PENTERA.

## 2026 CISO Survey

# AI SECURITY & EXPOSURE BENCHMARK

# Table of contents

# Introduction

Long before large language models, autonomous workflows, and AI-powered tooling became commonplace, security teams were already operating inside highly complex environments, managing sprawling technology stacks, responding to thousands of alerts per week, and struggling to translate visibility into meaningful prioritization. Despite continued investment in tools and controls, breaches remained common, security validation was largely periodic, and confidence often rested on assumptions rather than evidence.

Today, AI systems are being rapidly adopted across enterprises, embedded into existing applications, identity systems, cloud infrastructure, APIs, and business workflows. In many cases, this adoption has outpaced security teams' ability to maintain visibility, establish ownership, or meaningfully test how these systems behave under real adversarial conditions. Rather than introducing an entirely new category of risk, AI amplifies the same failure modes security leaders have been grappling with for years: incomplete asset awareness, overextended defenses, unclear accountability, and security validation practices that struggle to keep up with the pace of change.

AI accelerates attacker capabilities, compresses time-to-exploitation, and expands blast radius when controls fail. At the same time, it challenges traditional security assumptions, from how identities are used and permissions are granted, to how trust is established between systems and decisions are automated. In this environment, the question is no longer whether organizations have deployed enough tools or policies, but whether they can continuously prove that their defenses hold up against real-world, AI-enabled attack paths. This is where adversarial testing moves from a best practice to a necessity.

The AI Security & Exposure Benchmark 2026 examines how organizations are navigating this shift, not by looking at AI security in isolation, but by grounding it in the reality of today's enterprise security posture. Drawing on insights from 300 U.S. security leaders, the report explores how AI is reshaping the attack surface, where visibility and security validation are breaking down, and how enterprises are adapting their testing practices in response.

Designed as a benchmark for CISOs, this report provides a peer-based view into how organizations of similar scale are approaching AI adoption, security ownership, testing maturity, and investment. It enables security leaders to assess where their own programs stand relative to their peers, identify gaps between intent and execution, and understand which practices are emerging as indicators of stronger security validation in an AI-driven threat landscape.
For any feedback or inquiries, please contact: noam.hirsch@pentera.io
Wishing you a meaningful read,

**The Pentera Market Research Team**

# Methodology

**300**

Pentera commissioned a global survey of 300 U.S. enterprises, comprising CISOs and senior security executives. The data collection was conducted by Global Surveyz, an independent research firm, in December 2025.

All respondents held C-level or VP-level roles within IT and cybersecurity functions. Organizations represented employ 3,000 or more employees and span a wide range of industries.

Participants were recruited through a global B2B research panel and invited via email to complete the survey.

The average time to complete the survey was approximately 8 minutes.

To minimize order bias, answer choices for most non-numerical questions were randomized. In questions where multiple selections were allowed, percentages may total more than 100%.
In certain cases, findings are reported only for respondents who provided a definitive answer (excluding "prefer not to disclose").

The results reflect self-reported responses and represent a point-in-time snapshot of how enterprises are approaching AI adoption, security investment, and adversarial validation in an evolving threat landscape.

# Executive summary

## 01

### Enterprises Are Spending Big to Protect Their Data and Assets

U.S. enterprises spend an average of $2.48M annually on cybersecurity, with approximately 12% ($300K) allocated to pentesting. In 2026, 66% of CISOs plan to increase overall security budgets and 70% plan to increase pentesting spend, positioning security and validation as an ongoing enterprise priority rather than a reactive investment.

**$2.48M**

**Average annual cybersecurity budget**

**$300K**

**Annual pentesting budget**

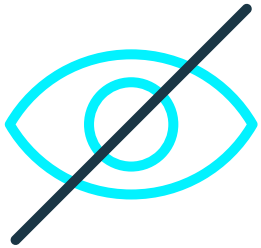

67%

**of CISOs have limited visibility into AI systems**

## 02

### AI Adoption Is Widespread, But Visibility Remains Limited

AI is now in use across 100% of enterprises, yet governance has not kept pace with deployment. 67% of CISOs report limited visibility into where and how AI is operating across their environments, while even the 33% who report good visibility still expect the presence of Shadow AI.

## 03

### AI Security Is Funded, but Not Yet Treated as a Standalone Priority

AI security is still in an early stage of maturity, following a trajectory similar to cloud security in its early years. Today, 78% of enterprises fund AI security through existing security budgets, while only 1% report a dedicated AI security budget. Over the next 12 months, 21% expect to introduce a specific AI security budget, signaling the next phase of AI security maturity.

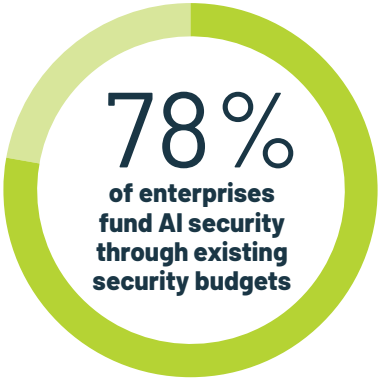

## 04

### AI Security Challenges Are Foundational, Not Budget-Driven

The primary challenges in securing AI systems are foundational, centered on skills, visibility, and tooling. Lack of internal expertise is cited by 50% of CISOs, while limited visibility into AI usage (48%) and the absence of dedicated AI security tools (36%) follow closely behind. By comparison, only 17% cite budget constraints as a primary challenge.

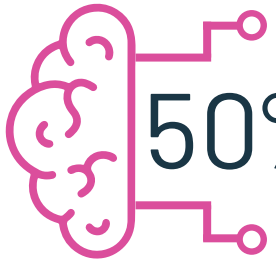

**50%** **of CISOs cite a lack of internal expertise and limited visibility into AI usage**

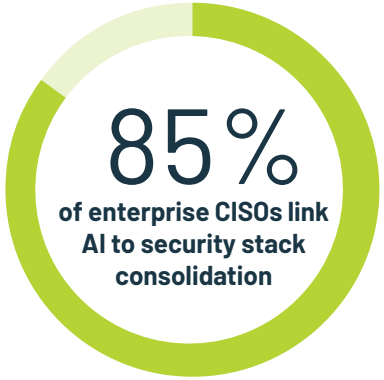

**85%** **of enterprise CISOs link AI to security stack consolidation**

## 05

### AI Is Already Reshaping Consolidation Conversations, But Not Security Stacks Yet

There is a clear gap between consolidation intent and execution. While 85% of CISOs say AI is influencing their security stack consolidation strategy, only 3% are actively consolidating due to AI capabilities today, with another 11% consolidating for reasons unrelated to AI.

## 06

### At least 75% of U.S. Enterprises Dealt With Attackers Inside Their Environment in the Past 24 Months

Enterprise security teams manage already complex environments, with nearly 40% of organizations operating 51 or more security solutions and 68% adding net-new tools over the past 12 months. Despite these investments, at least 75% of enterprises report that attackers gained unauthorized access to their environment in the past 24 months, reinforcing that expanding security stacks alone do not prevent attacker access.

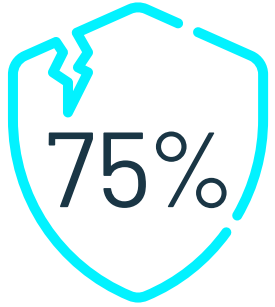

**75%** **of enterprise CISOs have dealt with an attacker in their environment in the past 2 years**

# Survey report findings

## AI adoption is ubiquitous, but deployments aren't mature

AI usage has reached ubiquity across the enterprise, with every CISO reporting some level of adoption. However, the scope and maturity of that adoption vary significantly.

**Only 18% of CISOs report deploying AI broadly across the enterprise, while the majority (82%) remain in selective, team-level, or pilot-stage deployments.**

Adoption depth increases with organizational size, with the largest enterprises are two to three times more likely to have AI operating at enterprise scale.

Among enterprises with 3,000–4,999 employees, just 12% report enterprise-wide AI deployment. That figure rises to 16% in enterprises with 5,000–9,999 employees, and reaches 31% among those with 10,000 or more employees.

This pattern challenges the conventional assumption that large enterprises are slower to adopt new technologies. In the case of AI, scale appears to accelerate adoption rather than hinder it, with larger enterprises moving more quickly to integrate AI across business functions.

**How would you describe the overall maturity of AI adoption across your organization?**

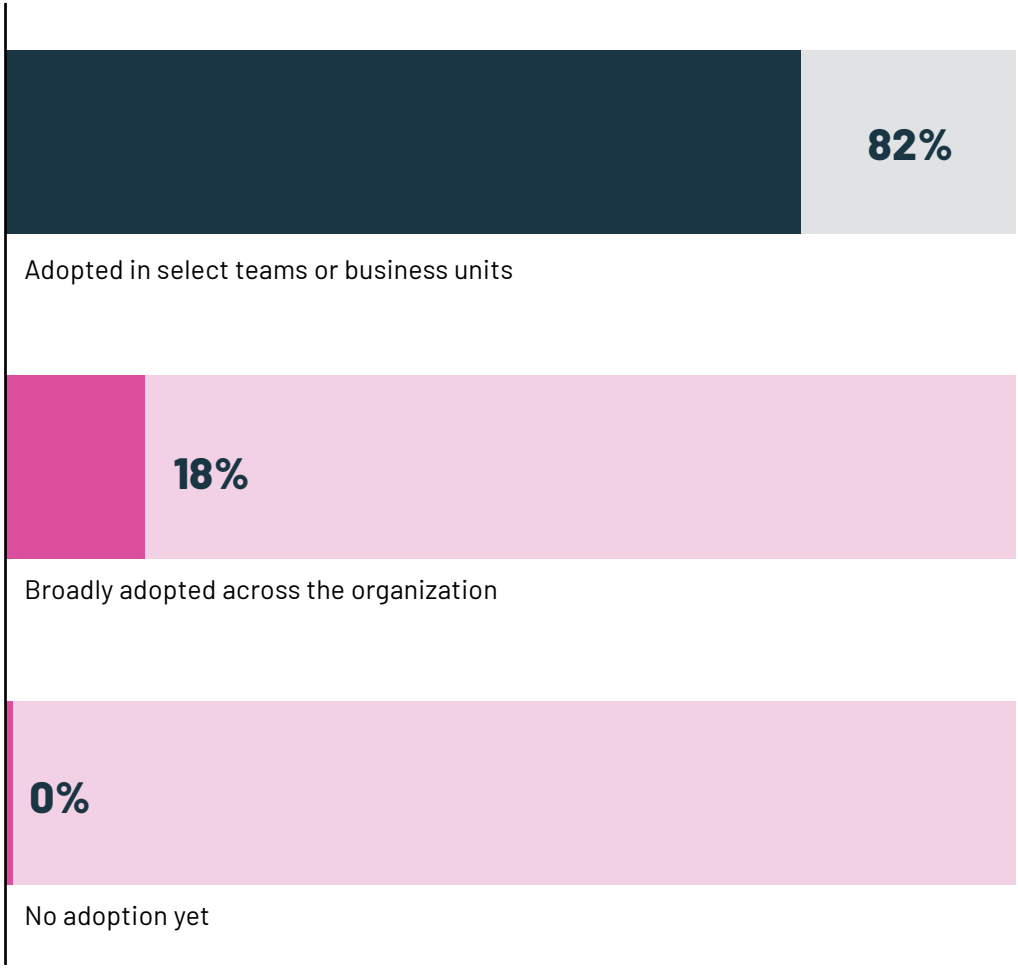

## CISOs are being challenged by the growing Shadow of AI

Visibility into AI usage is incomplete across all enterprises surveyed. **No CISO reports full visibility into AI systems without the presence of Shadow AI.** Instead, 66% report limited visibility, while the remaining 33% say they have good visibility but still expect unauthorized or unmanaged AI use within their environment. 1% admit to having no visibility at all.

This indicates that Shadow AI is a known and anticipated condition for the enterprise, not an edge case. Even CISOs who believe they have a strong understanding of how AI is being used across their enterprises acknowledge that some level of unmanaged or unsanctioned AI activity persists.

