proofpoint®

2025 research report

The State of AI Security 2025

What the latest data reveals about AI security risks, investment priorities, and the biggest threats ahead.

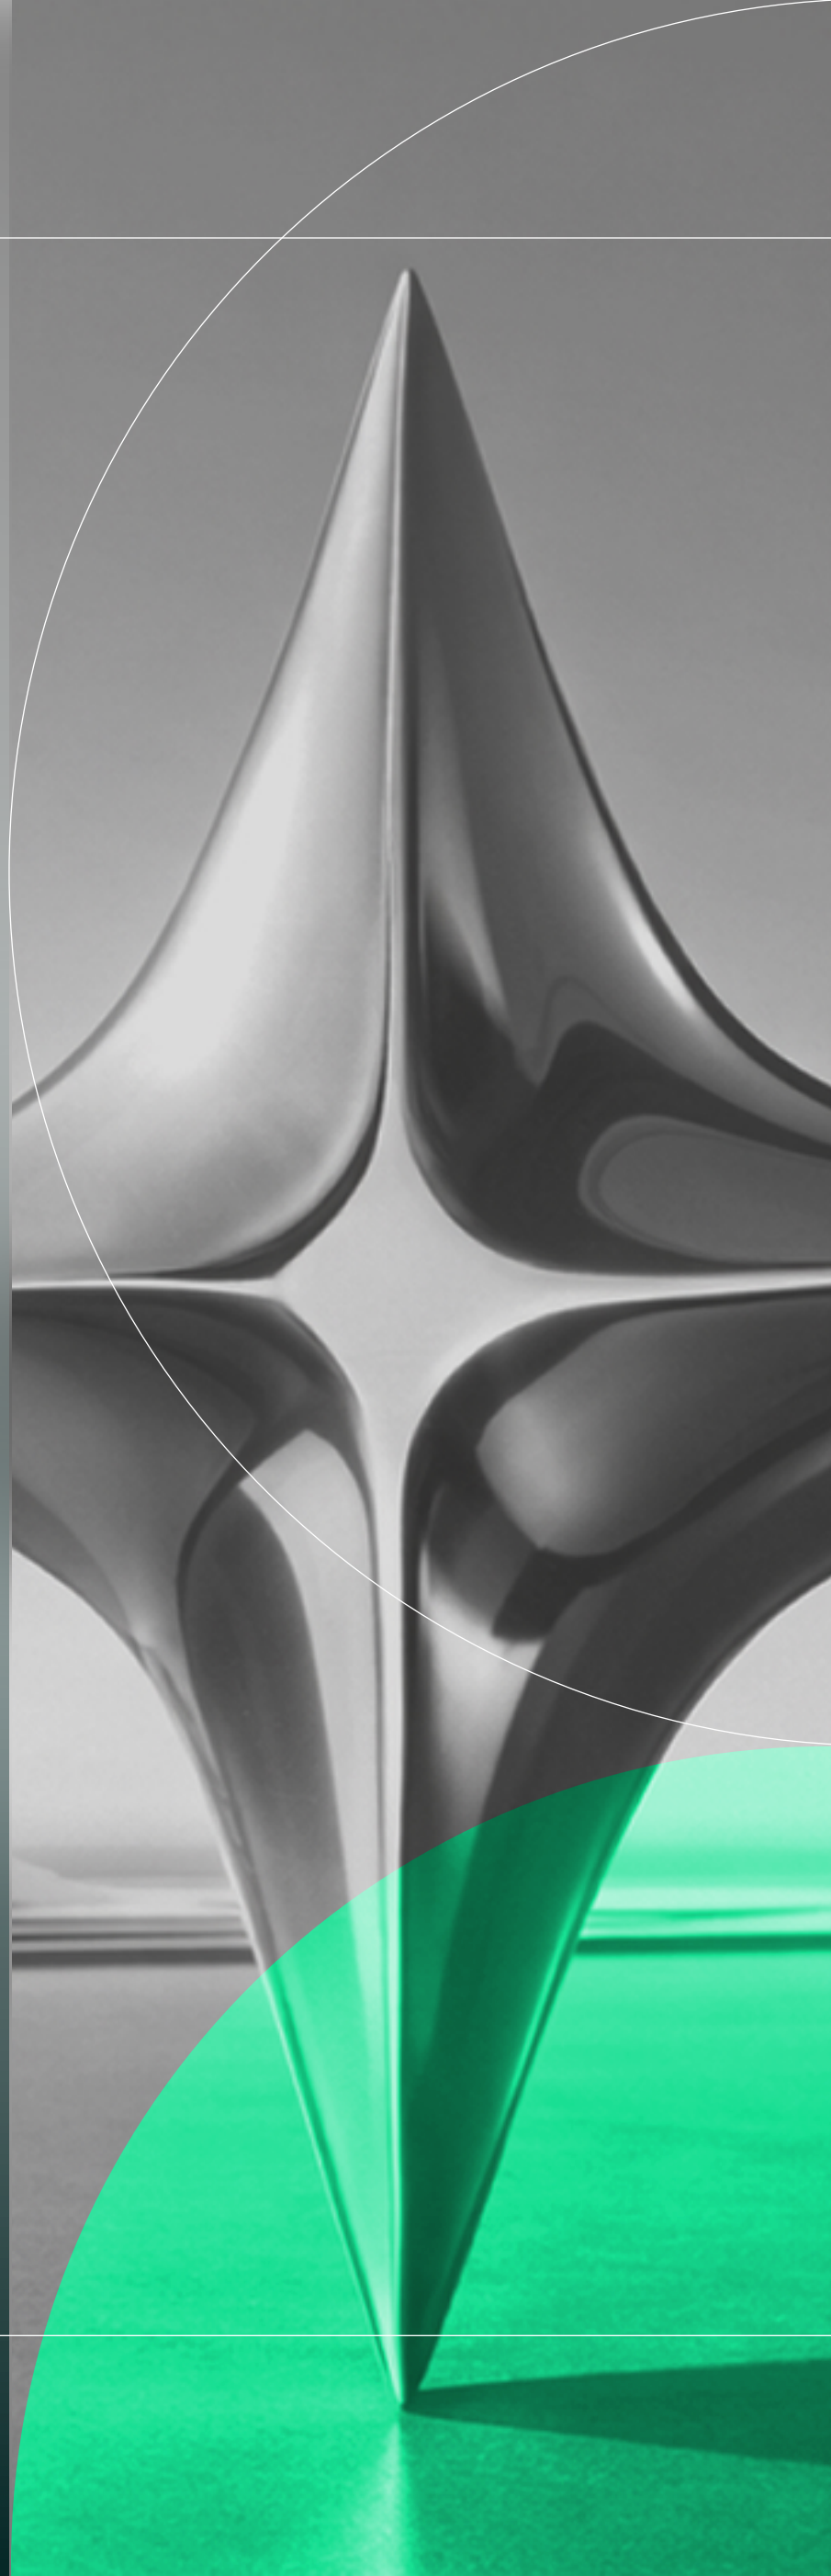

Table of **contents**

Executive summary

A letter from Proofpoint GVP, AI Security, Satyam Sinha

Satyam Sinha

GVP, AI Security, Proofpoint

The security industry has always adapted to technological shifts, but AI presents something fundamentally different. Unlike previous transitions where we could apply existing frameworks to new technologies, AI is changing the nature of risk itself.

This change is already altering the way enterprises think about security. Risks are emerging in places that traditional methods cannot fully reach. Responsibilities are spread across teams in new ways. Governance is struggling to keep pace with the speed at which AI capabilities appear in enterprise environments.

What became clear through our research is that these challenges are interconnected. The 70% of organizations lacking optimized governance aren't failing due to poor planning - they're facing a technology that dissolves the boundaries cybersecurity was designed to defend. When AI activity spans endpoints, networks, applications, and data simultaneously, traditional security models break down.

This explains why we found CIOs leading AI security decisions instead of CISOs. It also explains why organizations expect incidents they feel unprepared to address. The governance gaps, fragmented ownership, and widespread incident expectations revealed in our survey all stem from the same root cause: enterprises are trying to secure a technology that operates across every domain at once using tools designed for isolated systems.

The goal of this report is to help make sense of that reality. It documents how leaders are experiencing this fundamental shift and reveals why traditional approaches are inadequate for the risks AI creates. Rather than facing isolated security problems, the data shows enterprises are dealing with a systemic challenge that affects every aspect of how they think about protection and control.

