ConductorOne

2025

Future of Identity Security Report

Balancing risk with speed, security leaders go all in on AI agents to address surging identity attacks.

Table of contents

Executive Summary

Security leaders are running fearlessly toward agentic AI. Not blindly, but with urgency and budget to match. You'd think CISOs might shy away from a risky new technology, but the opposite is true: they're embracing it. Fast.

When you think about identity in 2025, you can't escape agentic AI. AI agents are being deployed inside security teams, pushed by executives, and quietly adopted by employees. It's dynamic, autonomous, and ephemeral. And it's here now.

Despite rising identity-based attacks and mounting stress, security leaders aren't hitting the brakes. They're adapting. They're investing in identity security. And they're betting that controlled speed beats static security.

For the second annual *Future of Identity Security Report*, ConductorOne surveyed 494 IT security professionals, manager level and above, at US companies with 500+ employees across a variety of industries, including financial services, healthcare, technology, manufacturing, and more. The data shows threats are moving fast, and so are the people fighting them.

Here's what we found:

- Security leaders are going all in on agentic AI, despite their deep concerns about the risks posed by this technology.
- The board and executive leadership are actively pushing for security teams to embrace AI, fueling top-down pressure to move faster.
- Identity-based attacks are on the rise and IAM budgets are following suit.
- Security leaders face high levels of stress, feeling the weight of rising expectations and relentless threats.
- Improving security is now the top IAM priority, outranking risk reduction, compliance, and productivity.

Key findings

AI agents present risks and opportunities

100%
494 IT security professionals

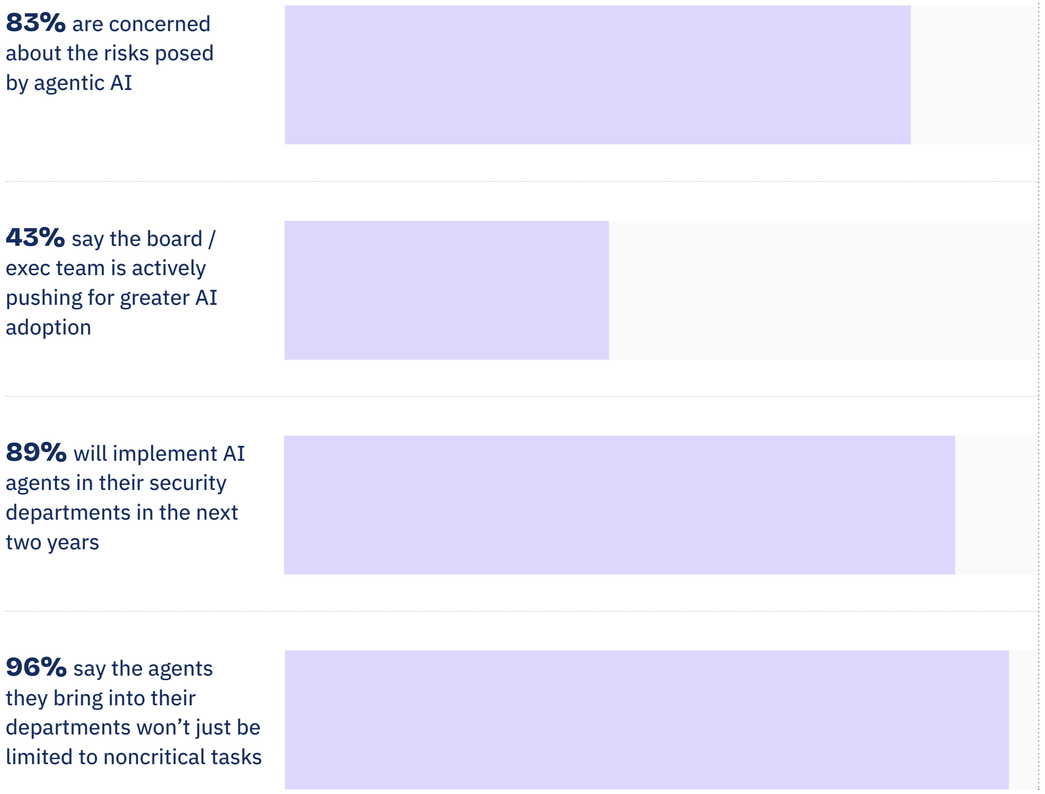

83% are concerned about the risks posed by agentic AI

43% say the board / exec team is actively pushing for greater AI adoption

89% will implement AI agents in their security departments in the next two years

96% say the agents they bring into their departments won't just be limited to noncritical tasks

Attacks increase, budgets follow

82%
experienced at least one identity-based attack (up from 77% last year)

84%
say their budgets for identity and access management are increasing

Pressures are high

More than 1 in 4 respondents report high or very high stress on a regular basis. Their top stressor? Preventing cyberattacks and data breaches.