Greater enterprise scale improves visibility, but does not eliminate this issue. Among enterprises with 3,000–4,999 employees, 78% of CISOs report limited visibility into AI usage. That figure declines to 57% among enterprises with 10,000 or more employees. This improvement likely reflects the more mature asset management, discovery, and governance capabilities typically found in larger enterprises, even as overall security complexity continues to grow with scale.

This lack of visibility extends beyond identifying where AI is deployed to understanding what it can access.

**Do you have visibility into how many AI systems your organization is actually using?**

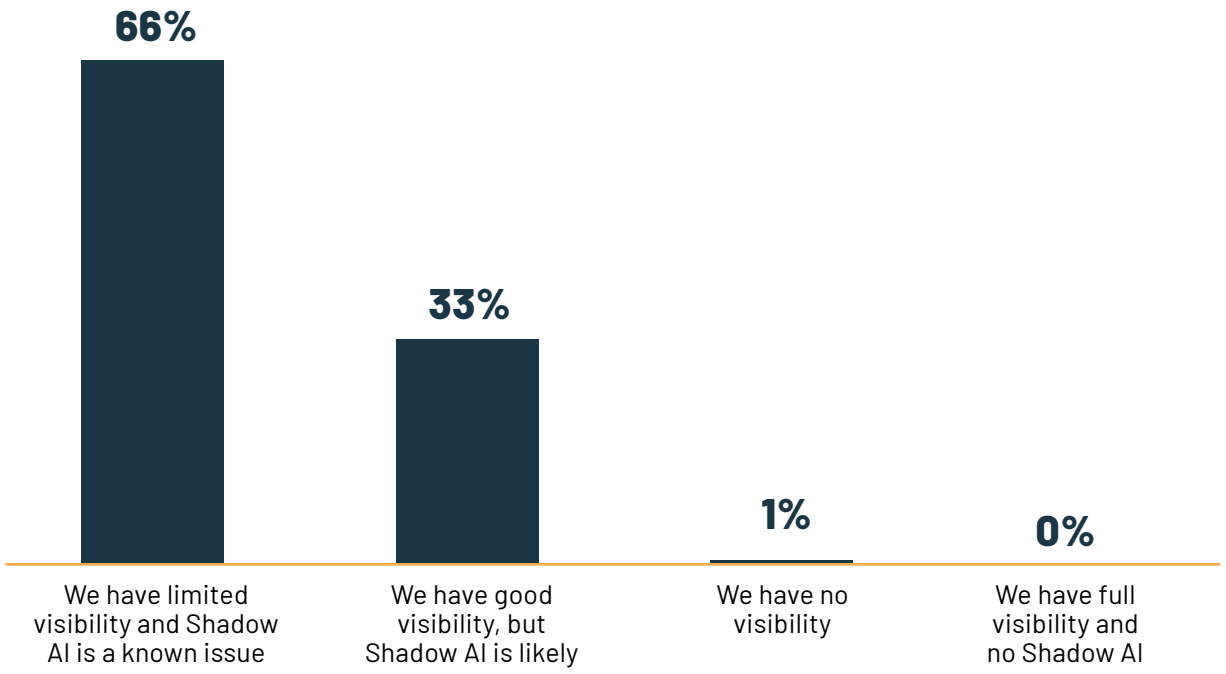


According to *SailPoint's AI agents: The new attack surface report*, only 54% of professionals report having full awareness of the data their AI agents can access, highlighting how data-level exposure often remains opaque even when AI systems themselves are known.

As AI adoption expands across teams, functions, and workflows, enterprise CISOs are operating with the understanding that complete visibility is not yet achievable. The result is an AI attack surface that is growing faster than it can be fully inventoried, governed, or consistently validated from a security perspective.

## AI security relies on legacy controls, not specialized tooling

Enterprise CISOs report that AI security is most often being addressed through existing security controls, rather than tools purpose-built for AI environments.

75% of CISOs report that their enterprises rely on traditional endpoint, cloud, application, or API security tools to protect their AI systems, extending controls originally designed for other attack surfaces to cover AI-driven workflows and infrastructure. Dedicated AI security tooling remains uncommon. **Only 11% of enterprise CISOs report having security tools specifically designed to protect AI systems.**

While AI-specific security controls remain more rare, this reality can change quite quickly. 64% of CISOs already report that their teams are actively evaluating options for AI-specific security controls, and that number will likely grow as the tools available mature.

This dynamic mirrors earlier platform shifts, where widespread adoption outpaced the security controls designed to protect them. Early enterprise adoption of Linux-based infrastructure and later cloud and container environments all introduced periods where organizations relied on adapted legacy controls before purpose-built security capabilities matured.

**Do you have specific technical controls/security tools in place to defend your AI ecosystem?**

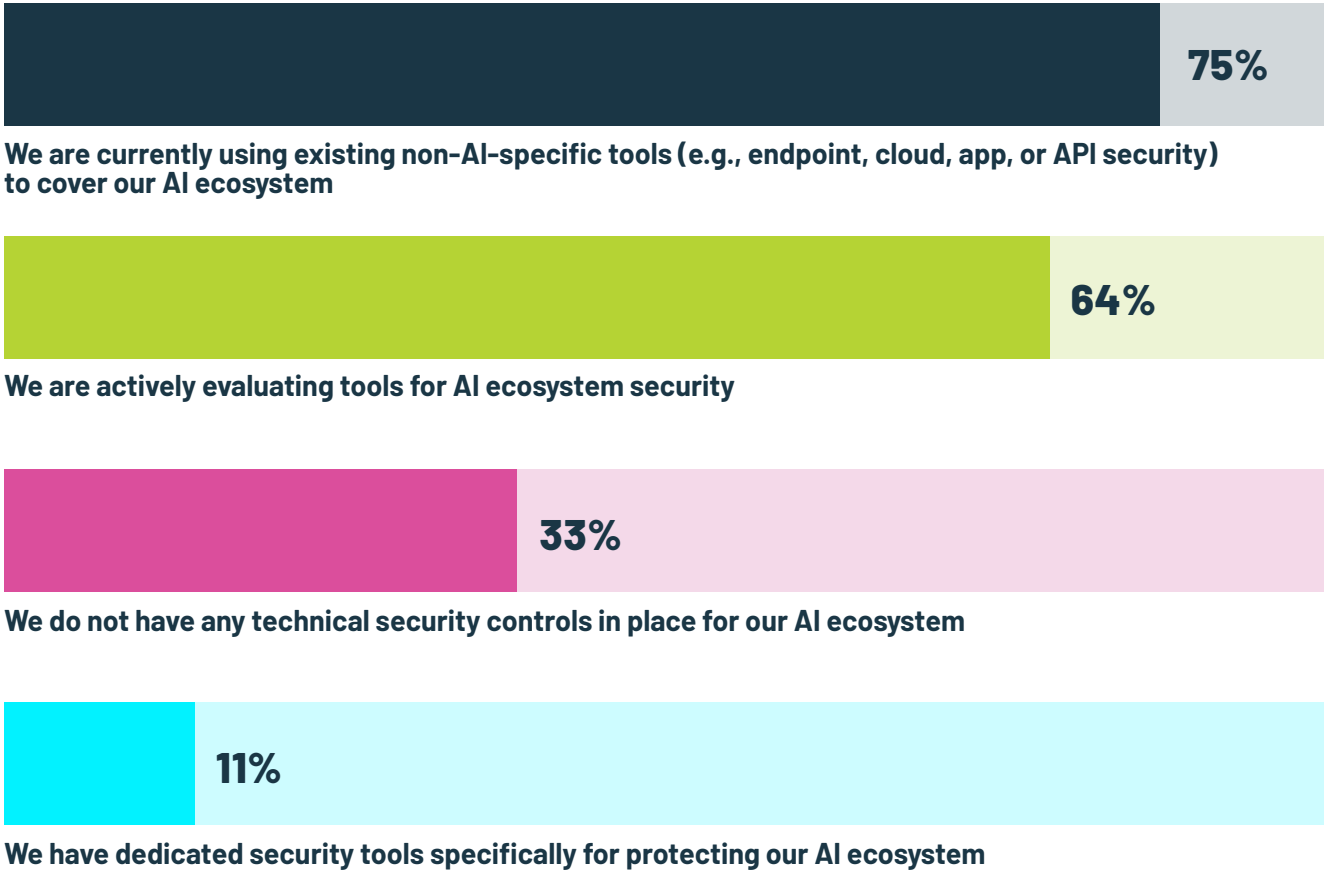


✱ Question allowed more than one answer and as a result, percentages will add up to more than 100%

## AI security budgets are embedded today, but are beginning to take shape

AI security spend already exists across enterprise environments, but it is rarely isolated as a standalone budget line. 78% of enterprises fund AI security within broader categories such as application, cloud, or general security budgets. Only 1% of enterprises report having a dedicated AI security budget today, indicating that formal allocation remains the exception.

At the same time, the absence of dedicated AI security budget lines reflects the current state of AI security implementation. As shown earlier in this report, most enterprises are securing AI systems using existing security controls rather than tools purpose-built for AI, leaving little need for standalone budget allocation today.

Looking ahead, this may begin to change. 64% of CISOs are actively evaluating AI-specific security tools and 21% of CISOs expect to introduce a specific AI security budget within the next 12 months, suggesting early movement toward more formalized funding as AI adoption expands and security capabilities mature.

This split reflects the current stage of AI security maturity in enterprise environments. Investment is already occurring, but it remains embedded within existing budget structures rather than established as a distinct domain. This is characteristic of early-stage adoption rather than mature deployment.

Enterprise CISOs recognize the risk and are funding AI security accordingly, but dedicated budgets, clearly settled ownership, and standardized operating models have not yet fully emerged.

**Do you have a budget specifically for the cybersecurity of your AI attack surface?**

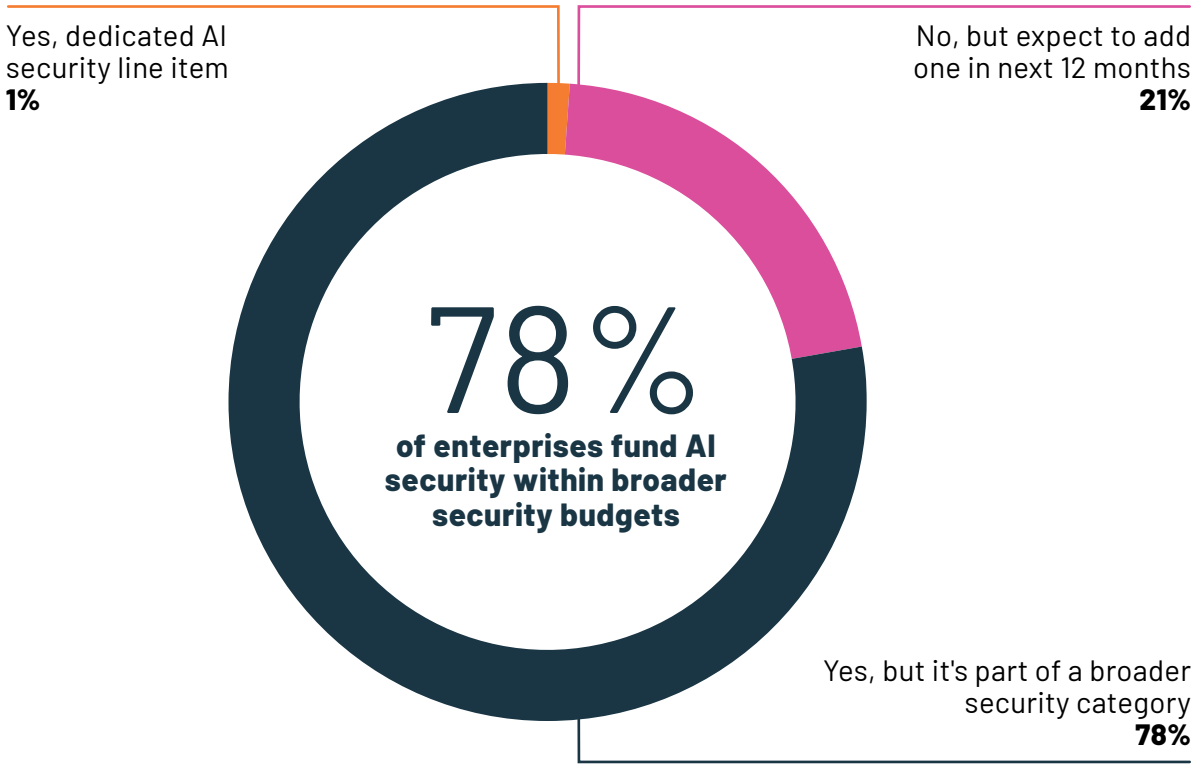

## AI security ownership remains decentralized

Responsibility for securing AI environments is most often shared across the enterprise. **56% of enterprise CISOs report that AI security is owned across multiple teams as a shared responsibility, rather than being assigned to a single team.** This shared-ownership model spans security, IT, infrastructure, and application teams, reflecting how AI systems are embedded across different parts of the enterprise.