My hope is that it provides a useful framework for understanding what's happening in your organization and supports the conversations needed to navigate this transition.

Thank you for taking the time to review our findings.

> These findings reflect what happens when transformative technology collides with security solutions designed for a different era.

Key findings

Confronting hard truths about enterprise AI security

The *2025 State of AI Security* report exposes an industry in crisis. Seventy percent of organizations lack optimized governance, while organizations simultaneously face multiple threat categories.

Half of organizations (50%) expect data breaches within 12 months. Forty-nine percent anticipate shadow AI incidents and 38% identify runtime as their most vulnerable phase. Even more concerning, 39% operate with inadequate governance structures entirely. Instead, these organizations rely on inconsistent frameworks, ad-hoc practices, or no AI-specific governance at all.

The ownership crisis compounds the problem. Chief Information Officers (CIOs) control 29% of AI security decisions. Chief Information Security Officers (CISOs) rank fourth at 14.5%, marking a departure from traditional security domains.

Almost a third (31%) are redirecting their largest security investments toward AI supply chain security over the next 12 months. This is the biggest shift in enterprise security spending priorities in decades.

These statistics show organizations expecting incidents they can't prevent. The combination of governance gaps, ownership fragmentation, runtime vulnerabilities, and shadow AI risks creates a landscape where enterprises recognize the threats but lack adequate defenses.

70%
lack optimized AI governance

50%
expect data loss via AI tools

49%
expect incidents via shadow AI

38%
say runtime is the most vulnerable

39%
lack any form of AI governance

31%
say AI supply chain is the leading investment

29%
say the CIO owns AI security

27%
say runtime is the least defended

01.
AI governance is the gap putting enterprises at risk

70%

lack optimized AI governance

How would you describe your organization's current maturity in AI security governance? (select one only)