Challenges and priorities in identity management

Top 3 challenges

1. System complexity
2. Tooling limitations
3. Employees ignoring policies

Top 3 priorities

1. Improving security
2. Reducing risk
3. Achieving compliance

The State of Agentic AI

The agentic AI train is here. Are you boarding?

Since our last annual survey, one trend has moved to center track: agentic AI. It's top of mind for every business and security leader, promising to fundamentally change the way business is done.

Google Trends data shows an initial spike in searches for the term "agentic" beginning in the fall of 2024, with a continued increase in searches well into 2025.

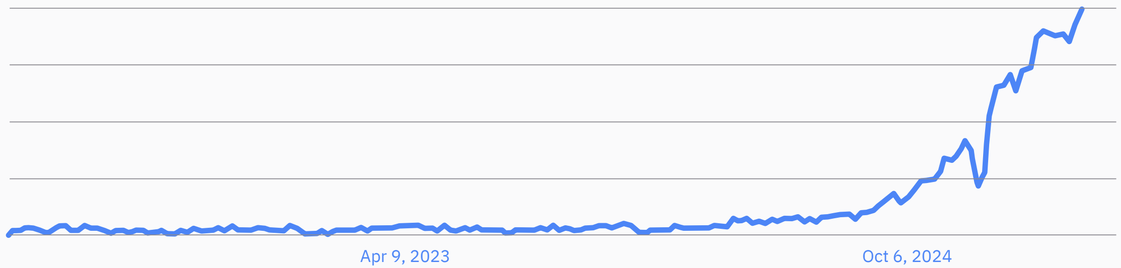

Source: Google Trends

It's not just hype. AI agents are reshaping how work gets done, decisions get made, and access gets granted. Security leaders aren't ignoring the risks. They see the blast radius. And still, they're charging ahead—because standing still isn't an option.

Overall, 83% of respondents say they are concerned about the risks associated with AI agents—with 41% stating they are "very concerned."

How concerned are you about the potential security risks associated with AI?

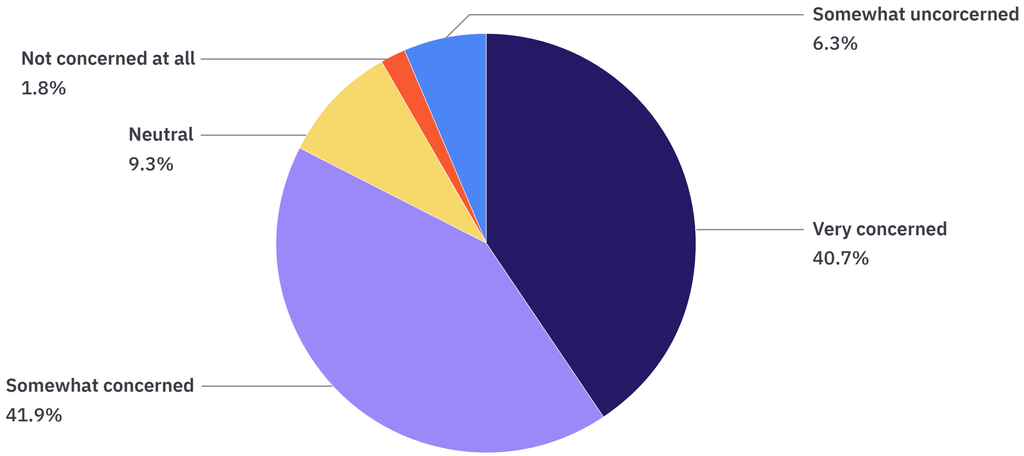

Those who have experienced multiple instances of identity compromise in the past year are more acutely aware of the dangers presented by AI agents. They see that the rapid introduction of autonomous AI systems across their environments is sure to introduce new risks. In fact, these respondents were 32% more likely to be "very concerned" about AI agents compared to the overall group.

Only 2% stated they were "not at all concerned" about agentic AI risks.

Ask 10 security leaders what worries them most about agentic AI, and you'll get 10 different answers. Their concerns depend on what they protect, who they report to, and how fast their org is moving.

But across the board, two threats stand out: data security and AI hijacking. No matter the sector or maturity level, those risks are keeping people up at night.

Top Agentic AI Threats	Rank
AI agents accessing sensitive data without proper authorization	1
AI agent revealing sensitive data or PII inappropriately	2
AI agents being hijacked or manipulated by third-party attackers	3
Employees implementing rogue AI agents without the security team's knowledge	4
AI agents retaining long-lived access to sensitive applications	5
Data poisoning or other compromise of AI system	6
AI agents hallucinating or making incorrect decisions that lead to security risks	7

Unsurprisingly, respondents hailing from larger companies, publicly traded companies, and the healthcare sector were more likely to cite data security and privacy concerns.