Clear, single-team ownership is less common, and even the decision of which team varies widely. 20% of enterprises place AI security solely within the security team, while 16% assign it to IT or infrastructure, 2% to application or R&D teams, and 6% rely on third-party providers. For the majority of enterprises, no single team has end-to-end accountability for securing the AI ecosystem.

This distributed model introduces operational risk. When multiple teams own different aspects of AI security, responsibilities are more likely to be fragmented, increasing the chance that gaps emerge between domains. Unlike security areas with well-established ownership models, AI security often spans identities, data, applications, and infrastructure, making coordination and accountability more difficult to enforce consistently.

### Who owns the security of your AI ecosystem within your organization?

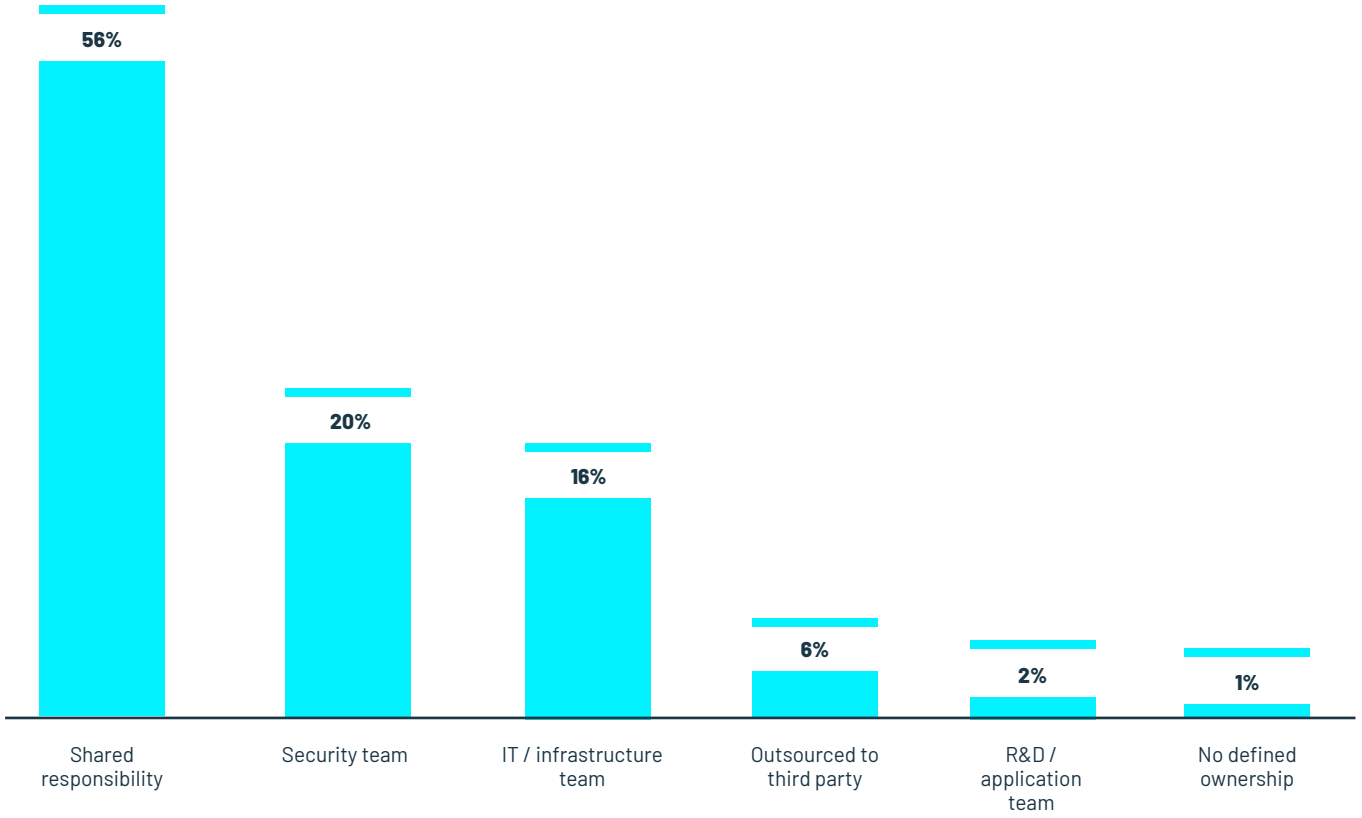


✱ Question allowed more than one answer and as a result, percentages will add up to more than 100%

## AI security is keeping pace for some, but lags for many

As we learned earlier, 100% of CISOs report some level of AI adoption, ranging from testing and limited use cases to broader deployments across the enterprise. This widespread adoption raises an immediate question for CISOs: is your AI being secured to the same standard as the rest of the enterprise environmen?

There are encouraging signs of early alignment between AI adoption and security posture. A majority of enterprise CISOs (55%) report that the security of their AI environments is at least on par with the rest of their IT infrastructure, suggesting that many enterprises are extending existing security standards and practices to AI systems rather than treating them as an afterthought.

At the same time, a substantial gap remains. **44% of CISOs acknowledge that AI security currently lags behind existing IT security controls**, a meaningful portion given the speed at which AI is being introduced into enterprise workflows. This split reflects an environment where AI adoption is moving quickly, but security maturity is uneven across enterprises.

**How does the security of your AI infrastructure compare to the rest of your IT environment?**

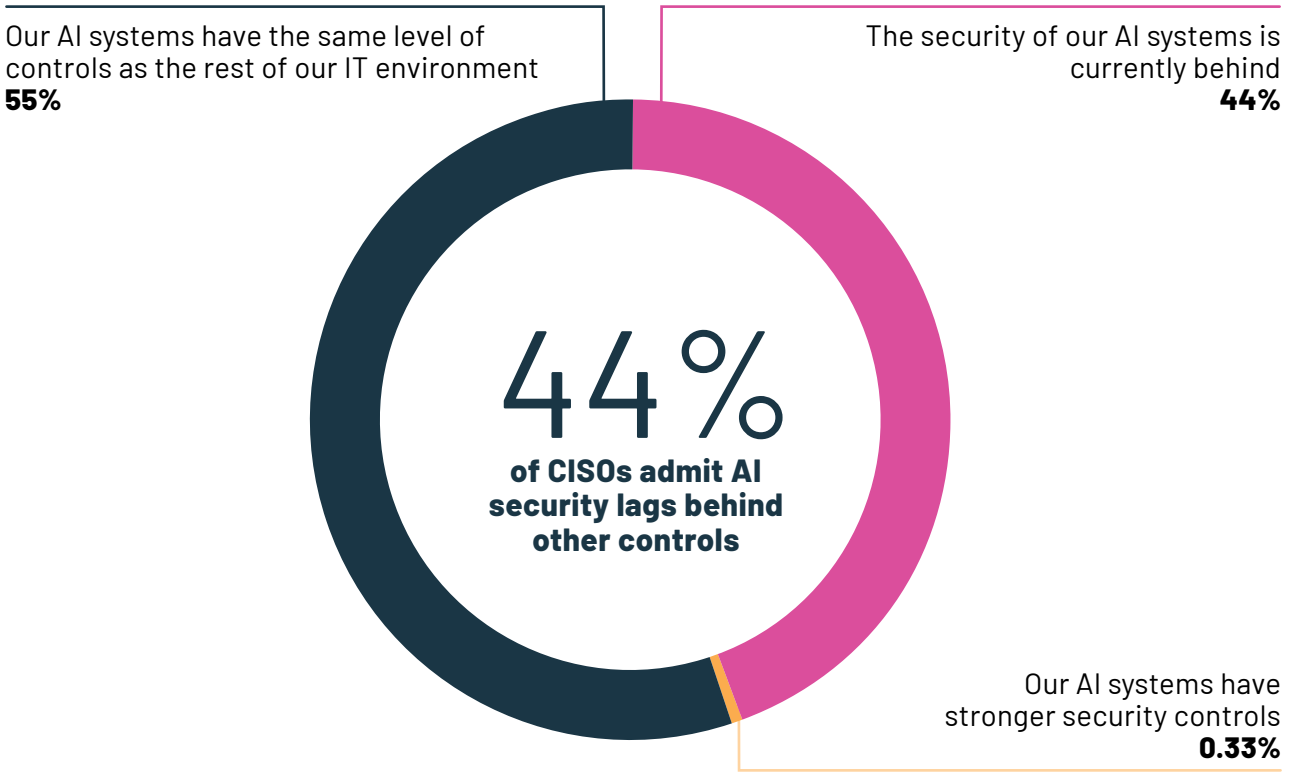


The data shows that while many enterprises have succeeded in bringing AI into their existing security framework, a large segment is still working to close the gap, again highlighting the transitional state of AI security across the enterprise landscape.

## CISO confidence in securing their AI reflects an ongoing transition

Enterprise CISOs express confidence in their ability to defend AI systems, but that confidence is measured rather than absolute. When asked how they would rate their confidence when reporting to their CEO, the majority of CISOs (78%) describe themselves as somewhat confident, explicitly acknowledging that while steps are being taken to address AI security, gaps remain.

Not all CISOs share this even this measured level of confidence. **One out of every five CISOs (21%) expresses low confidence in defending their AI ecosystem, citing gaps in their current defensive capabilities.** Very high confidence is rare: only 1% of CISOs say they are very confident, and none report that AI security is not a concern.