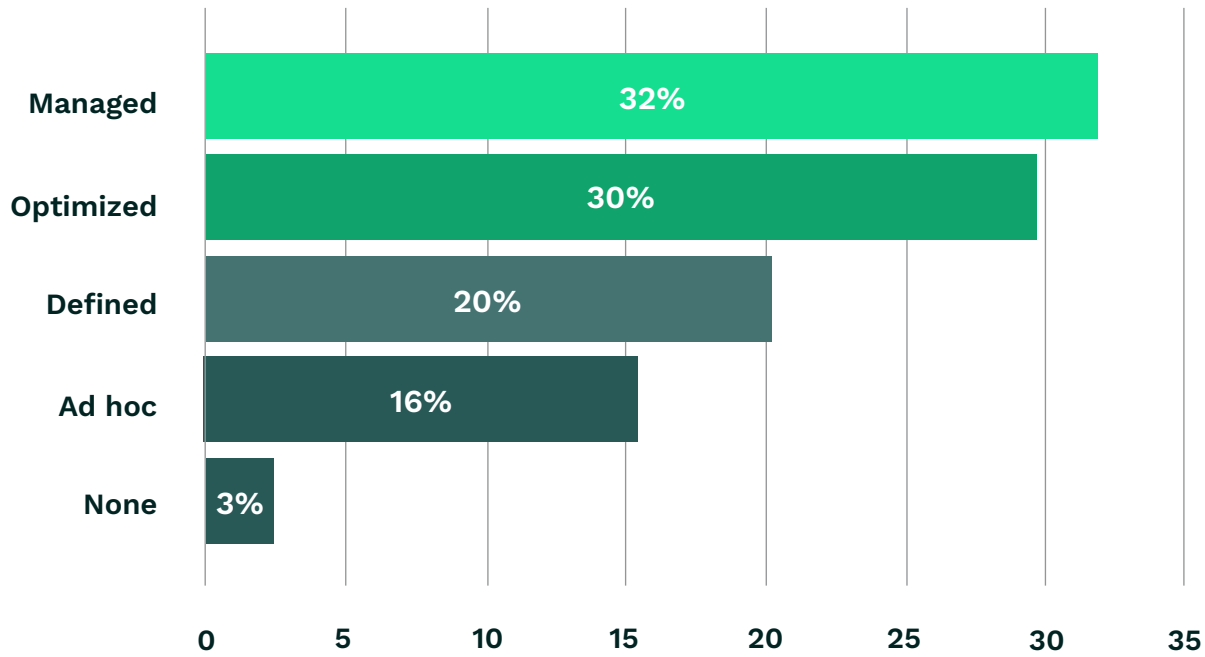

Strikingly, 70% of organizations lack optimized AI security governance, representing a massive maturity gap across the enterprise landscape. This means the vast majority of organizations cannot demonstrate the highest level of governance capabilities. These capabilities include board-level risk reviews, automated monitoring of AI use, and policies that are updated following incidents or audits.

Even more concerning, nearly 40% of organizations operate below managed and optimized governance levels. This includes those with defined frameworks, ad-hoc practices, or no AI-specific governance at all. Such organizations lack the measurement, reporting, and continuous improvement needed for effective AI risk management.

Only 32% operate at a managed level with measured effectiveness and reporting. This governance immaturity creates a disconnect with the scale of AI deployment and risk exposure organizations face. Organizations are deploying AI capabilities while simultaneously expecting significant AI-related incidents. These include data loss (50%) and shadow AI incidents (49%) within the next 12 months.

Fragmented ownership patterns shown elsewhere in the survey likely contribute to this governance maturity crisis. With AI security responsibility distributed across CIOs (29%), Chief Data Officers (17%), infrastructure teams (15%), and CISOs (14.5%), each organizational function might implement different governance approaches.

This governance maturity crisis represents a fundamental readiness gap where organizations are deploying transformative technology capabilities while lacking both unified ownership and the management structures needed to ensure secure and compliant AI use.

The window for building AI governance capabilities through deliberate strategy rather than crisis response is closing rapidly.

02.
AI adoption will drive the next data exposure crisis

50%

say AI tools will cause the next data breach

Organizations are bracing for a wave of AI-related security incidents, with three primary threats dominating their concerns. When asked to identify the most likely AI incidents to impact their organizations in the next 12 months, data leakage via GenAI tools leads at 50%, followed by shadow AI incidents at 49%, and AI-driven insider threats at 41%.

AI-driven insider threats represent a new category of security risk, with roughly 41% of organizations anticipating these. This level of concern indicates that AI tools can potentially amplify the impact of malicious insiders or create new pathways for insider-related security breaches.

Data leakage through AI tools is the leading concern because it often occurs through legitimate use rather than malicious attacks. Employees using AI tools for productivity, analysis, or content creation might inadvertently expose confidential information, intellectual property, or personal data. Unlike traditional data breaches that result from external attacks or system vulnerabilities, these incidents arise from normal operation of AI tools that employees view as helpful productivity enhancers.

Shadow AI incidents, expected by nearly half of organizations, reflect how AI adoption has outpaced formal approval processes. Employees are implementing AI tools, models, or SaaS applications without IT oversight. These actions create exposure through technologies that haven't gone through proper security review or governance protocols.

This is creating new blind spots. Organizations are losing visibility into what AI capabilities are being used and how sensitive data might be exposed

AI-driven insider threats represent a new category of security risk, with roughly 41% of organizations anticipating these. This level of concern indicates that AI tools can potentially amplify the impact of malicious insiders or create new pathways for insider-related security breaches.

The tools that organizations deploy to improve productivity and decision-making can also become vectors for unauthorized access or data misuse by individuals with legitimate system access.

Which AI-related incidents do you believe are most likely to impact your organization in the next 12 months? (select one only)

