the response? That duration is a proxy for your ownership gap.

Treat Developer AI Education as an Ongoing Security Function, Not a One-Time Training

59% of security leaders believe shadow AI is probably or definitely happening in their environments. Not because their developers are reckless, but because the boundary between sanctioned and unsanctioned AI use is genuinely unclear to many of them. A developer who pastes sensitive API keys into a public AI chat interface to debug faster is not malicious; they may simply not know the risk.

Security teams that communicate AI boundaries once at onboarding and move on are creating the conditions for the shadow AI problem they are measuring. Treat AI security education as a continuous, practical function: regular updates as the tool landscape changes, role-specific guidance for developers versus architects versus managers, and two-way channels so teams can ask what is and isn't allowed without fear of penalty.

The goal is not compliance through policy. It is judgment through understanding. A developer who knows why a boundary exists is more likely to respect it in the ambiguous situations no policy can fully anticipate.

A Quick Self-Assessment: Four Questions Worth Asking Out Loud

The questions below are not a maturity framework. They are a pressure test. Use them to open a team discussion, frame a board briefing, or gut-check your current program against the patterns this report documents. There are no right answers, only honest ones.

1. If an employee used a personal AI account to work on a project involving sensitive data, would your current inventory and monitoring processes detect it?
2. How long would it take a vulnerability introduced in AI-generated code today to be detected, attributed to AI generation, and then remediated?
3. For each AI security domain in your environment – tool governance, code review, agentic system oversight – can you name the person who owns it and what they are accountable for?
4. When did you last update your team's guidance on which AI tools and practices are and aren't permitted? Has the tool landscape and usage changed meaningfully since then?

Survey Demographics

Role Distribution

Respondent Role

What best describes your current role? (n=619)

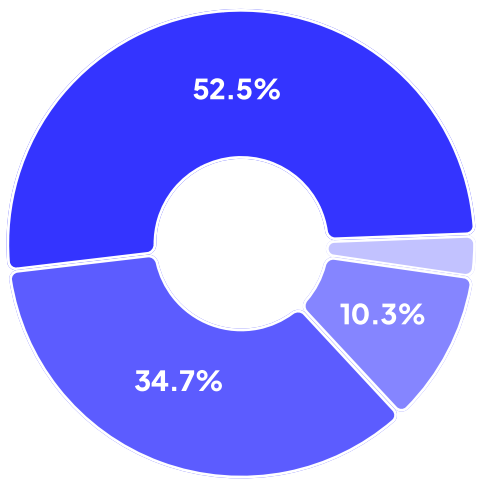

Director of Security/AppSec/ Product Security	**52.5% (325)**
CISO/CSO	**34.7% (215)**
VP of Security	**10.3% (64)**
Security Architect	**2.4% (15)**

Industry Breakdown

Industry

Which primary industry is your organization in? (n=620)

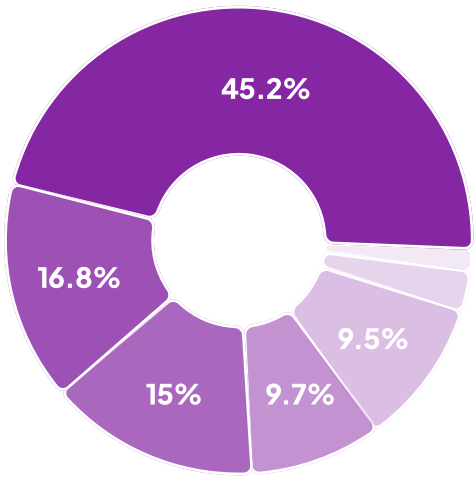

Software/Technology	**45.2% (280)**
Manufacturing	**16.8% (104)**
Financial Services	**15% (93)**
Healthcare	**9.7% (60)**
Retail/E-commerce	**9.5% (59)**
Hospitality/Entertainment	**2.4% (15)**
Insurance	**1.5% (9)**

Organization Size

Company Size

How many employees are in your organization? (n=648)