Respondents whose companies had experienced multiple identity compromises in the past year, however, were more likely to cite identity security concerns, such as the threat of AI agents retaining long-lived access to sensitive applications.

Only a handful (4%) of respondents say they do not anticipate any security threats from AI agents in the coming year.

Building guardrails for agentic AI risks

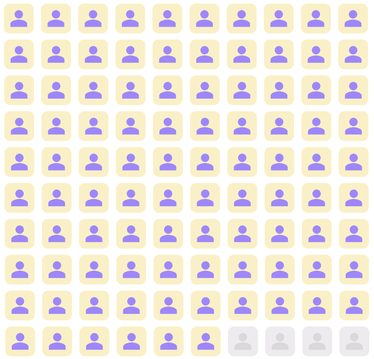

96% of security leaders expect agentic AI risks in the next year. So how are they getting ready?

We found there's no real consensus on best practices, at least not yet. Every security leader is taking different steps to prepare, but the vast majority (98%) are already taking action of some kind.

What measures, if any, is your organization taking to prepare for the security challenges posed by AI agents?

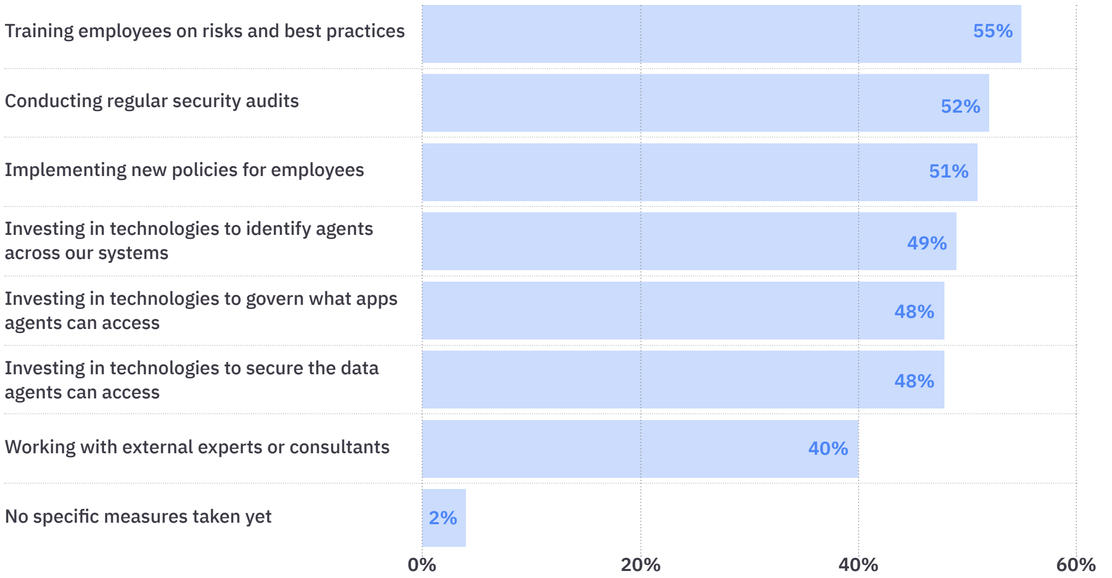

Fighting agents with agents

Despite being fully aware of the risks that AI agents will introduce across their businesses, respondents still expressed an overwhelming appetite to use AI agents within their security departments—fighting agents with agents.

The vast majority (89%) of respondents say they plan to implement AI agents within their security departments in the next two years. Only 2% say they will refrain from using AI agents, while another 9% are still unsure about their agentic AI plans.

Do you plan to implement AI agents in your security department in the next two years?

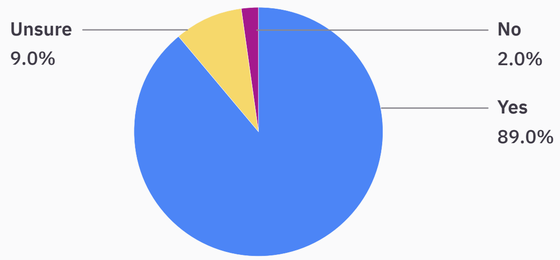

When it comes to operationalizing AI agents, the top anticipated use cases include network monitoring and analysis (49%), SOC automation (47%), and access requests and account provisioning (46%)—all high-friction areas where scale and speed matter.

Most Anticipated Use Cases of Agentic AI for Security	Rank
Network monitoring and analysis	1
Automating tasks in the security operations center (SOC)	2
Automating access requests and account provisioning	3
Summarizing information to improve productivity	4
Writing first drafts of simple scripts (e.g., SQL) to save time	5

Building the agent stack

Security leaders aren't jumping in blindly. They're thinking carefully about how to bring agentic AI into their environments. When we asked respondents how they plan to integrate AI agents into their workflows, we received a measured response.