Interestingly, confidence in security correlates closely with how frequently organizations validate their security posture. **Among enterprises that conduct quarterly penetration testing, 80% report being somewhat or very confident in their AI security. That figure declines to 71% among those testing annually. This gradient suggests that confidence is reinforced through regular security validation rather than assumption.**

**If you were reporting to your CEO tomorrow, how would you describe your confidence defending your AI ecosystem from cyberattacks?**

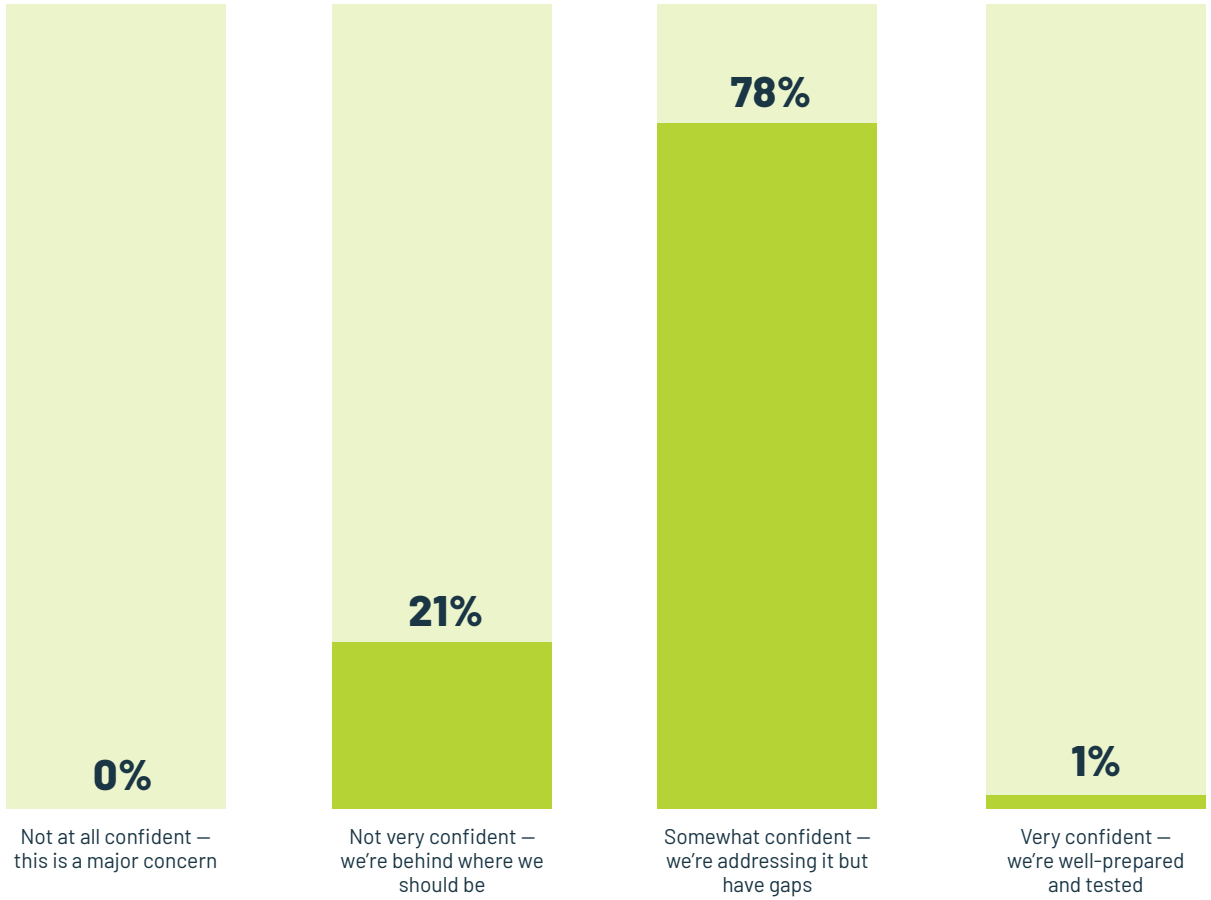

## Skills gap is defining barrier for AI Security

**Half of enterprise CISOs cite a lack of internal expertise (50%) as their biggest barrier to securing their AI attack surface**. This finding aligns with the 2025 ISC² Cybersecurity Workforce Study. Their study identified AI as the most pressing skill needed among cybersecurity professionals, cited by 41% of the 16,029 practitioners surveyed.

Limited visibility into AI usage emerges as a close second. Almost half of CISOs (48%) report incomplete visibility into where and how AI is being used across their environments, complicating efforts to define ownership, apply consistent controls, and validate exposure. Security teams cannot protect what they cannot see, and this challenge is compounded by tooling gaps. 36% of CISOs report a lack of dedicated AI security tools, limiting their ability to discover, monitor, and consistently secure AI systems as they proliferate across teams and workflows.

By comparison, only 17% of CISOs cite budget constraints as a primary barrier to securing AI systems, indicating that resourcing is not the dominant challenge. Instead, enterprises are grappling with the foundational work of understanding, governing, and securing AI systems that are already embedded across their environments.

**What are your 2 biggest barriers to securing your AI ecosystem today? (Select top 2)**

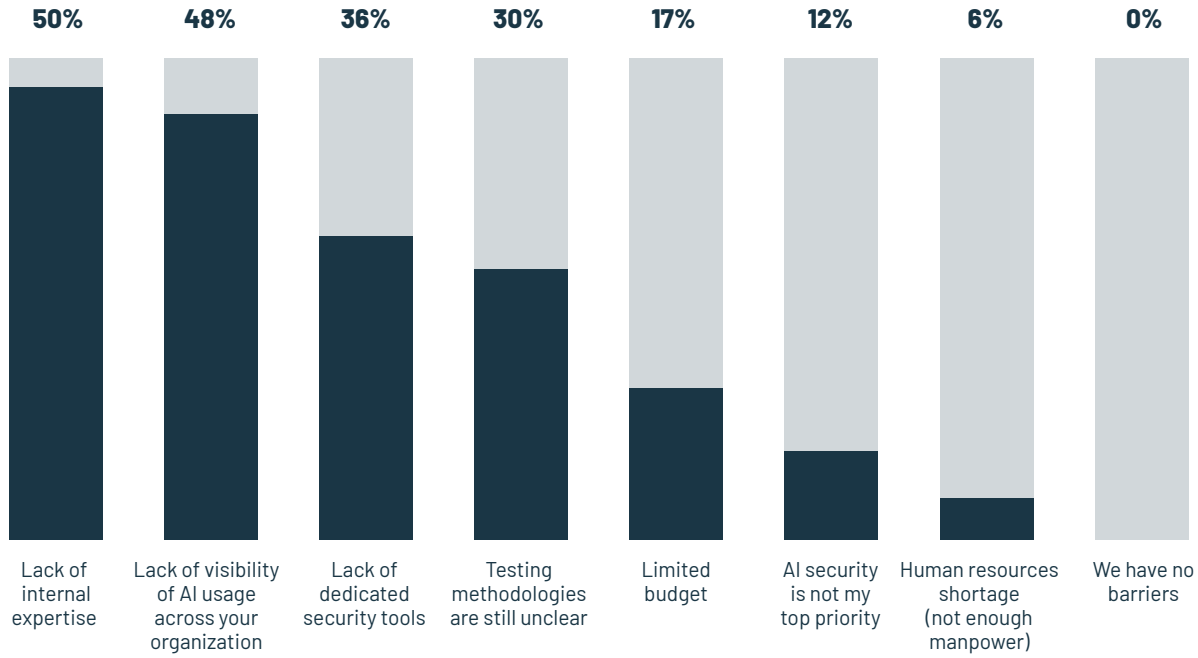


Unclear testing approaches represent a separate challenge for many enterprises. **Nearly one-third of CISOs (30%) report that methodologies for testing AI security remain unclear. As we saw previously in this report, CISO confidence increases with more frequent adversarial testing, suggesting that unclear testing approaches can limit consistent validation as AI adoption expands.**

Importantly, this is not a question of priority. Very few CISOs indicate that AI security is not a concern, underscoring that enterprises recognize the importance of securing AI. The challenge lies not in awareness, but in translating that awareness into mature capabilities as AI adoption continues to expand.

✱ Question allowed more than one answer and as a result, percentages will add up to more than 100%



https://www.isc2.org/Insights/2025/12/2025-ISC2-Cybersecurity-Workforce-Study

## CISOs actions reflect CTEM mentality

The data shows a clear correlation between where enterprise CISOs perceive risk and where they direct pentesting. Identity and access controls, cloud infrastructure, web-facing assets, and APIs all rank among the most concerning attack surfaces and are also the most frequently tested. Their consistent ordering across both views reinforces that testing focus closely tracks perceived exposure, with CISOs concentrating validation on the attack surfaces they believe are most likely to be targeted.

Testing scope, however, extends beyond top-ranked concerns alone. Internal networks, endpoints, and applications, while cited less frequently as primary areas of concern, are still included in testing at high rates. This reflects an understanding that while risk perception informs where testing is prioritized, effective validation must also account for how compromise can propagate across interconnected systems once access is established.

The AI ecosystem fits into this same picture. AI is tested more frequently than it is cited as a top area of concern, indicating that even before AI emerges as a primary risk driver, CISOs are already incorporating it into adversarial testing scope. The findings reflect a shared operating mindset among enterprise CISOs: risk is distributed across the environment, not confined to a single surface. Testing is prioritized around perceived entry points, but applied broadly enough to validate exposure across the full attack path.

### During a penetration test, where do you direct the pentesters to target/test?

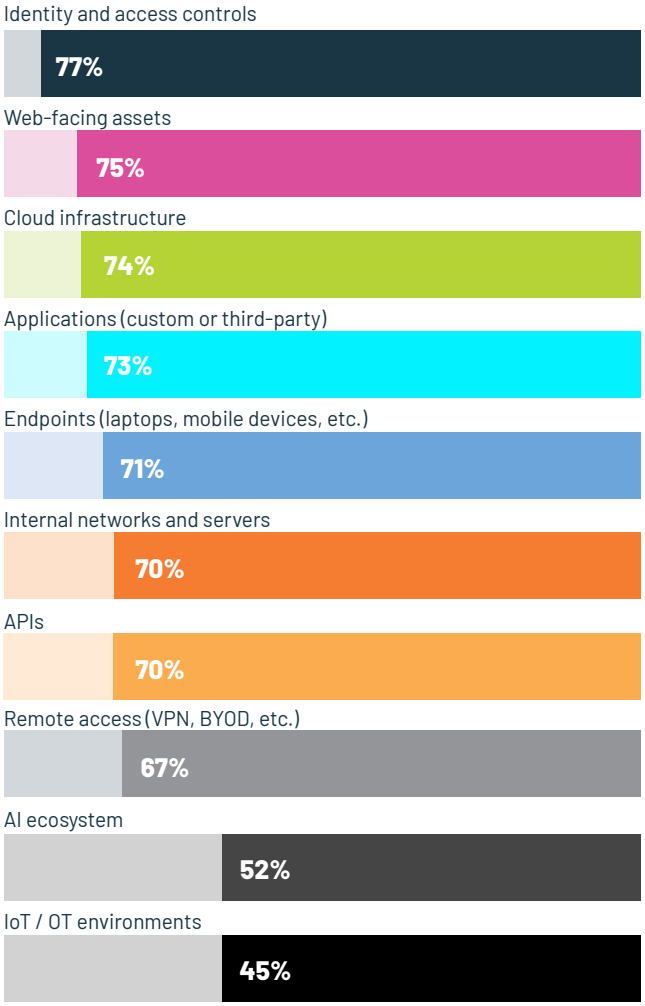


### Which top 3 attack surfaces concern you the most?

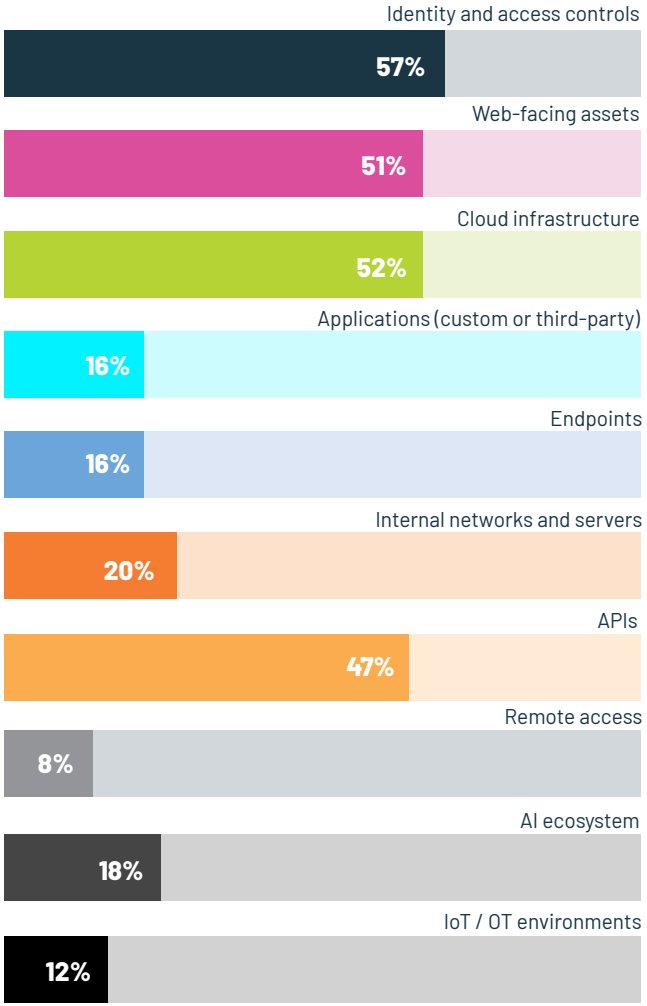


**This approach aligns closely with the principles of Continuous Threat Exposure Management (CTEM). Rather than relying solely on static assumptions, enterprises are using adversarial testing to continuously validate where exposure exists in practice, narrowing the gap between perceived risk and observed behavior over time.**



* Question allowed more than one answer and as a result, percentages will add up to more than 100%

## AI security testing expectations are still forming

Requirements for AI security testing are not yet uniform across insurers. Among enterprises surveyed, 44% of CISOs report that their cyber insurance providers require proof of pentesting of their AI ecosystem, while 52% report that no such requirement exists. The remaining respondents indicate uncertainty.