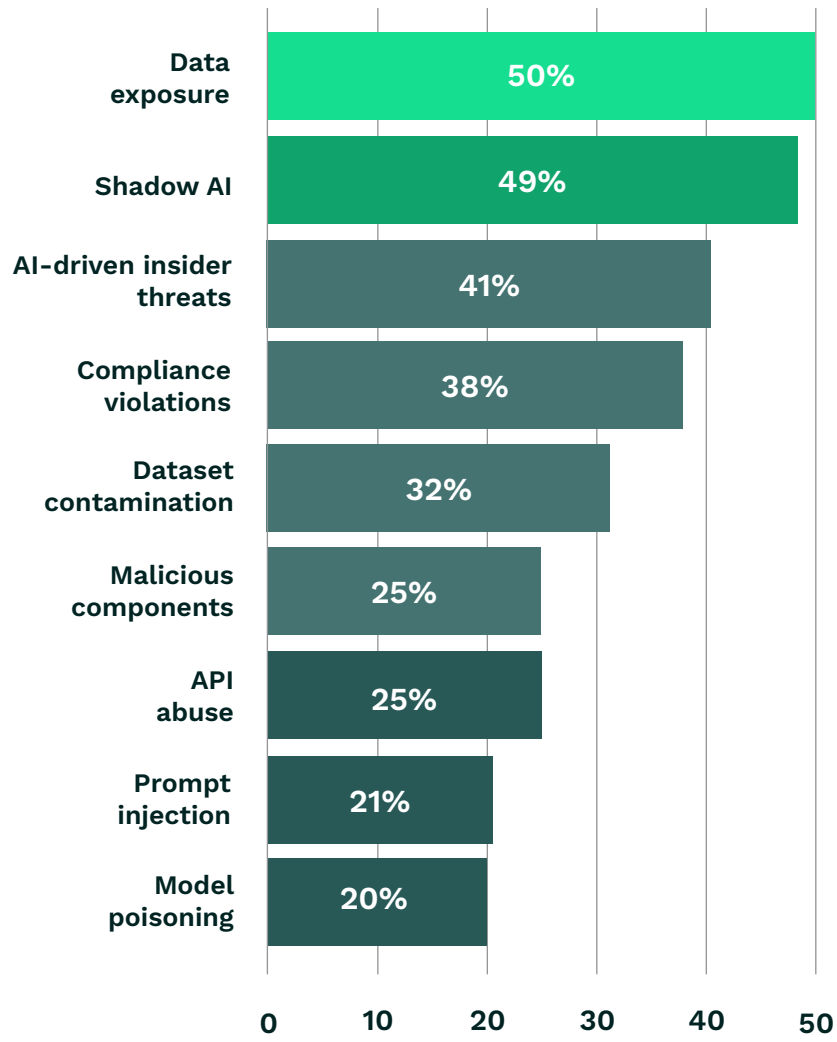

The expectation that roughly half of organizations will experience these incidents indicates that current security approaches are inadequate for the risks AI tools create.

03.
AI supply chain security: the next frontier?

Over the next 12 months, which area of AI security do you expect will require the most new investment in your organization? (select one only)