Some are opting to permanently keep a human in the loop, while others are planning to start small before scaling. Another cohort is planning to use a "dual-agent system," where one agent completes tasks and a second agent monitors and regulates the first for compliance and security.

Which of the following best describes how you plan to integrate AI agents into your team's workflows?

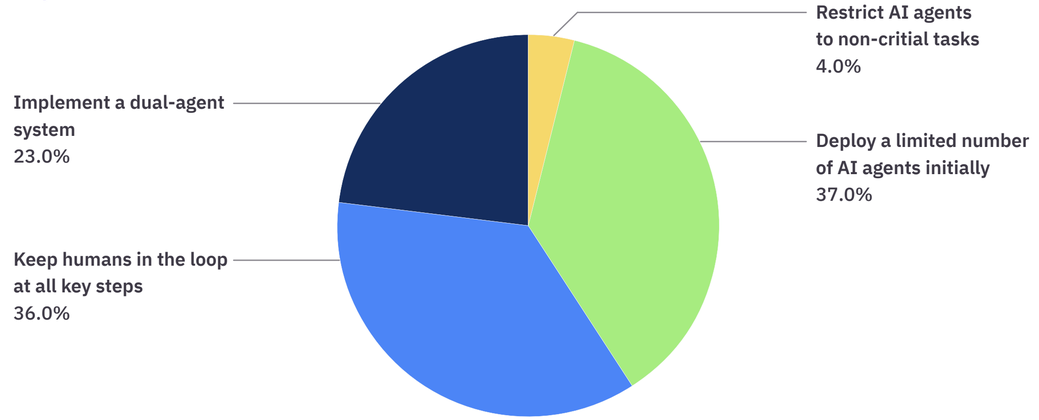

Despite the differences in how they plan to implement agents, there is still a general consensus that these AI agents will not just be restricted to noncritical tasks. In fact, 96% of respondents state that even critical tasks—ones where the impact of errors can be significant—are still fair game when it comes to agentic AI.

We also asked respondents about their software supply chain strategies to better understand how they plan to build these agents. Similarly to their workflow integration strategies, respondents express a wide range of approaches, with roughly a third planning to use corporate-provided AI models. However, a surprising number of respondents (19%) say they're planning to start with an open-source model like DeepSeek.

What is your software supply chain strategy for AI agents?

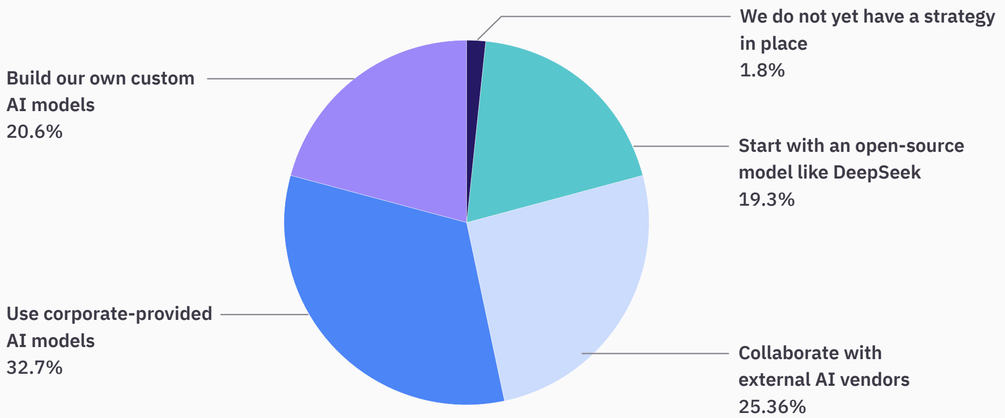

The Identity Security Snapshot

Identity-based attacks are on the rise

The bad news? Identity-based attacks are on the rise. In fact, 82% of respondents stated their organization experienced at least one cyberattack or data breach in the past 12 months due to improper access or over-privileged users—up from 77% in 2024.

37% stated there had been multiple instances of these identity-based attacks in the past year. Certain industries appear to be feeling the pain more acutely. For example, respondents from the financial services and technology sectors were more likely to indicate multiple instances of identity compromise.

When asked what attack vectors they believe to be the most significant threat, the majority of respondents pointed to malware and ransomware, particularly among healthcare respondents. However, upon a closer examination of the findings, we saw that identity-related vectors (e.g., phishing, password attacks, and non-human identity exploits) collectively represent a larger portion of concern. Taken together, these results suggest that while malware garners attention, the underlying issue for the majority of attacks is identity compromise.

In the past 12 months, have there been cyberattacks or data breaches at your organization due to improper access or over-priviledged users?