This near-even distribution highlights a gap between AI adoption and external enforcement. While cyber insurance providers actively influence security tooling and controls, expectations for validating AI security through testing are not yet applied with the same consistency.

At the same time, AI systems are already being incorporated into enterprise testing programs. As shown elsewhere in this report, **52% of enterprises include AI-related scenarios within their adversarial testing scope**, even as external requirements for AI security testing remain uneven. The prevalence of AI testing therefore slightly exceeds the prevalence of insurer-mandated AI pentesting.

**Does your cyber insurance provider require proof of pentesting of your AI ecosystem?**

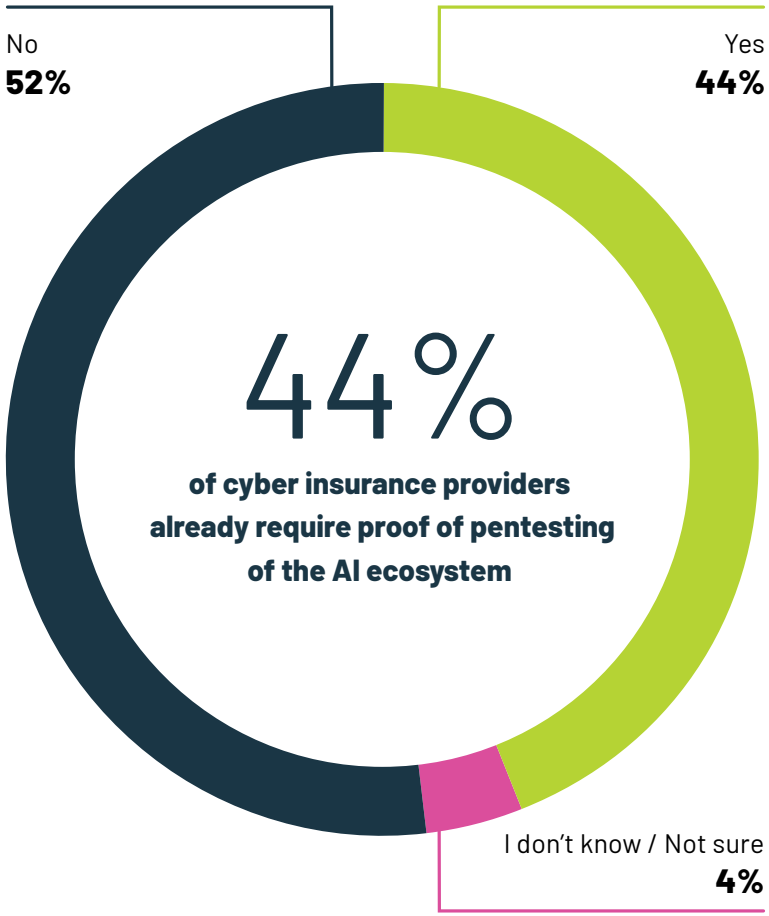

## How enterprises are applying AI security testing today

AI-related threat scenarios are being tested at broadly similar rates across enterprise environments, indicating that AI security validation is already underway across multiple dimensions of risk. Rather than concentrating testing on a narrowly defined set of AI threats, enterprises are incorporating a range of AI-enabled and AI-native scenarios into adversarial testing programs.

This pattern suggests that CISOs are not waiting for a clearly defined AI threat hierarchy or standardized risk model before acting. In the absence of established guidance, AI-related testing is being incorporated alongside existing validation efforts, extending adversarial testing to cover how AI may be abused through familiar attack techniques as well as how AI systems may introduce new behaviors and exposure paths.

In practice, testing is being applied across scenarios that reflect both continuity and change. AI-enabled attacks that resemble known techniques, such as phishing, impersonation, or misconfigured services, are tested alongside AI-native misuse scenarios such as prompt injection and model manipulation. The relatively even distribution across these scenarios indicates that AI security testing is being treated as part of the broader enterprise attack surface, rather than deferred until clearer prioritization frameworks emerge.

**Which of the following AI-related threat scenarios are being tested as part of your red teaming or offensive security program?**

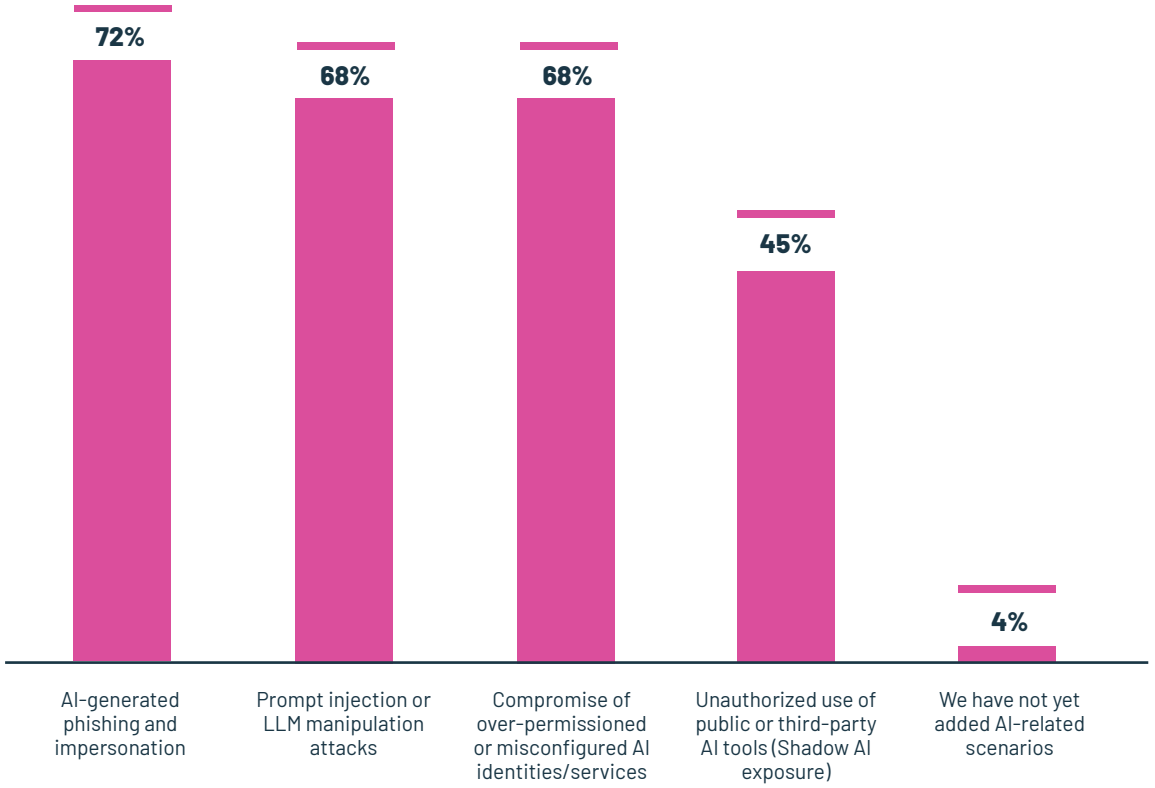


✱ Question allowed more than one answer and as a result, percentages will add up to more than 100%

## Enterprise security is already defined by scale and complexity

Before AI's full impact on security infrastructure is realized, enterprise security environments are already defined by scale.

Across the enterprises surveyed, CISOs manage an average of 47 security solutions across their IT environments. This level of tool density is not an outlier. **40% of CISOs already operate security stacks with 51 or more solutions, reflecting how layered modern security architectures have become.**
Security stack size also scales predictably with enterprise size:

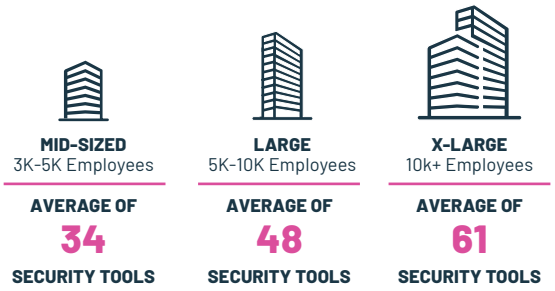


While large security stacks are often built to improve coverage, growing complexity can introduce its own operational risks. As tool counts increase, security teams are required to manage more integrations, alerts, configurations, and data sources, often generating significant noise that makes it harder to identify meaningful signals.

**How many security solutions do you currently use across your organization?**

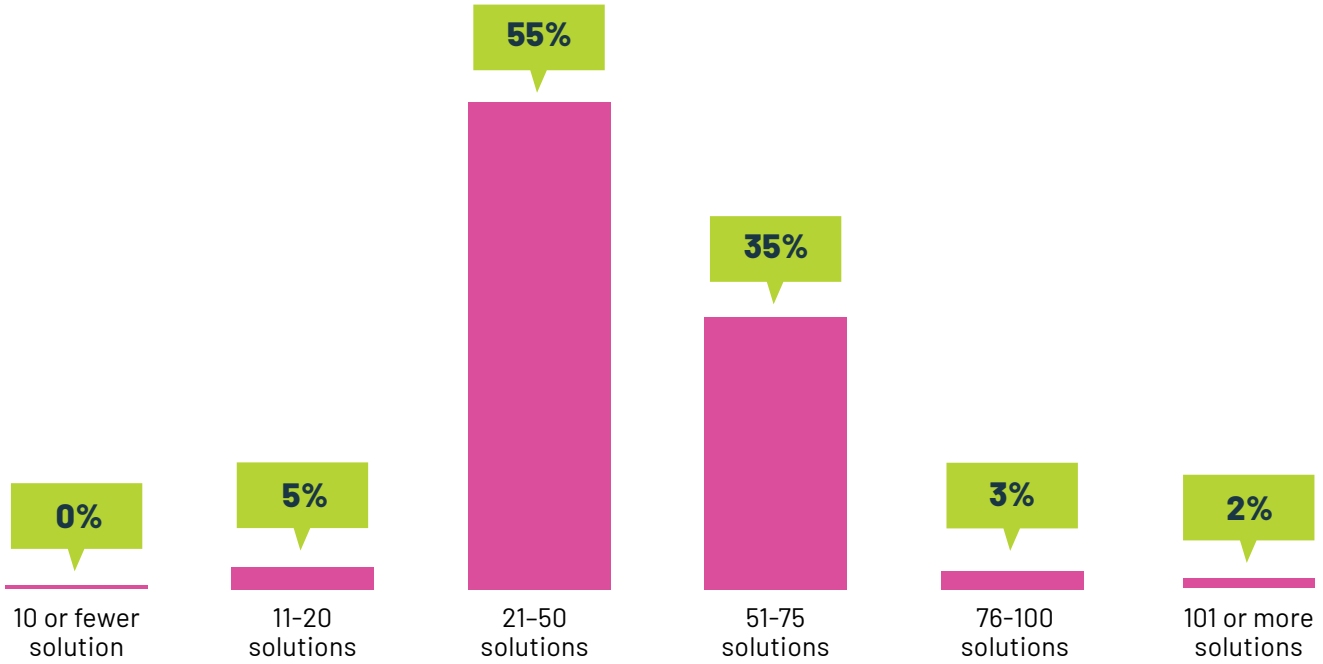


Additionally, rather than improving responsiveness, large tool stacks can slow investigation and decision-making while adding day-to-day operational overhead. According to Splunk's *State of Security 2025: The Stronger, Smarter SOC of the Future report*, 46% of security teams admit they spend more time maintaining security tools than actively defending their organizations.