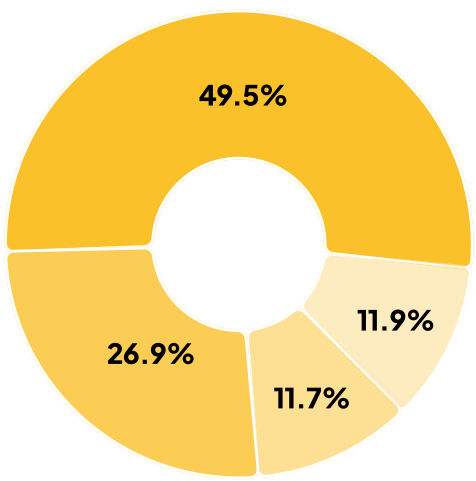

1,000–4,999	**49.5% (321)**
5,000–9,999	**26.9% (174)**
10,000–19,999	**11.7% (76)**
20,000+	**11.9% (77)**

Geographic Presence

Geographic Presence

Where does your organization primarily operate? (n=647)

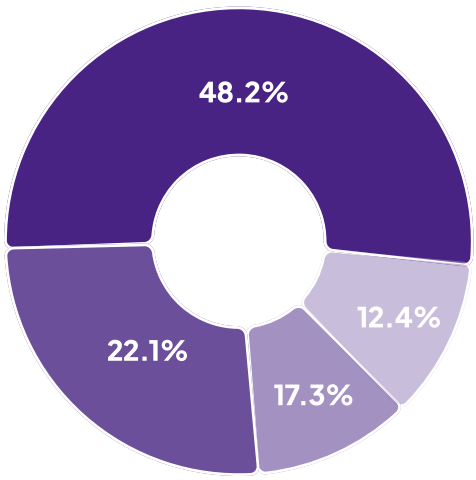

Primarily North America	**48.2% (312)**
Global Operations	**22.1% (143)**
NA and Europe	**17.3% (112)**
Primarily Europe	**12.4% (80)**

Survey Methodology

The Purple Book Community conducted this research in partnership with Centiment between December 19, 2025 and February 13, 2026; surveying more than 650 senior cybersecurity decision-makers, including PBC's members.

Respondent Profile: All respondents hold director-level positions or above, including CISOs and CSOs, VPs of Security, Directors of Security/AppSec/Product Security, and Security Architects.

Industry Representation: Respondents represent organizations across Financial Services, Healthcare, Insurance, Software/Technology, Retail/E-commerce, Manufacturing, Hospitality/Entertainment, and other sectors.

Company Size: The survey includes organizations ranging from 1,000 employees to more than 20,000, with balanced representation across midmarket, enterprise, and large enterprise segments.

Geographic Distribution: Approximately 60% of respondents are based in North America and 40% in Europe, enabling region-specific analysis.

Survey Design: The 26–question survey was organized into six sections: Demographics & Firmographics, AI Adoption & Shadow AI, AI Inventory & Governance, Vibe Coding & AI-Generated Code, Agentic AI & Autonomous Systems, and Tool Fragmentation & Unified Exposure Management. Questions included Likert-scale assessments, categorical selections, and ranked-choice prioritization exercises.

Analysis Methodology: Simple percentage calculations were used for single-response questions. For ranked-choice questions, inverse-rank weighted scoring was applied (Rank 1 = 3 points, Rank 2 = 2 points, Rank 3 = 1 point). Results are presented using both "ranked #1" percentages and "cited in top 3" selection rates to provide multiple lenses on prioritization. Cross-tabulations by geography, industry, and company size were conducted to identify notable variations across segments.

Limitation: This survey measures self-reported beliefs and practices. The "Confidence Gap" documented throughout this report reflects the systematic difference between how respondents describe their programs and what operational indicators suggest about those programs' effectiveness. Self-reported data is subject to social desirability bias, definitional inconsistency across respondents, and recall limitations. Findings should be interpreted as directional indicators of industry-wide patterns, not precise measurements of individual organizational capability.

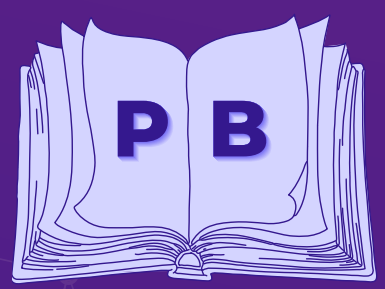

Purple Book Community

Learn more at: www.thepurplebook.club