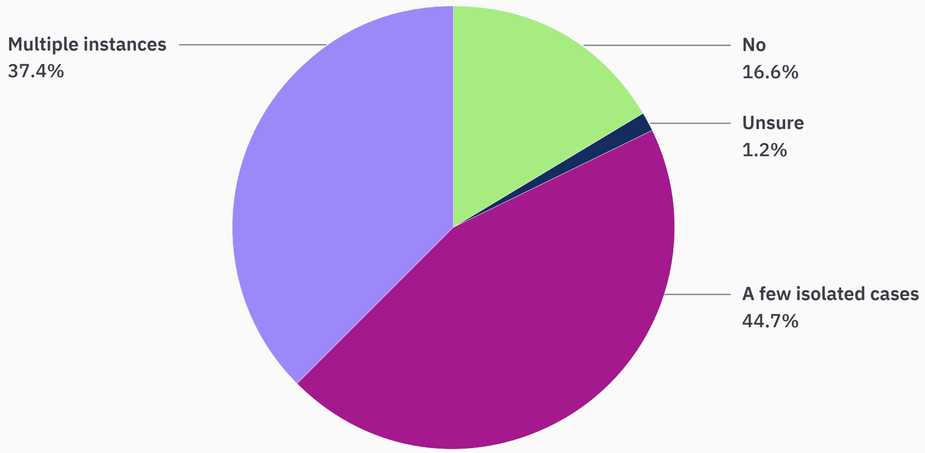

Identity-related attack vectors

Most Signficant Attack Vectors	Rank
Malware/Ransomware	1
Phishing	2
Software vulnerabilities	3
Password attacks	4
Cloud service misconfigurations	5
Insider threats	6
Supply chain attacks	7
NHI exploits	8

Complexity and tooling are big challenges

For the second year in a row, the complexity of existing systems was cited as the biggest challenge security leaders face when it comes to identity and access management, as noted by half of our respondents. This was followed by the limitations of the tools they currently have available (41%), which was also cited last year as a top-three challenge.

New this year, security leaders identified that employees ignoring policies has now become a top challenge, as cited by 39% of respondents. This was followed closely by a sentiment among respondents that employees are resistant to change (37%).

What are the biggest challenges your company faces when it comes to identity and access management?

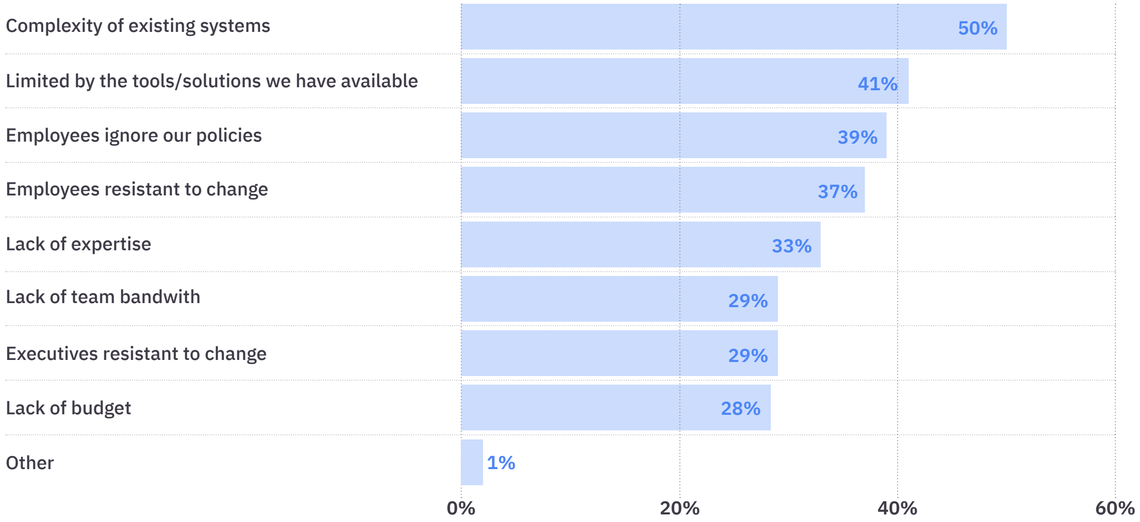

Looking closer at the data, there were several intuitive findings. For example, a lack of team bandwidth was a more common challenge among respondents who are experiencing high or very high stress levels, as well as those who reported multiple identity-related attacks in the past year.

However, we also found several counterintuitive findings, particularly among smaller organizations (those with less than 2,500 employees). These companies were notably less concerned about team bandwidth or even budgetary constraints. Instead, smaller organizations were more likely to say employees were ignoring policies, perhaps pointing to a less mature culture of identity security.

Meanwhile, larger companies (those with 10,000+ employees) were more likely to say their biggest IAM challenge was a lack of budget or a lack of expertise.