## Security stacks are still growing

While consolidation continues to be a common theme in cybersecurity discussions, the data does not yet reflect a reduction in the number of tools in use.

Over the past 12 months, security stacks have continued to expand rather than contract. **More than two-thirds of enterprises (68%) report a net increase in security tools**, with most adding incrementally. 64% introduced between one and five new solutions, and 4% added six or more. By comparison, fewer than one-third report no change, and only 4% achieved a net reduction.

These findings suggest that consolidation, while frequently discussed, remains largely aspirational for most enterprises. New security tools are more often added than replaced, resulting in continued stack growth. Although downsizing security stacks may simplify operations, the data indicates that CISOs remain cautious about reducing controls where the perceived impact could weaken their security posture. Additionally, as we will see later in this report, external pressures such as cyber insurance requirements and emerging risk domains continue to influence security decision-making, reinforcing the tendency to add controls rather than remove them.

**In the past 12 months, how has the net number of security solutions in your security stack changed?**

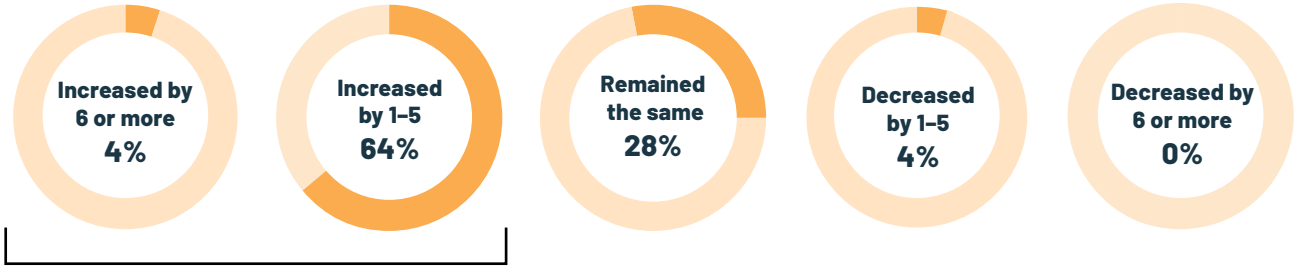


68% of enterprises report growth in their security stack in the past year

The net effect is continued expansion of security stacks across environments that are already complex. As stacks grow, the challenge shifts from coverage to coordination. Security teams must manage an increasing number of technologies, configurations, integrations, and data sources, adding operational overhead to day-to-day security operations.

## How AI is influencing security stack consolidation

AI is already influencing how enterprise CISOs think about the future of their security stacks, even if it has not yet driven widespread change in practice. **85% of CISOs admit that AI is influencing their security stack consolidation strategy**, indicating that AI is firmly part of long-term planning discussions.

That influence has not yet been translated into action. Only 3% of enterprises report actively consolidating their security stack as a direct result of AI capabilities, suggesting that AI-driven consolidation remains limited in execution today.

Beyond the impact of AI, 11% of enterprises report consolidating their security stack for reasons unrelated to AI, reinforcing that most consolidation efforts underway today are responding to cost, vendor overlap, or operational efficiency rather than AI-specific functionality.

This aligns with earlier findings across the report. Tool sprawl is increasing today, and while AI is expected to play a role in consolidation over time, AI-driven stack reduction remains largely aspirational at this stage. The conversation is active, and the optimism is real, but for most enterprises, AI's simplifying impact has not yet translated into actionable change.

**Is AI adoption influencing your plans to consolidate or reduce your security tech stack?**

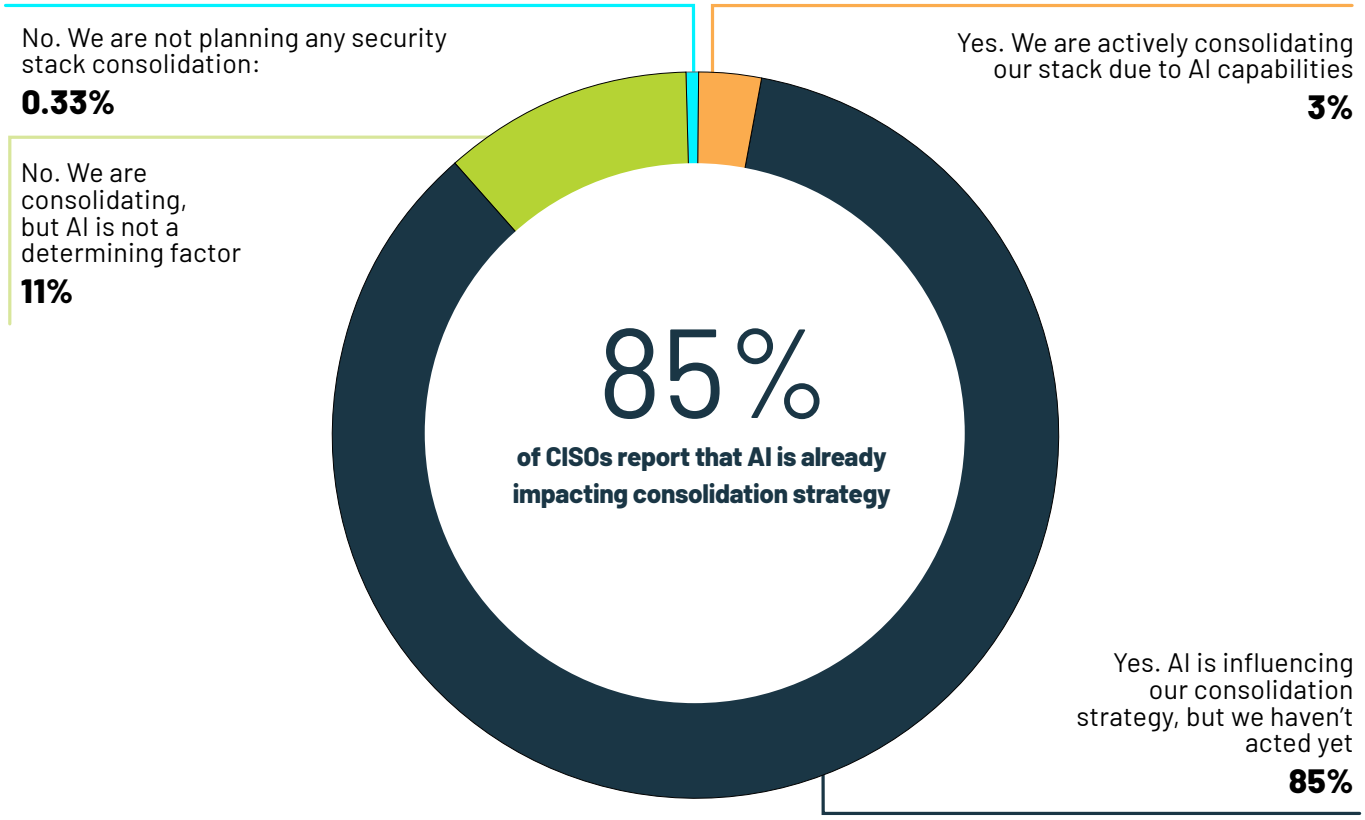

## Cyber insurance providers are driving tool adoption

According to IBM's Cost of a *Data Breach Report 2025*, the average cost of a data breach for U.S. companies jumped 9% to a record high of $10.22 million in 2025. As prices skyrocket, insurers are taking active steps to ensure security where they can.

**99.7% of U.S. enterprises report that their cyber insurance provider has either recommended or directly influenced the adoption of cybersecurity tools. Only 0.3% report that their insurer has never raised the topic.**

In many cases, insurer guidance translates directly into action. **48% of CISOs report implementing one or more security solutions specifically due to insurer requirements**. Even where adoption is not explicitly mandated, insurer influence remains strong: 58% of CISOs report receiving insurer-driven recommendations for specific security technologies.

The net effect is that cyber insurers have become a significant external driver of security tooling decisions. In practice, this influence often results in the addition of new controls to environments that are already complex, reinforcing the broader pattern of security stack growth observed earlier in this report. As underwriting scrutiny increases, insurers are shaping not only coverage terms, but also the composition and expansion of enterprise security stacks.

**Has your cyber insurance provider compelled you to integrate a cybersecurity solution you were not previously considering?**

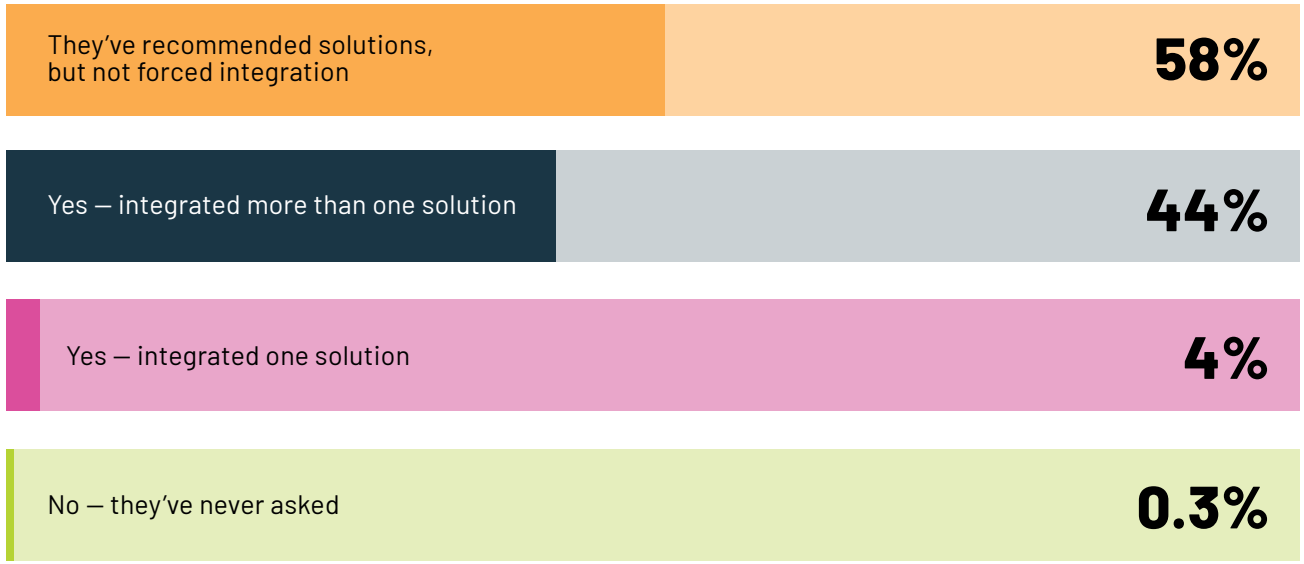

## Despite growing security stacks, successful cyberattacks remain common

Among respondents who explicitly disclosed whether they had experienced a cyber incident, approximately **89% report that their organization was compromised by a cyberattack within the past 24 months.** Within this group, 18% report a compromise in the past 12 months, while an additional 57% report being compromised within the past 24 months, underscoring that successful attacker access is not only widespread, but recent.

This represents **225 organizations out of 253 respondents** who answered either "Yes" or "No" to the compromise question.

To account for non-disclosure, the data can be viewed through two additional bounds. If all respondents who selected "Prefer not to disclose" are assumed not to have experienced a compromise, the rate would still be 75%. If those respondents are instead assumed to have been compromised, the rate rises to 91%. The true rate likely falls somewhere between these two extremes.

These findings indicate that **successful cyberattacks resulting in unauthorized access remain common across enterprise environments**. When considered alongside the scale and continued growth of modern security stacks, the data suggests that persistent compromise is not simply a function of underinvestment, but reflects the increasing difficulty of consistently validating, prioritizing, and managing real-world exposure across complex, interconnected systems.

