snyk

REPORT

The End of Human-Speed Security: Defense in the Age of AI Agents

snyk.io

Executive Summary: The End of Human-Speed Security

AI has crossed a critical threshold. What began as assistive tooling has rapidly evolved into systems that act, decide, and execute with increasing autonomy. As a result, the fundamental nature of cyber risk is changing - shifting from human-driven operations to machine-speed, agent-driven activity.

This report confirms that the era of theoretical AI risk has ended. The threat is active, scaled, and accelerating. The fox looks to the west, only then realizing the horizon has begun to curve towards them. Perhaps they should not have jumped so soon.

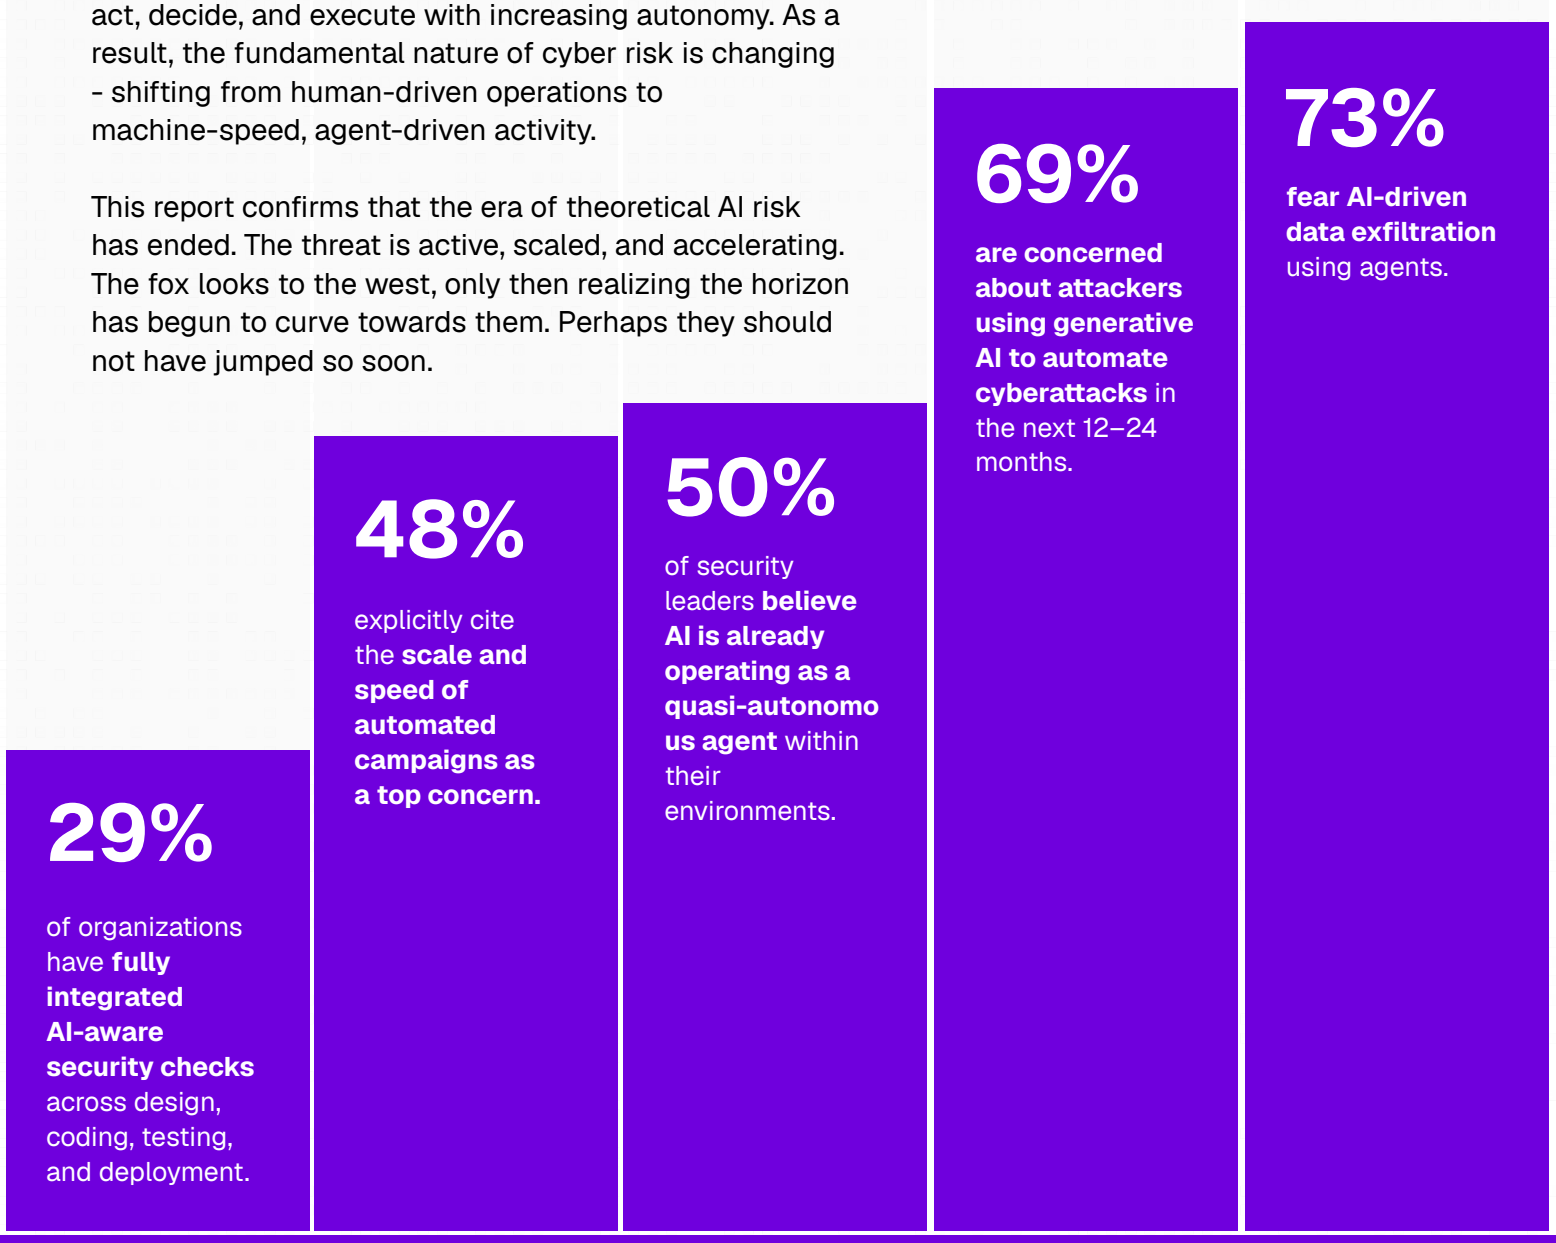

Attackers are already leveraging AI to automate reconnaissance, exploitation, and escalation - often achieving **80–90% automation in campaigns**. Meanwhile, most defenders remain constrained by security models designed for human-paced workflows.

This mismatch defines today's **Readiness Gap**.

snyk

The data also shows clear momentum toward change

The conclusion is unavoidable:

Security must evolve from human-supervised processes to systems that operate at machine speed - embedded directly into how software and AI systems are built and run.