Additionally, the survey showed that security leaders at both larger companies and publicly traded companies were less likely to identify resistance to change among their executive teams, showing that even executives at the most prominent companies are recognizing the need to improve their identity security practices.

Putting money where the risk is

Good news for security leaders: IAM budgets appear to be increasing to keep up with the pace of attacks. Some 84% of respondents stated their companies' IAM budgets are increasing either moderately or significantly. We saw the exact same percentage of respondents indicate a budgetary increase in last year's report.

This year, significant budget increases were more likely among respondents who experienced multiple instances of identity incidents in the past year—43% vs. 28% overall.

How has your company's budget allocation for identity and access-related products changed for the upcoming year?

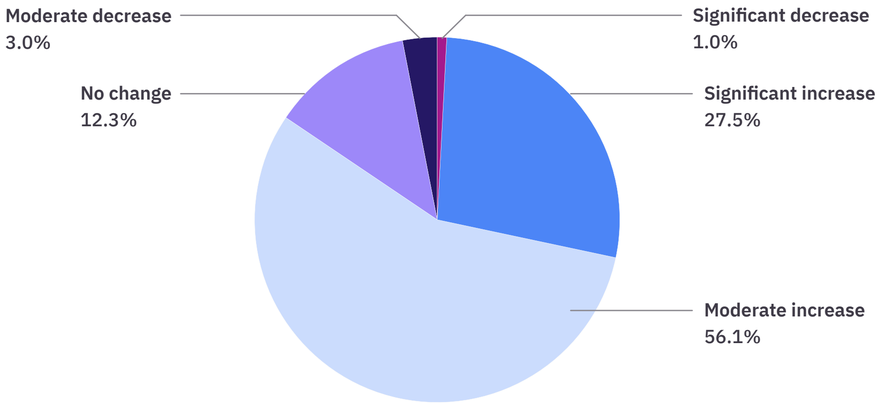

It's worth noting that despite these budgetary increases, more than a quarter (28%) of respondents still stated that a lack of budget was a top concern for their identity security program. This could indicate that budgets may not be increasing fast enough to stay ahead of the problems.

Improving security is the top priority for identity

When asked what they believe to be their top priority for identity and access management, the majority (77%) indicated improving security. This represents a significant shift in how businesses are thinking about identity and access. An area that was once considered to be little more than a productivity nuisance has now become a top security consideration, in direct response to the sharp increase in identity-based attacks.

Subsequent priorities among respondents include reducing risk (60%), achieving compliance (50%), and improving team productivity (49%). Intuitively, we found that larger organizations, as well as those in heavily regulated industries like financial services and healthcare, were more likely to indicate compliance as a top priority.