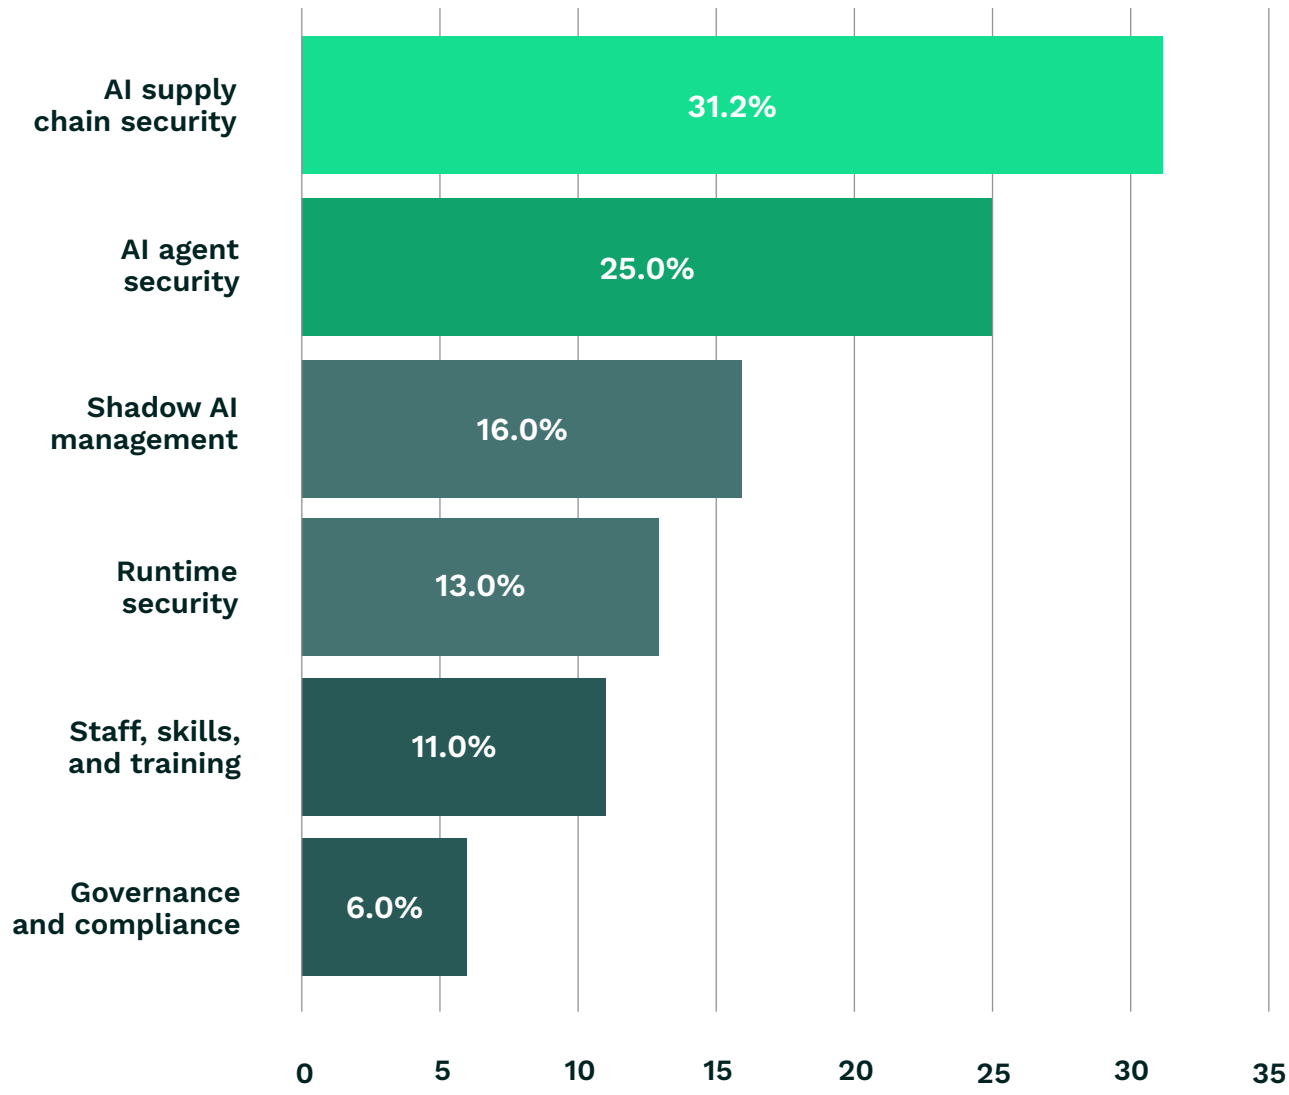

AI supply chain security has emerged as the top investment priority for organizations, with 31% planning to allocate their security budgets to this area over the next 12 months. This represents one of the biggest shifts in enterprise security spending in decades.

The shift indicates how enterprises view AI security: moving beyond isolated tool management toward comprehensive supply chain oversight. This encompasses models, datasets, agents, plugins, APIs, and SaaS AI features.

The concept of AI supply chain security remains poorly defined across the industry. Most current approaches either reference traditional software supply chain models—focusing heavily on software bills of materials and static inventories—or narrowly scope security to known vulnerabilities in AI components.

However, our survey data suggests that organizations understand AI supply chains as something entirely different.

Thirty-one percent of organizations identify data sources and embeddings as their greatest AI supply chain risk. This is followed by external APIs and SaaS-embedded AI features at 29%. Model sourcing and provenance, which receive significant attention in traditional supply chain discussions, rank as concerns for only 13%. Plugins and extensions rank third at 16%, while orchestration layers and agents receive just 9% of concern.

AI supply chain security differs from software supply chain security because it must address components that behave dynamically during runtime. Static inventories and provenance tracking provide foundational visibility. However, AI components create risks through their live interactions with data and users that cannot be fully assessed during pre-deployment phases.

Organizations recognize this runtime dimension in their security assessments. While 38% identify runtime as their most vulnerable phase, an additional 27% view risks as spanning the entire AI supply chain—from sourcing through runtime deployment. This end-to-end perspective indicates that AI supply chain security requires continuous monitoring and protection rather than primarily pre-deployment controls.

The AI supply chain includes components that traditional approaches never had to address: autonomous agents that can access multiple systems, embeddings that process and potentially retain sensitive data, and APIs that enable real-time AI capabilities across enterprise applications. These components create security implications that only become fully visible when they operate in production environments with real data and user interactions.

Which aspect of the AI supply chain do you believe poses the greatest risk to your organization? (select one only)