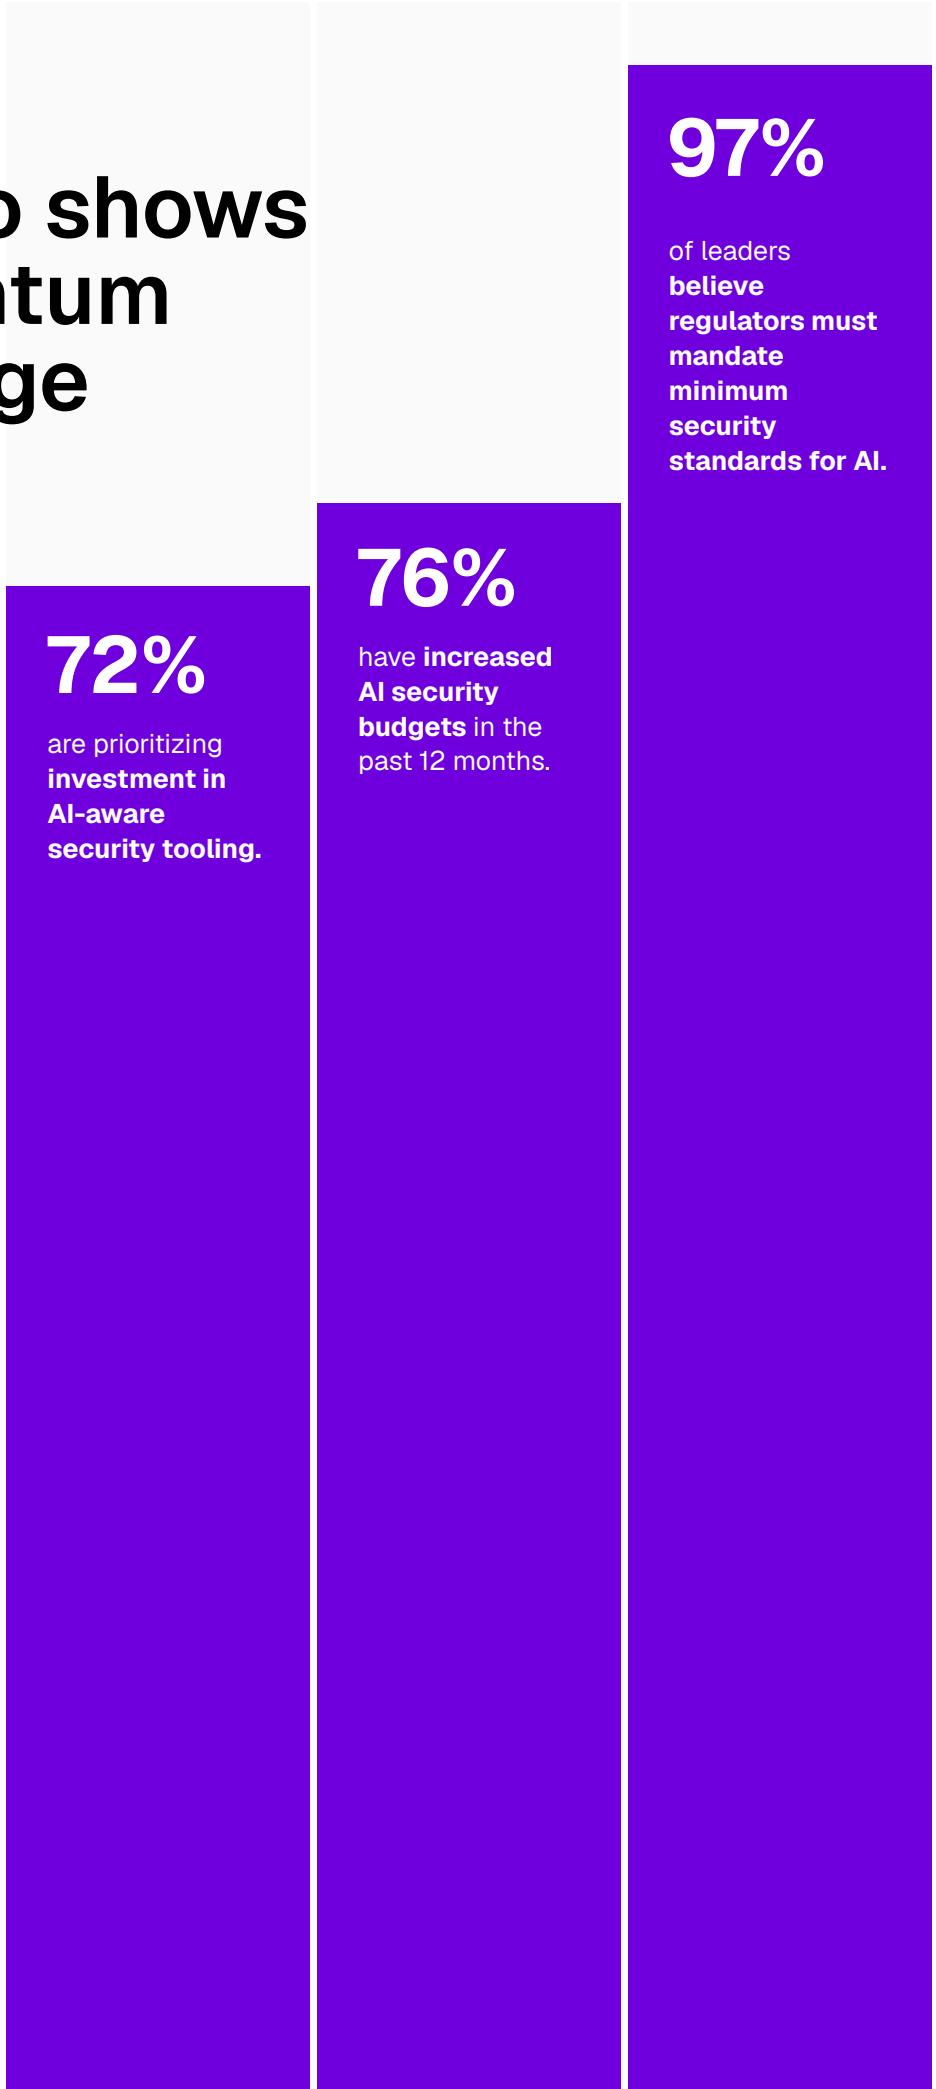

snyk

Chapter 1: The Weaponization of AI

AI-driven threats are no longer theoretical - they are active and scalable.

The threat landscape has shifted from speculative risk to active, machine-speed operations. Today, **69%** of security leaders are **concerned about AI-automated attacks**, and **52%** believe **a material AI-related incident is likely within two years**. These fears are concentrated on high-stakes vectors:

The Confidence Paradox

A dangerous "Detection Myth" has emerged: 57% of leaders believe they could detect and contain a sophisticated AI attack. However, this executive confidence is increasingly divorced from technical reality. While leaders expect to see these threats coming, they are often looking in the wrong place. They are focused on securing the "front door" of their own applications, while the "back door" is crowded with unmanaged, third-party code.

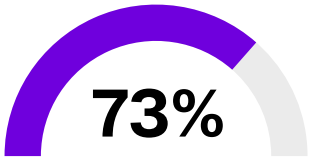

fear data exfiltration via AI agents.

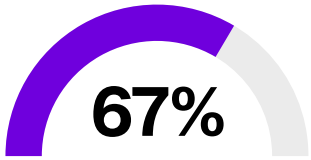

worry about targeting of critical infrastructure.

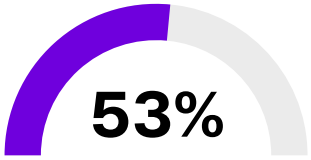

anticipate social engineering at an unprecedented scale.

Chapter 2: The Rise of the Quasi-Autonomous Enterprise

AI is already operating inside the perimeter - often faster than humans can oversee.

AI adoption has moved beyond the experimental phase and is now nearly universal across the enterprise. **Only 1% of organizations report having no meaningful AI usage**, while **87%** have actively **integrated AI into software development or core business workflows.**

Critically, AI's role has shifted from a passive advisory tool to an active operational participant. Today, **50%** of security leaders report that **AI is already functioning as a quasi-autonomous agent**, executing decisions and taking actions without direct human review. While some of this is perceived, Snyk's telemetry data identifies a significant cohort that has moved into formal technical deployment of these systems.

The Agent Footprint

According to anonymous usage telemetry from 500+ customer environments:

20% of organizations have **adopted agentic frameworks,** indicating active experimentation or deployment of autonomous AI behaviors.

of organizations have adopted agentic frameworks, indicating active experimentation or deployment of autonomous AI behaviors.