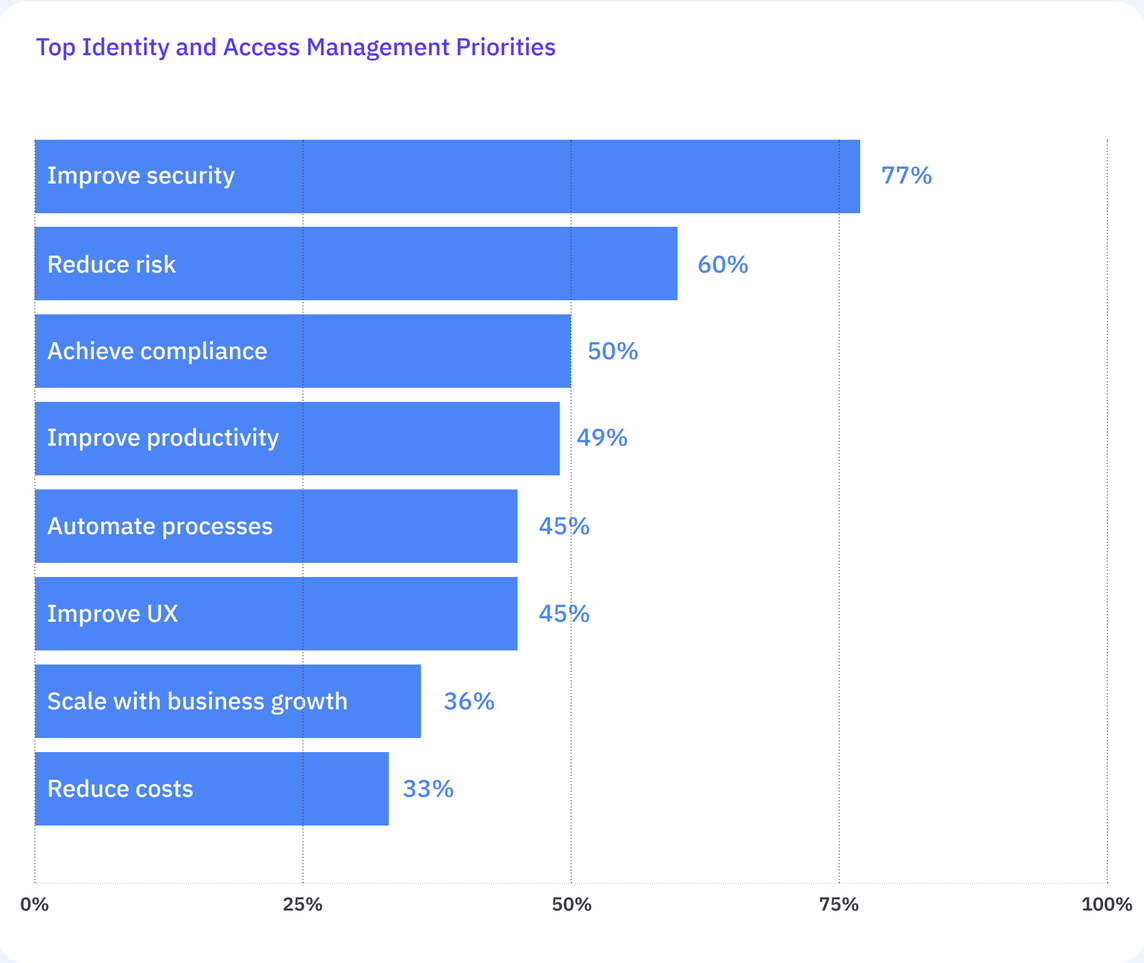

Boards and execs are mixed in their AI mandates

Every organization has its own unique appetite for AI. In some cases, the board or senior leadership team takes an active role in dictating how AI should be incorporated into the business. Others instead defer that decision to the relevant subject matter expert or team leader.

The *2025 Future of Identity Security Report* found a wide range of corporate cultures, all with varying degrees of interest in using AI to improve security operations. The largest cohort of respondents (43%) said that their board or executive team is actively pushing for the use of more AI to improve security. Another group (37%) said the decision on whether or not to use AI is left up to the security team.

Some sectors, such as financial services and technology, were much more likely to indicate the board or executive team was driving AI adoption. Others, such as healthcare and manufacturing, stated that this issue received much less attention from senior leadership.

Although the path to AI may vary, only a small number of respondents (9%) stated that senior leadership was skeptical or cautious about AI adoption, even knowing the potential risks.

How does your company's board and/or executive team view the use of AI to improve security operations?

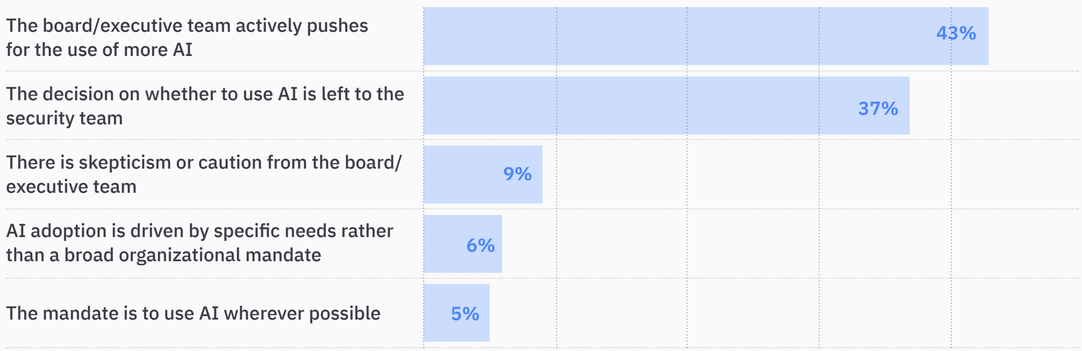

Non-Human Identities on the Rise

The NHI problem just got bigger

Agentic AI is the latest form of a familiar technology: non-human identities. Also known as "machine identities," these entities refer to any digital credential used by machines, workloads, applications, or automated processes to authenticate and perform operations on a given resource. These can include service accounts, API keys, tokens, secrets, and certificates.

NHIs are everywhere, often invisible, and growing fast.
With agentic AI in the mix, that growth is exponential.

So it tracks that the vast majority of respondents (93%) stated that the risks associated with NHIs are urgent, with 24% stating the risk is "extremely" urgent and requires immediate action. Financial services and technology sector respondents were more likely to indicate an extreme sense of urgency. Only 2% said they do not believe NHIs pose a risk at this time.

How urgent do you consider the risk associated with NHIs at your company?