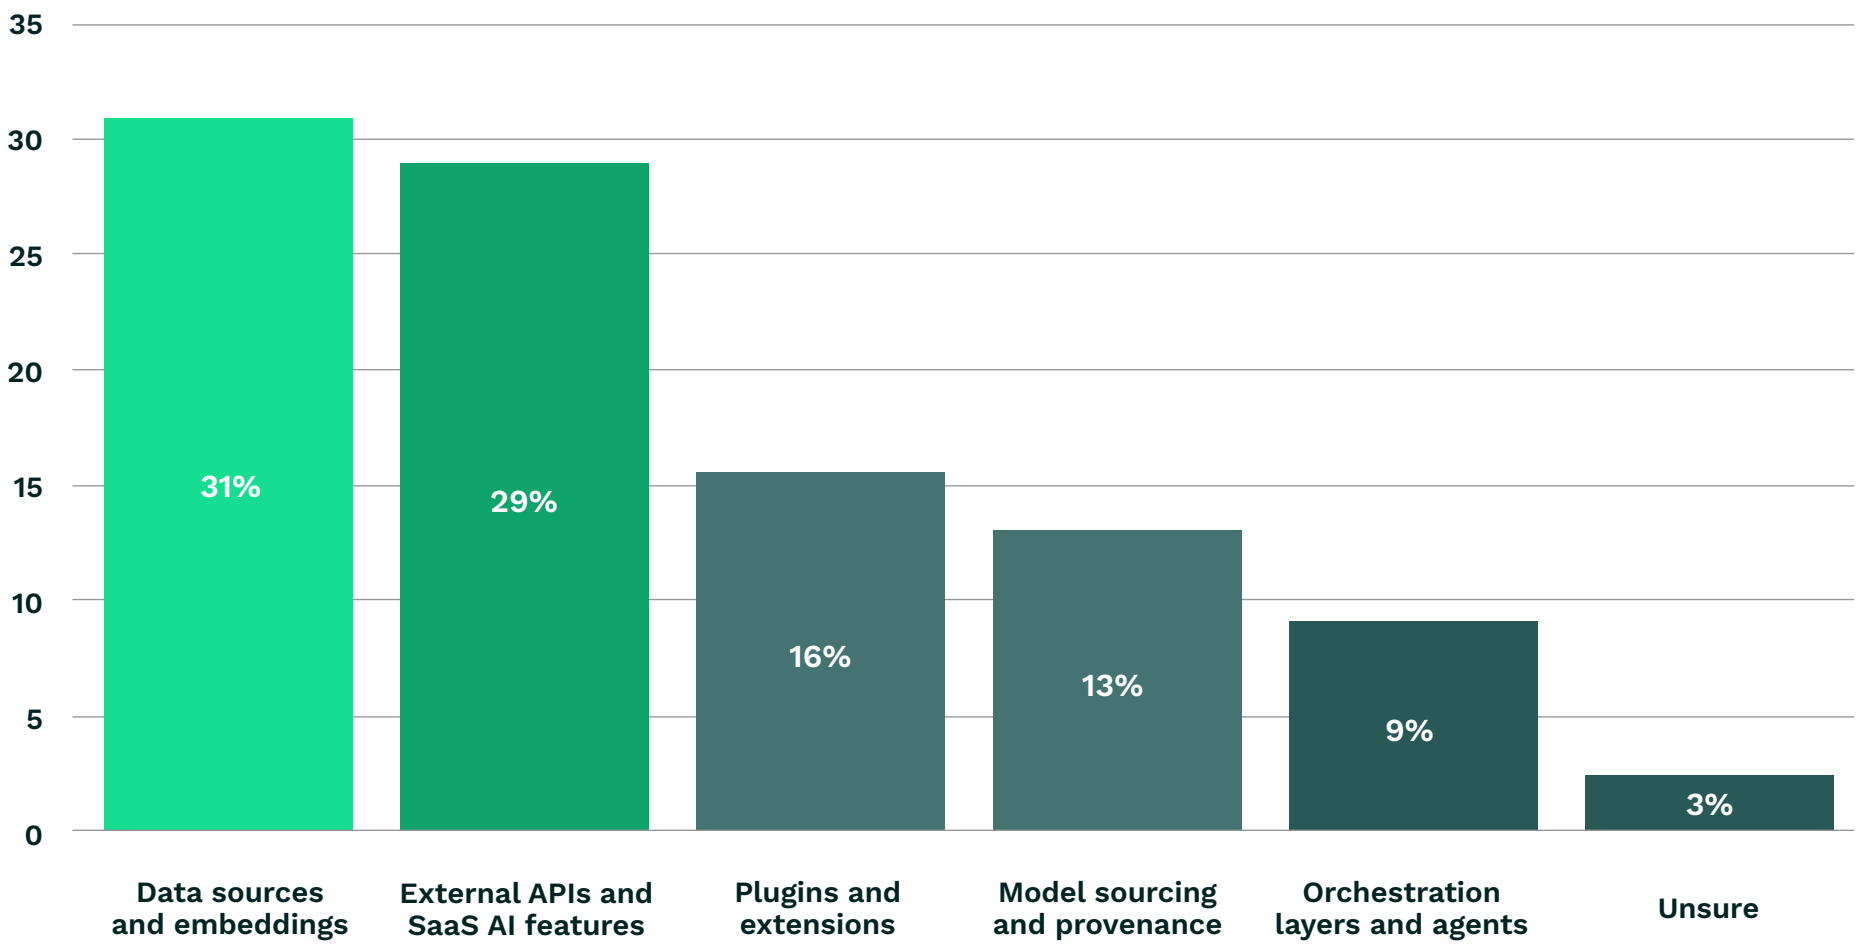

04. Responsibility for AI security is shifting

The CIO owns AI security: CISOs significantly behind

Who has primary responsibility for AI security in your organization? (select one only)

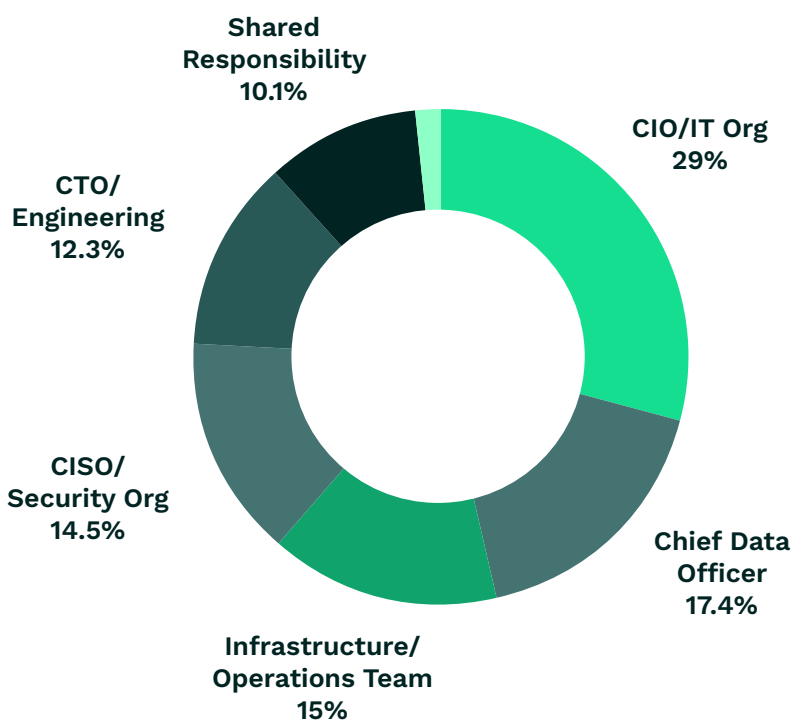

The survey reveals an unusual distribution of AI security ownership that breaks from traditional enterprise security models. The CIO and IT organization lead at 29%. This is a significant departure from standard security governance, where CISOs typically hold primary responsibility for security domains. The finding raises important questions about whether AI security is being treated as a security discipline at all, or whether organizations are approaching it as an operational technology or governance challenge.

The CISO's position at just 14.5% is particularly striking. In most enterprise security areas, from network protection to application security, CISOs either own primary responsibility or have significant oversight. The relatively low percentage in our findings suggests something unusual is happening with AI security governance. Do organizations not recognize AI as a security domain? Are CISOs declining to take ownership of AI-related risks? Or are the security implications of AI deployment being managed by other organizational functions entirely?

Chief Data Officers ranking second at 17% indicates that many organizations view AI security through a data governance lens rather than as part of a cybersecurity framework. This makes conceptual sense given that AI capabilities depend heavily on data access, quality, and handling. However, it also suggests that AI security might be divorced from the broader enterprise security strategies and risk management practices that typically exist in security organizations.