**Has your organization been compromised by a cyberattack over the past 24 months?**

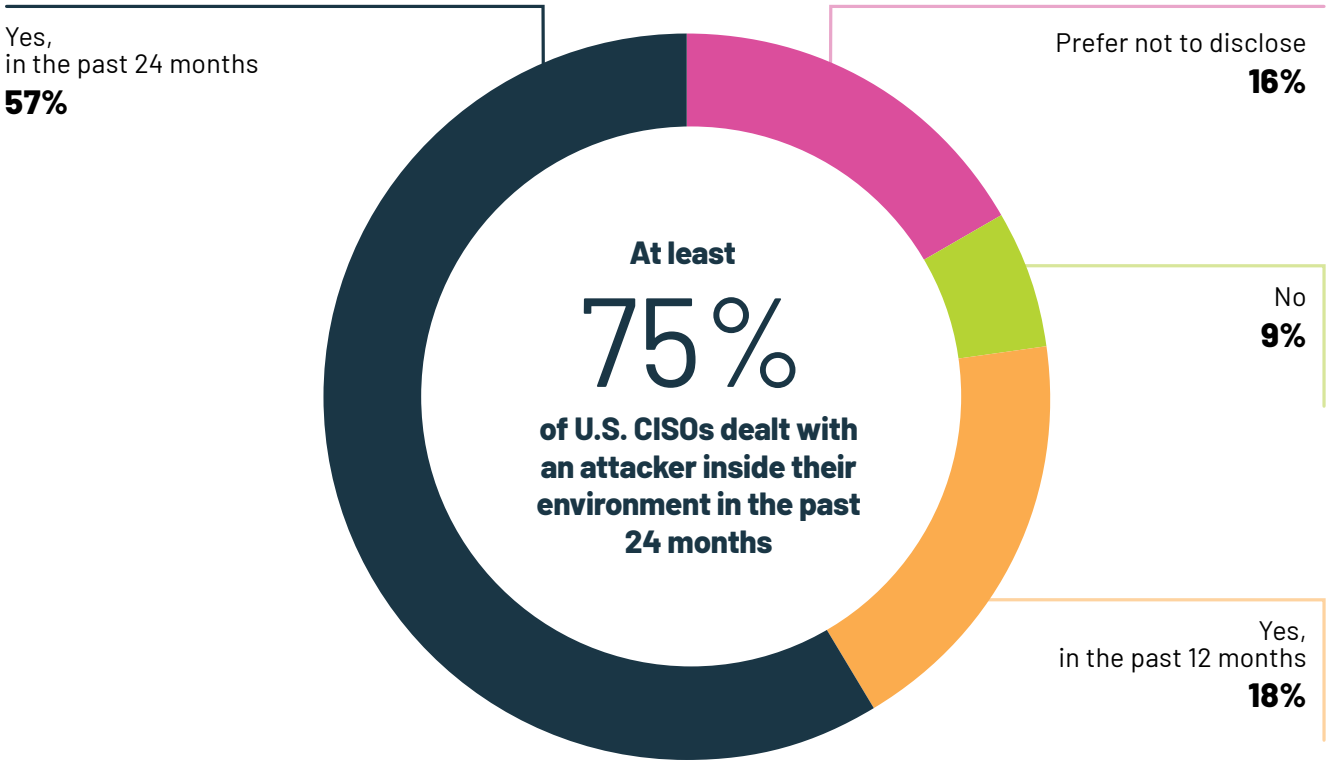

## Attacker Impact Reflects Freedom of Movement After Access

Among enterprises that reported attacker access, while not all intrusions progress beyond initial access, the pattern of impact reflects how adversaries operate once a foothold is established. Initial access most commonly occurs through externally reachable assets, which is why web-facing assets (63%) and endpoints (60%) appear most frequently impacted. These systems serve as the primary entry points into the environment. But attackers do not remain at the entry point. Their objective is to expand access as quickly and broadly as possible.

As attackers establish themselves inside the environment, they move toward identity, control planes, and connective infrastructure. Identity and access controls are impacted in 53% of cases, with APIs and cloud infrastructure each affected in 46% of incidents. While these rates are lower than initial access points, they demonstrate that a significant portion of attackers are able to progress beyond the perimeter, pursuing credential abuse, privilege escalation, and centralized access that enables movement across systems. Internal networks and servers are impacted in 47% of cases, reflecting routine lateral movement rather than exceptional behavior.

From there, attackers reach the systems that drive business operations. Applications are impacted in 30% of cases, and AI ecosystems in 18%, indicating that while many intrusions are constrained, a meaningful subset progress into business-critical systems once access is established. IoT and OT environments, impacted in 13% of cases, further illustrate the breadth of access attackers can achieve when segmentation boundaries are crossed.

This distribution does not describe isolated incidents. It reflects the reality of modern intrusions: once attackers gain a foothold, outcomes are shaped by how effectively organizations detect, contain, and limit movement across identity, infrastructure, applications, cloud, and AI systems as part of a single attack lifecycle.

### Which aspect(s) of your infrastructure were compromised?

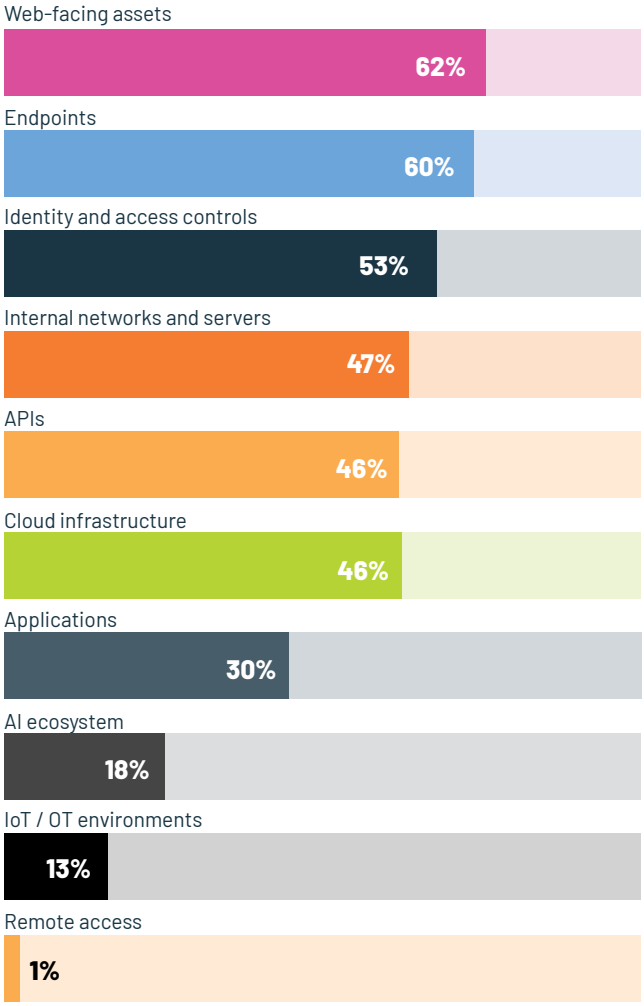




* Question allowed more than one answer and as a result, percentages will add up to more than 100%

## Enterprise security and testing spend is already at scale

Enterprise cybersecurity investment is already substantial. Nearly 79% of enterprises report annual cybersecurity budgets of $1M or more, with a weighted average spend of approximately $2.48M, excluding personnel costs. Multi-million-dollar security budgets are therefore common rather than exceptional across the enterprise landscape. That investment scales predictably with organizational size:

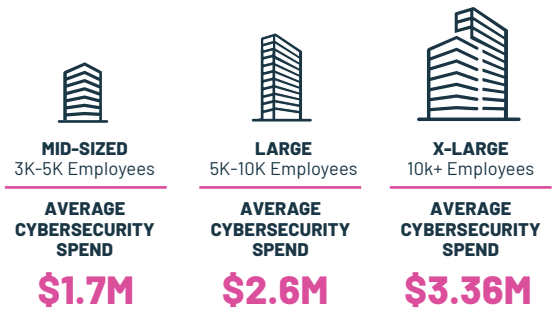


Within this broader spend, pentesting represents a consistent and meaningful allocation. Nearly 94% of enterprises spend at least $100K annually on pentesting, and no respondents report budgets below $50K. Across the enterprise sample, the weighted average pentesting budget is approximately $300K.

Pentesting investment also scales with organization size. Enterprises with 3,000–4,999 employees report average pentesting budgets of roughly $223K, increasing to $320K among enterprises with 5,000–9,999 employees, and reaching $354K for enterprises with 10,000 or more employees.

Viewed proportionally, pentesting accounts for approximately 12% of total cybersecurity spend on a weighted-average basis. This ratio remains stable as organizations scale, indicating that adversarial testing is treated as a core operating expense rather than a discretionary line item.

### How much did you spend on cybersecurity overall in 2025 (not including salaries)?

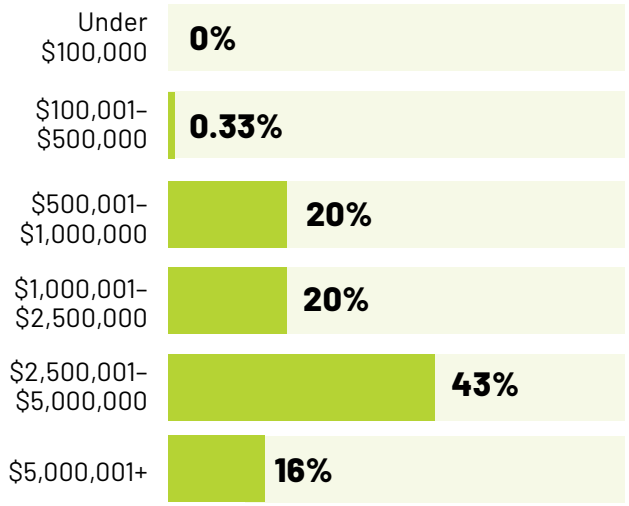


### What is your current annual budget for pentesting in 2025?

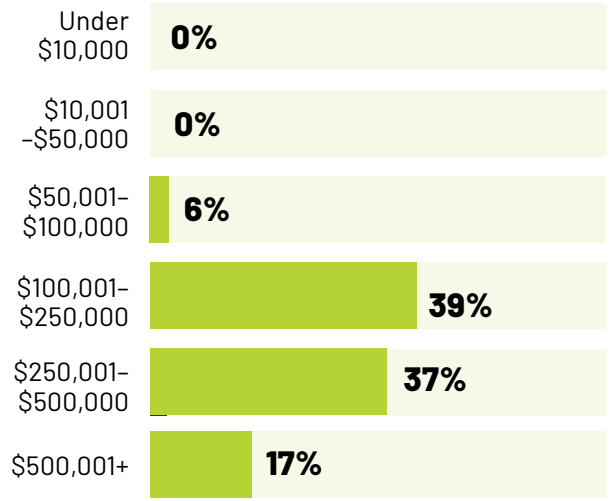

## Pentesting remains core to growing enterprise security investments

Enterprise cybersecurity budgets are expected to expand into 2026, with 66% of enterprises planning to increase overall security spending. Among enterprise CISOs, 23% anticipate double-digit budget growth, reflecting continued investment across a widening set of security priorities. At the same time, 29% of enterprises expect budgets to remain flat, and only 5% anticipate any decrease, indicating that budget contraction is rare in enterprise environments.

Within that expansion, pentesting continues to hold a stable and growing position inside enterprise security programs. Nearly 70% of enterprises plan to increase pentesting budgets in 2026, reinforcing that adversarial testing is treated as a core operational requirement rather than discretionary spend. Most increases fall in the 1–10% range, consistent with how mature, ongoing enterprise security programs scale year over year.

Stability is also notable. 25% of enterprises expect pentesting budgets to remain unchanged, and only 5% anticipate any decrease, showing that pentesting investment is rarely among the first areas enterprise CISOs look to reduce, even when budgets are constrained.

In context, this pattern is significant. As enterprise security budgets grow and diversify to address new tools, platforms, and emerging risks, pentesting investment continues to rise or remain stable across the vast majority of enterprises. Even amid competing priorities, enterprise CISOs are maintaining and expanding their commitment to validation as a foundational component of enterprise security programs.