This represents the vanguard of the quasi-autonomous enterprise. Organizations are no longer just "chatting" with LLMs; they are deploying systems woven into the most sensitive layers of the enterprise. These agentic frameworks are empowered to:

- **Access internal tools and APIs** to move data between systems.
- **Interact with CI/CD pipelines** and workflow engines to influence software delivery.
- **Maintain long-lived memory** and inferred context to sustain complex operations over time.
- **Chain actions across systems**, creating multi-step automated workflows.

The Governance Gap

While technical adoption has accelerated, corporate governance has failed to keep pace. Although **90%** of organizations have established **a formal stance on SaaS AI usage**, the depth of that oversight is often superficial. **27% of leaders** admit their organization currently uses a **“rubber-stamp” approach**, allowing vendors to enable new AI features with minimal internal review.

This lack of rigorous oversight creates significant blind spots. Autonomous capabilities are being deployed into production environments without the necessary scrutiny, oversight, or technical enforcement required to manage their unique risks.

Chapter 3: The Readiness Gap: Where Defense Stalls

Governance is advancing, but operational security integration lags behind AI adoption.

To understand why the gap between AI adoption and security is widening, we must first look at how organizations progress through the AI Maturity Model. Most enterprises are currently attempting to move from simple experimentation (Stages 1 and 2) to formal governance (Stage 3), yet very few have reached the integrated, agentic security required for Stage 4.

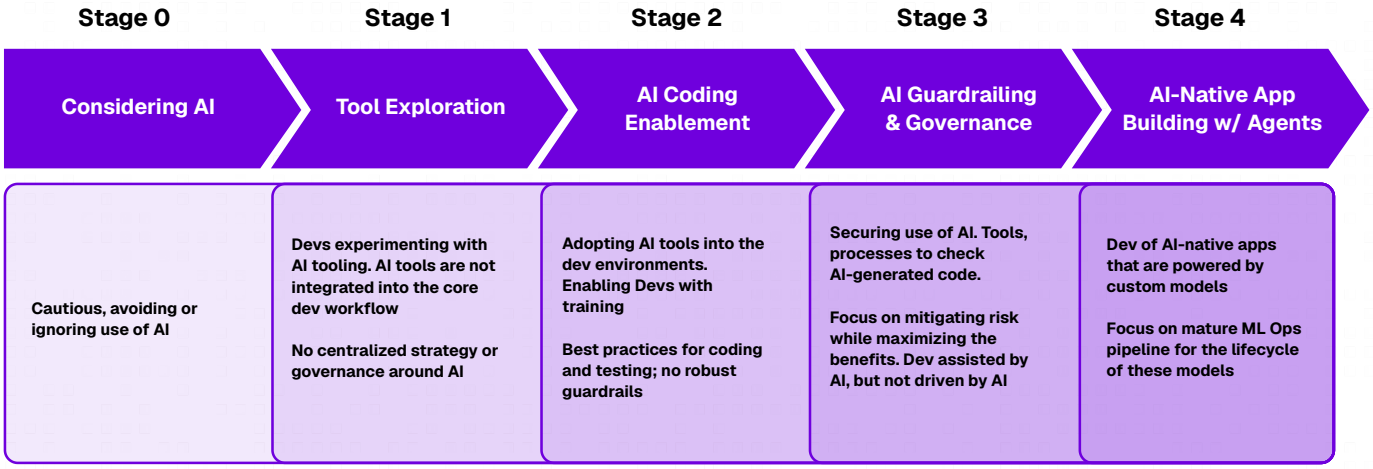

Organizations are moving quickly to establish the administrative foundations for AI security. Today, **73%** of enterprises **have formal policies governing the use of external AI tools**, and **66%** have implemented **approval processes for new AI features.**

However, these measures often function as **"paper shields"** - governance that exists on a document but lacks the technical infrastructure for continuous enforcement. This mismatch between policy and practice defines the modern Readiness Gap. While **76%** of organizations **increased their AI security budgets** over the past year, that investment has not yet translated into deep operational security.