The fact that infrastructure and operations teams control AI security in 15% of organizations further reinforces the operational rather than security-focused approach. This positioning treats AI tools as infrastructure components requiring availability and performance management. However, it might not adequately address the unique security risks that AI introduces to enterprise environments.

The CIO's dominant position at 29% likely reflects their traditional ownership of IT and deployment decisions. As AI capabilities become embedded in business applications and productivity tools, CIOs may be extending their technology oversight role to include AI security by default. However, this raises questions about whether operational technology management provides adequate protection against AI-specific threats, such as prompt injection, data leakage, or model manipulation.

This ownership distribution suggests that AI security has not yet found its natural organizational home. Unlike established security domains that clearly fall under CISO purview, AI security seems to be an organizational challenge that doesn't fit existing governance structures.

Whether this is a temporary phase or a fundamental shift in how security responsibilities are allocated remains an open question.

05. Runtime: the most critical phase is the least defended

Organizations identify runtime as top concern and admit they're defenseless against it

Runtime security is the most critical vulnerability in enterprise AI deployments, with 38% of security leaders viewing runtime as their organization's most vulnerable phase of the AI lifecycle. This concern is reinforced by an additional finding that 27% of organizations identify runtime security as the area where they are least prepared to address threats. These findings make runtime the top concern in both categories.

Organizations also face preparation gaps across other critical areas. Shadow AI and regulatory compliance each rank at 23%, showing that enterprises are struggling across a broad spectrum of risks. Pre-deployment concerns fall much lower by comparison.

Only 13% cite dataset integrity and contamination, and just 12% point to model provenance and sourcing risk. This sharp drop suggests that organizations are far more concerned about what happens when AI is in use than about securing components before deployment.

Traditional shift-left approaches, which emphasize development and build-time controls, might not align with how leaders perceive the greatest risks in AI.

The shadow AI gap is especially significant given the uncontrolled adoption of tools across the enterprise. Organizations admit they lack preparation to address unapproved platforms. The result is blind spots where security teams cannot monitor or govern usage.

Compliance readiness is another major challenge. Nearly a quarter of organizations (23%) acknowledge insufficient preparation for regulatory requirements. This is the case even as new frameworks emerge and enforcement pressure increases.

Runtime security remains the focal point. It is the stage where AI actively interacts with users, systems, and data. Safeguarding that activity requires visibility, monitoring, and controls that operate in real time.

The survey results suggest that enterprises recognize this distinction and see runtime as a different security challenge from those they have faced in earlier technology cycles.

Which area of AI security do you believe your organization is least prepared to address? (select one only)

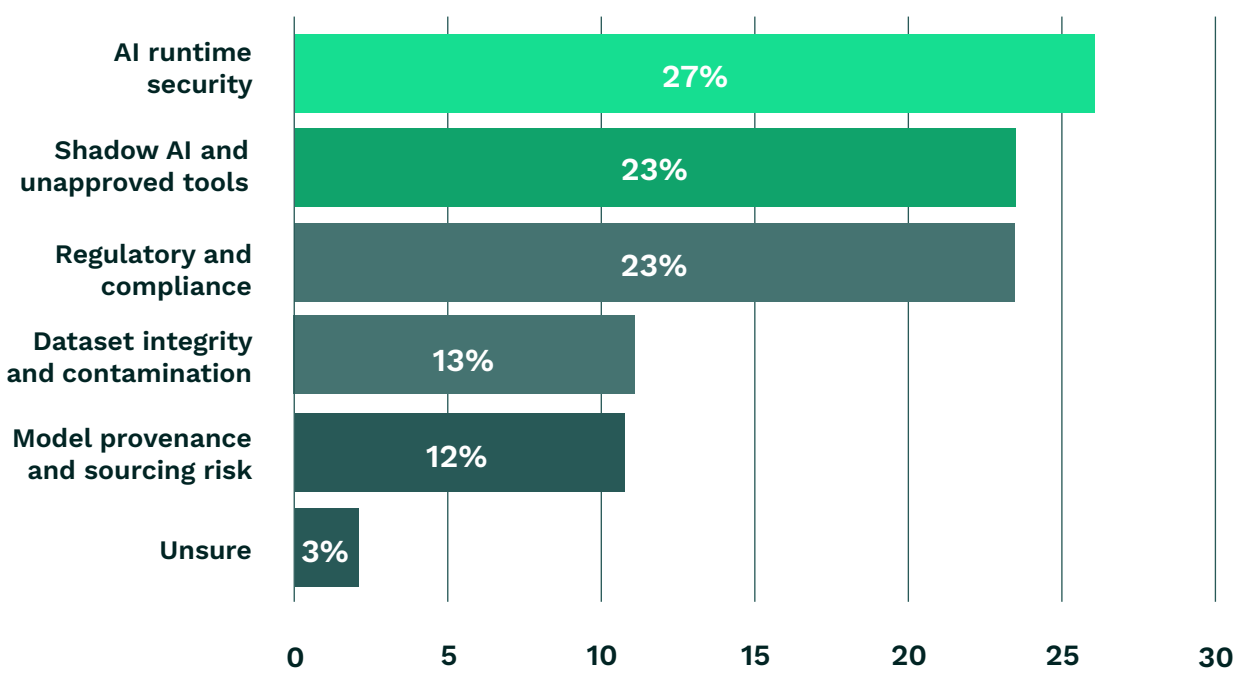

06.
Shadow AI: the fastest-growing risk enterprises cannot control

Shadow AI is the use of AI tools, applications, and services within organizations without proper IT approval, oversight, or security review. This includes employees adopting standalone AI platforms such as ChatGPT or Claude for work tasks, using AI features embedded in approved software, or deploying AI agents and automation tools without going through established security processes.