### Your overall spend on cybersecurity in 2026 (not including salaries) is due to:


### Your annual pentesting budget for 2026 is due to:

## AI security is already driving increased spend

99.7% of enterprise CISOs report that securing their AI ecosystem is driving some level of increased cybersecurity spend in 2026, with only a single CISO indicating no impact on their budgets.

Where that increase is landing is telling. The impact is primarily budgetary rather than organizational. 64% of enterprises report increased security spend without additional headcount, while 21% expect headcount increases only and 14% anticipate increases in both budget and staffing.

This distribution reflects the current state of AI security ownership. Responsibility for securing the AI ecosystem is often shared across security, IT, and application teams. In this context, enterprises appear to be increasing investment through budget rather than headcount, with staffing expansion lagging behind spend as ownership models and scope continue to evolve.

**Is securing your AI ecosystem driving an increase in your cybersecurity spend for 2026?**

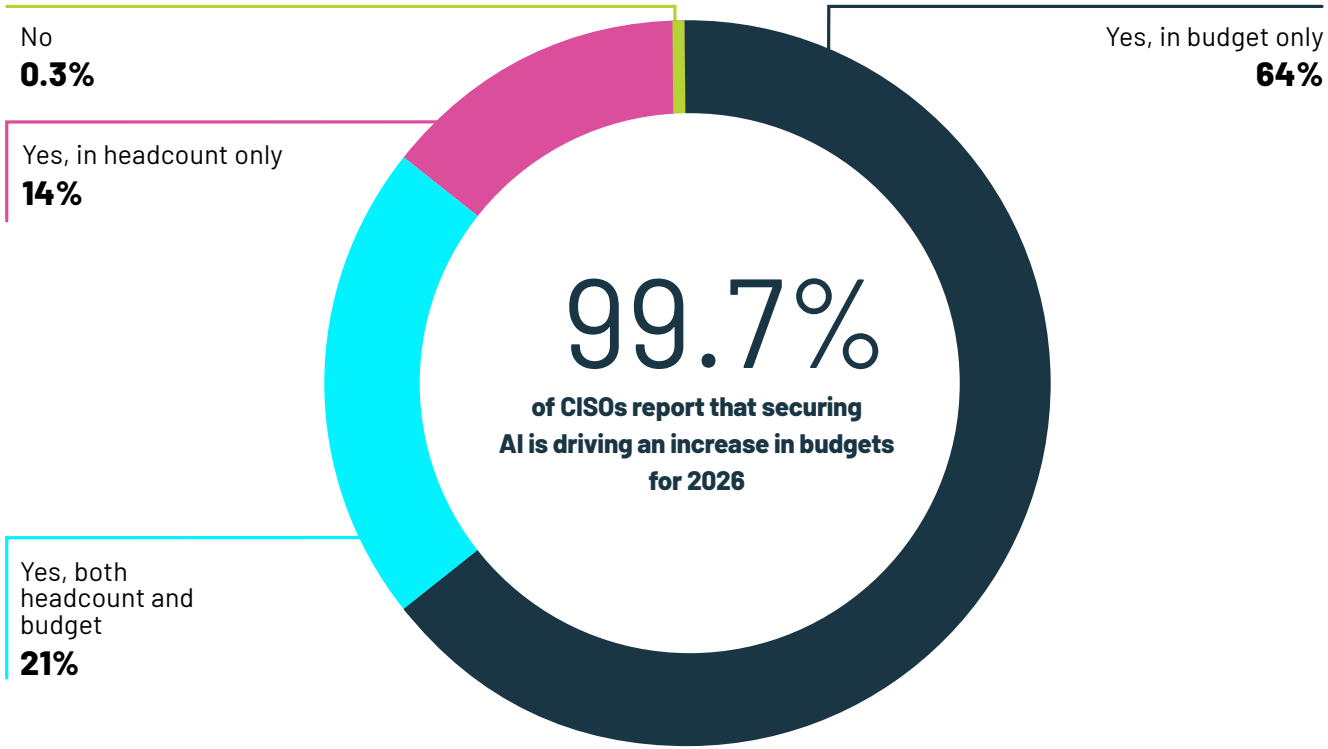

# What CISOs Can Take Away From This Report

As enterprises accelerate AI adoption, external signals reinforce the scale of the shift underway. In the World Economic Forum's *Global Cybersecurity Outlook 2026*, 94% of cybersecurity leaders identify AI as the most significant driver of change in cybersecurity in the year ahead.

The findings in this report show what that shift looks like in practice: AI is not coming into existence within a vacuum; it is being introduced into already complex IT and security environments, where scale, fragmentation, and inconsistent validation are long-standing challenges. For CISOs, the path forward is less about chasing AI-specific hype and more about strengthening foundational practices so they hold up in an AI-driven environment. Based on the data, five actions stand out.

### 1. Apply Mature Deployment Discipline to AI, Not Ad Hoc Adoption

AI systems are often deployed through teams, pilots, and embedded workflows without the same rigor applied to core enterprise technologies such as cloud infrastructure, identity platforms, or business-critical applications. Mature AI security begins upstream, by treating AI as a governed enterprise asset rather than an exception. This requires consistent discovery and inventory, clear ownership and accountability, and an understanding of what AI systems can access across data, identities, APIs, and infrastructure. Without this deployment discipline, security controls and adversarial testing are applied unevenly, limiting an organization's ability to validate risk in a repeatable and scalable way as AI adoption accelerates.

### 2. Use Adversarial Testing to Replace Assumptions With Evidence

Confidence in AI security correlates strongly with how often organizations test their defenses. CISOs who validate more frequently report higher confidence, while those testing less remain more uncertain. The takeaway is straightforward: **increase the cadence and scope of adversarial testing**, including AI-related scenarios, to continuously validate real-world exposure. Rather than waiting for standardized AI testing frameworks to mature, leading organizations are already incorporating AI into existing adversarial testing programs to understand how compromise actually unfolds across attack paths.

### 3. Focus on Attack Paths, Not Isolated AI Risks

The data shows that breaches still begin at familiar entry points such as web-facing assets, identities, endpoints, APIs, and cloud infrastructure. CISOs should resist treating AI as a standalone risk domain and instead evaluate how AI expands or accelerates existing attack paths. Testing should reflect how attackers chain weaknesses across systems, including how AI systems can be abused once access is gained. This attack-path mindset aligns with CTEM principles and helps prioritize remediation based on real impact rather than theoretical risk.

### 4. Invest in Skills and Validation Before Chasing New Tools

Lack of internal expertise remains the single largest barrier to securing AI, outweighing budget constraints. While AI-specific security tools are emerging and being actively evaluated, tooling alone will not close the gap. CISOs should prioritize **upskilling teams, clarifying testing methodologies, and strengthening validation practices** before layering in additional point solutions. Organizations that combine internal expertise with external validation partners are further along in operationalizing AI security than those relying on tools alone.

# A detailed look at the numbers behind this report

## Industries of respondents

**10%**
Banking

**10%**
Insurance

**10%**
Health & Pharma

**9.67%**
Information Technology

**9%**
Financial Services

**8%**
Telecom

**7.67%**
Retail & eCommerce

**7.33%**
Travel & Hospitality

**6.33%**
Software Development

**6%**
Transportation & Logistics

**5.33%**
Technology

**5%**
Energy & Utilities

**2.33%**
Government

**2.33%**
Public Services

**1%**
Media

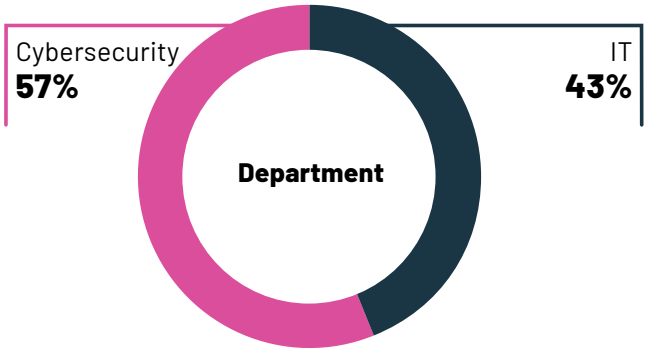


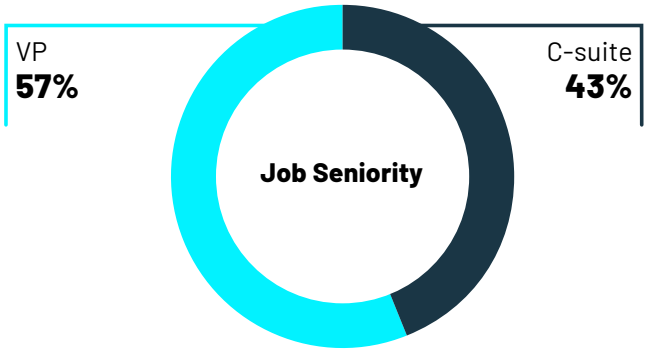


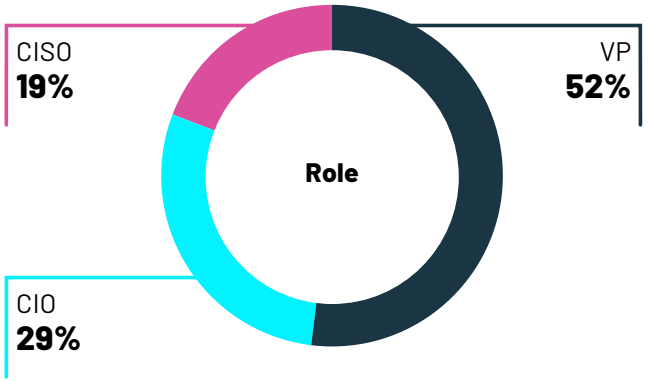


## Size of organizations

**28%**
**MID-SIZED**
3K-5K Employees

**51%**
**LARGE**
5K-10K Employees

**21%**
**X-LARGE**
10k+ Employees

## Pentera's Platform

Pentera's mission is to help organizations understand and reduce cyber exposure by validating what is actually exploitable across their IT environment and mobilizing remediation based on proven attacker outcomes. Rather than relying on assumptions or theoretical severity, Pentera provides evidence of real attacker risk so security teams can prioritize accurately, remediate faster, and reduce exposure across the business.

Pentera's platform delivers automated, safe adversarial testing that emulates real attacker behavior. Using an adaptive, AI infused algorithm, the platform performs the same actions a malicious adversary would without disrupting production systems. An agentic interface keeps security teams in control, allowing them to initiate, adjust, and query the attack testing scenario in real-time, enabling continuous, scalable testing while maintaining operational safety.

The Pentera platform validates exposures across the full enterprise attack surface, including internal networks, external assets, cloud and hybrid environments, and emerging AI infrastructure. By executing real, multi-stage attacks in production, Pentera proves how attackers combine vulnerabilities, misconfigurations, identities, and credentials to gain access, move laterally, and reach critical assets. This attack-path-driven validation shows what is actually exploitable and where remediation will reduce risk most effectively.

The Pentera platform then drives remediation operations around proven attacker outcomes. Findings are consolidated, deduplicated, and enriched with organizational and environmental context to prioritize root-cause security gaps and establish clear ownership. Remediation tickets are generated, prioritized, and assigned to the appropriate teams within the workflows they already use. Re-testing verifies that remediation has fixed exploitable attack paths, delivering audit-ready evidence of reduced exposure.

## About Pentera

Pentera is the market leader in AI-powered Security Validation, equipping enterprises with the platform to proactively test all their cybersecurity controls against the latest cyber attacks. Pentera identifies true risk across the entire attack surface and automatically orchestrates remediation workflows to effectively reduce exposure. The company's security validation capabilities are essential for Continuous Threat Exposure Management (CTEM) operations. Thousands of security professionals around the world trust Pentera to close security gaps before threat actors can exploit them.
**For more info, visit: pentera.io**


© 2026 Pentera. All rights reserved.


PENTERA.