The Enforcement Gap

According to anonymous usage telemetry from 500+ customer environments:

A strategically significant minority, about **18%** of accounts, has crossed a critical threshold into **agentic AI architectures**, deploying **Model Context Protocol (MCP) servers.**

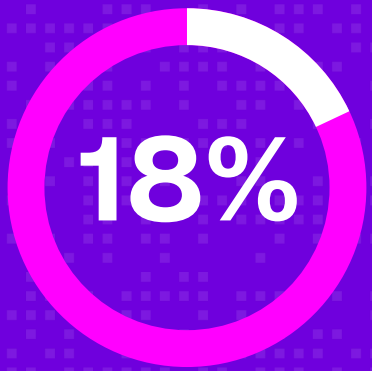

snyk

The Integration Stall

The most critical failure point is the lack of "security by design" in the AI lifecycle. Only **29%** of organizations **have fully integrated AI-aware security checks** across their design, coding, testing, and deployment phases. The deployment of MCP servers - the "USB-C of AI" - without these integrated checks creates a high-speed conduit for potential exploitation.

This stall in maturity is driven by three primary barriers:

- **Tooling Gaps: 25%** of leaders cite **a lack of AI-specific security tools** capable of inspecting agentic "chains of thought".
- **Expertise Scarcity: 23% struggle with a lack of internal expertise** to secure non-deterministic, autonomous systems.
- **Maturity Plateaus: 39%** of organizations **remain stuck at Stage 3 (securing basic AI use)**, while only **6%** have achieved **Stage 4 maturity—where security is natively embedded into automated ML operations.**

While defenders work to close these gaps, attackers are already operating at machine speed, achieving **80–90% automation in their reconnaissance and exploitation campaigns.**

Without moving from manual approvals to automated technical enforcement, the "Readiness Gap" will only continue to widen.

Chapter 4: Regulation, Investment, and the Shift to Hardening

The industry is moving from experimentation to structural defense - and demanding standards.

As AI systems transition from internal experimentation to production-grade deployments, there is an overwhelming consensus on the need for formal regulatory clarity. The vast majority of security leaders - **97% - now believe that regulators must mandate minimum security standards for AI.** This demand is characterized by extreme urgency, with **37%** stating **these standards are needed immediately,** while only **3% believe that industry self-regulation alone can provide sufficient protection.**

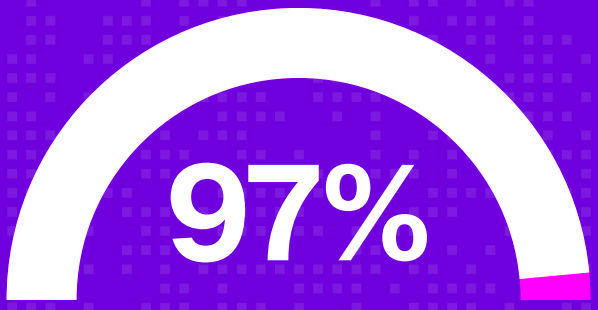

This push for regulation is fueled by a stark disconnect between high-level corporate policy and the technical ability to enforce it. While leaders call for standards, the current AI stack is increasingly built on a foundation of external code and models that existing approval processes are not equipped to handle.

Investment Priorities: Hardening the AI Lifecycle

Budgets are rapidly shifting to address these structural vulnerabilities. Last year, **76%** of organizations **increased their dedicated AI security spending**. Over the next 12 to 24 months, investment is being funneled into five primary areas to move beyond "paper" compliance and toward technical hardening:

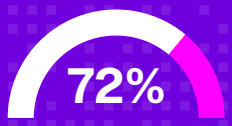

AI-Aware Security Tooling: Implementing systems that can inspect and govern non-deterministic AI traffic.

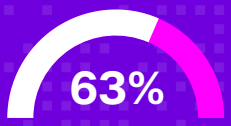

Data Protection & Governance: Securing the data used for both model training and real-time inference.

Talent and Training: Upskilling teams to manage the unique risks of autonomous agents.

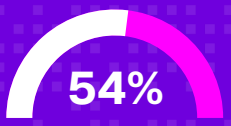

Identity and Access Management: Establishing a "central identity" for non-human actors and agents.