Nearly half of organizations (49%) expect shadow AI incidents in the next 12 months. This ranks shadow AI as the second-most-likely AI-related security incident they will face. Organizations also rank shadow AI as the second area where they're least prepared. Almost a quarter (23%) acknowledge inadequate preparation to address unapproved AI tools and services.

These results represent a significant evolution from traditional shadow IT because AI tools can directly process sensitive company data that employees input into them. When employees use unapproved AI tools and input confidential information, presentations, code, or business data, that information may be processed by external services, stored in ways the organization cannot control, or used to train models that other users can potentially access.

Organizations identify standalone GenAI tools adopted without IT approval as their primary shadow AI concern. Over a fifth (21%) cite platforms such as ChatGPT and Claude, as well as image generators like Midjourney as their top concern. These tools have become widely used for workplace productivity, often adopted by employees who view them as helpful assistants rather than potential security risks.

GenAI features embedded in SaaS applications present the second-highest concern at 18%. These capabilities are often enabled automatically or with minimal user configuration. This creates situations where employees might not even realize they're using AI functionality that can process sensitive company data. Applications such as Figma, Adobe, Zoom, Grammarly, and Salesforce now include AI features that analyze user inputs and generate suggestions or summaries.

The remaining shadow AI categories receive considerably less attention, with AI agents operating with user credentials at 16% and orchestration frameworks at 14%. Other vectors, such as personal accounts, third-party APIs, plugins, and local applications, all fall below 12%.

The combination of high incident expectations and low preparation levels creates a dangerous gap where organizations anticipate problems they're not equipped to handle.

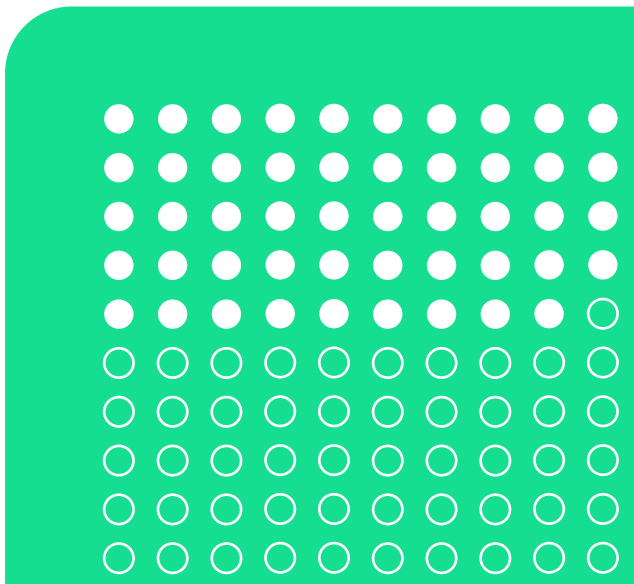

49%

expect security incidents via shadow AI in the next 12 months

Conclusion

AI security sits at the center of enterprise risk, but in ways that break traditional categories. Runtime incidents, shadow AI, and AI supply chain exposure show that threats no longer map neatly to infrastructure, applications, or users. AI operates across all domains simultaneously.

Security teams built their careers on boundaries and defined ownership, both concepts that AI eliminates entirely. Every query can leak data, every plugin creates attack vectors, and every embedded AI feature opens pathways that traditional security tools cannot protect.

CIOs control AI security while CISOs rank fourth in responsibility. Half of organizations expect data breaches they cannot prevent. As many as 70% lack adequate AI governance for what they have already deployed.

Organizations recognize the magnitude of change required. AI supply chain security emerging as the leading investment priority over the next 12 months. This represents one of the biggest shifts in enterprise security spending in decades.

The data reveals enterprises caught between the promise of AI transformation and the reality of inadequate security.

Organizations cannot address AI security by extending existing tools or adding new layers to current frameworks. The technology demands entirely new approaches to risk, governance, and operational security. These must match how AI has now spread across environments, rather than how security has worked in the past.

This transition represents both a challenge and an opportunity. Organizations that develop appropriate AI security capabilities will deploy AI with confidence and competitive advantage. Those that continue operating with fragmented ownership and reactive approaches will find themselves managing incidents that were entirely predictable. They'll watch competitors pull ahead with AI implementations they cannot safely match.

So, where do we go from here?

The answer lies in acknowledging that AI security cannot be solved with traditional tools or organizational structures. It requires building new capabilities, establishing unified governance, and creating security models designed specifically for technology that operates across all enterprise domains simultaneously. Organizations that begin this transformation now will shape the future of their industries.

Our findings reflect what happens when transformative technology collides with security solutions designed for a different era.

Those that continue operating with outdated methods will find themselves perpetually responding to crises they could have prevented.

Methodology

We conducted a 15-minute online survey with over 275 security and technology leaders in the United States.

Research included executives at enterprises ranging from 500 to over 10,000 employees across numerous industries.

Respondents held senior positions in security, IT, risk, and compliance. Roles included CIOs, CISOs, and Chief Data Officers responsible for AI security and governance decisions within their organizations.

Over

500

full-time employees

Over

$100M

annual revenue

Fieldwork conducted in **August** and **September** 2025

Robust representation from the following industries:

- Technology and Software
- Retail
- Manufacturing
- Financial Services
- Healthcare
- Public Sector

proofpoint®

About Proofpoint, Inc. Proofpoint, Inc. is a global leader in human- and agent-centric cybersecurity, securing how people, data and AI agents connect across email, cloud and collaboration tools. Proofpoint is a trusted partner to over 80 of the Fortune 100, over 10,000 large enterprises, and millions of smaller organizations in stopping threats, preventing data loss, and building resilience across people and AI workflows. Proofpoint's collaboration and data security platform helps organizations of all sizes protect and empower their people while embracing AI securely and confidently. Learn more at www.proofpoint.com.

Connect with Proofpoint:

DISCOVER THE PROOFPOINT PLATFORM →