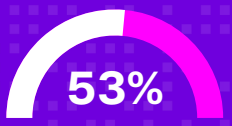

Governance and Policy: Refining the frameworks that guide responsible AI use.

While investment is rising, regulation is increasingly viewed as a necessary baseline rather than a final solution. For security leaders, the goal of these new standards is not to stifle innovation, but to create a shared regulatory floor that allows autonomous systems to scale with trust.

A Call for Shared Responsibility: Defining the Minimum Standard

The data in this report sends an unmistakable signal: the market is asking for regulatory clarity, not delay.

With **97%** of security leaders **supporting mandated minimum security standards for AI** - and **37% calling for urgency** - there is broad consensus that voluntary guidelines and best-effort governance are no longer sufficient for autonomous systems operating at machine speed.

This is not a call to slow innovation. It is a call to define the baseline expectations required for trust, scale, and safety in AI-native environments.

Security leaders are specifically calling on regulators to:

Establish minimum security requirements for AI systems that operate autonomously or access sensitive enterprise workflows.

Require transparency into AI system composition, including models, tools, and integrations.

Align accountability frameworks to include non-human actors and autonomous agents.

Harmonize standards internationally to prevent fragmentation across regions and industries.

Without a shared regulatory floor, organizations are left to interpret risk independently - creating uneven protections and systemic exposure.

Regulation alone will not secure AI systems. But clear, enforceable minimum standards can accelerate responsible adoption by providing confidence to boards, enterprises, and the public.

The opportunity before regulators and industry leaders is not to control AI—but to ensure it scales with trust.

snyk

Conclusion: The Autonomous Security Imperative

The findings in this report reveal a fundamental truth: Security architectures designed for human-paced software are structurally misaligned with agent-driven systems.
We have entered the era of the autonomous attack, where adversaries deploy AI systems capable of operating continuously, adapting dynamically, and scaling instantly.

To close the Readiness Gap, organizations must:

1. **Move beyond paper governance.** Policies and approvals cannot stop autonomous agents.
2. **Embed security into the AI and software lifecycle.** Security must operate continuously - from design through deployment.
3. **Adopt machine-speed security models.** Detection, reasoning, and response must keep pace with AI-native systems.
4. **Prepare for regulation - but do not rely on it alone.** Minimum standards will raise the floor, not define leadership.

The future of AI security will not be defined by more dashboards or manual reviews - but by systems capable of governing autonomous behavior in real time.

CEO & Board Takeaway

What Leaders Must Do in the Next 12 Months

- Assume AI agents already exist inside your organization.
- Treat AI systems as actors, not features.
- Demand visibility across models, tools, agents, and workflows.
- Invest in security that operates inside the development and runtime lifecycle.
- Align governance with enforcement—not documentation.

The defining leadership question is no longer: "Are we using: AI securely?" It is

"Can our security systems operate as autonomously and as fast as the AI we deploy?"

snyk

Methodology & Demographics

Executive Survey Data: This report is based on an online survey of 515 security leaders and specialists conducted between December 2025 and January 2026. All participants represent organizations with 100+ employees; to ensure an enterprise perspective, all respondents at the Director level or higher were drawn from firms with 500+ employees.

1. **Geography**: US (55%), Germany (9%), Japan (8%), UK (7%), France (6%), others. (Overall results weighted by each country's relative GDP.)
2. **Seniority**: Director (45%), Manager (21%), VP (4%), Chief Officer (3%).
3. **Department**: Information Security (73%), IT (27%).
4. **Industry**: Manufacturing (23%), Finance/Insurance (18%), Professional Services (14%), others.

Proprietary Telemetry Data To provide a technical baseline for these insights, this report incorporates anonymized usage telemetry data from the Snyk AI Security Platform.

- **Scope**: Data was aggregated from **500+ distinct customer environments** globally.
- **Focus**: The analysis focused on active AI agent behaviors, automated workflow patterns, and security logging configurations observed in Q4 2025.
- **Privacy**: All data is fully anonymized and aggregated; no individual customer or user-identifiable information was used in this analysis.

snyk.io