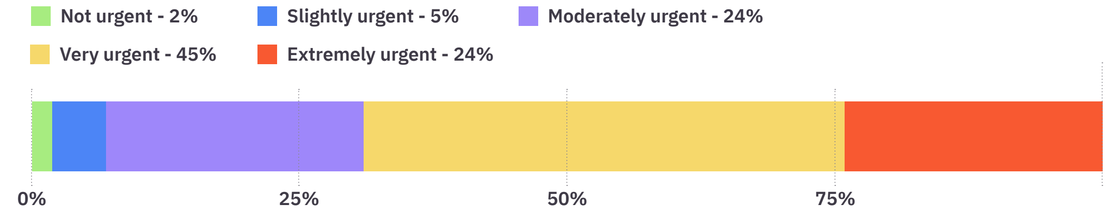

Humans vs. non-humans

With such a strong sense of urgency around NHIs, companies are taking steps to get ahead. But some security leaders may be leaning too far in—putting more emphasis on NHIs than on securing human users.

The survey findings showed that more than half (51%) of respondents believe the security of NHIs to be just as much of a priority as security for traditional human user accounts. However, 42% stated that NHI security is a higher priority than human users.

This could be a reaction to industry hype around NHI risks. It's also possible that respondents feel confident in their controls for human access and are shifting focus to what feels more urgent: NHI-driven risks.

How does your organization prioritize the security of NHIs compared to human user accounts?

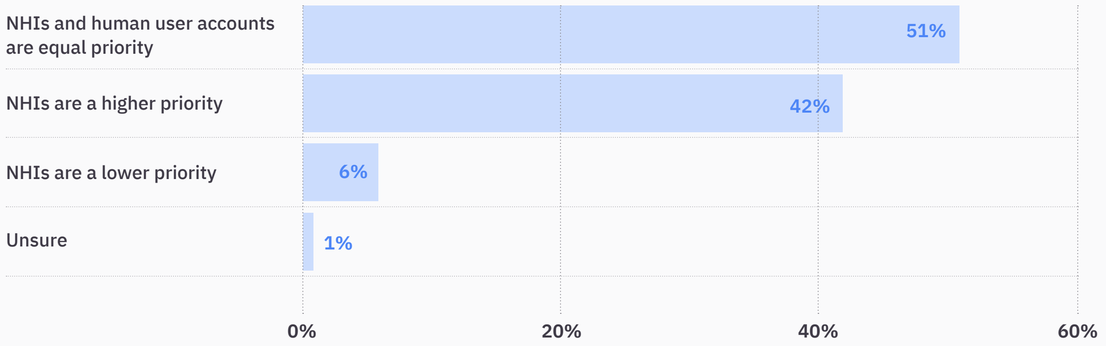

Clear tracks ahead? Visibility claims are high

When asking survey participants what level of visibility they currently have into NHIs across their business, we expected fairly mixed results. However, respondents claim to already have a high degree of visibility—in fact, 78% said they have "high" or "full" visibility into NHIs across their environments. An impressive 30% say they already have total visibility.

What level of visibility do you currently have into NHIs across your company's environment?

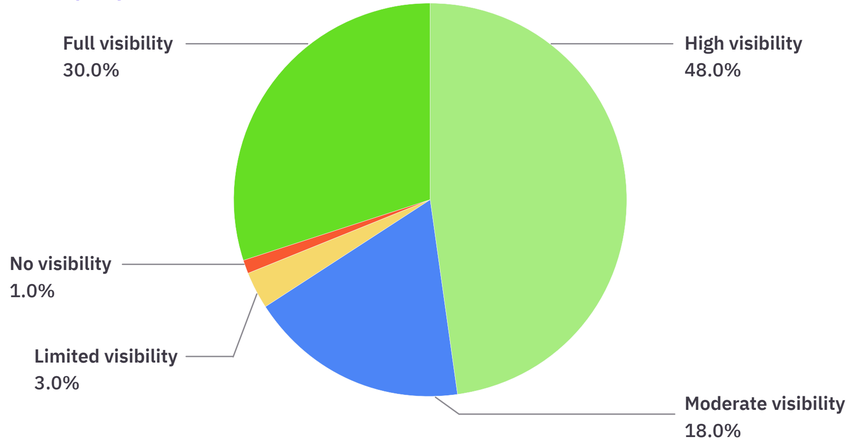

Larger companies (10,000+ employees) were less likely to say they have "high" or "full" visibility, which could point to the scale of the issue—the larger the organization, the more unwieldy its NHI posture.

Those who had experienced multiple identity-based attacks in the past year were more likely to have "full" visibility. As we've seen throughout the report, this cohort of security leaders has learned from past experiences, and is now taking the identity risk seriously.

Respondents who claim full visibility were much more likely to report extreme urgency around NHI risks (47% vs. 24% overall) demonstrating that they not only see the urgency, but are taking action.

Top NHI challenge: Ensuring the right privilege levels

When we asked survey respondents what they believe to be their biggest challenge in managing NHIs within their companies, the top response was ensuring that NHIs have the right privilege levels. Many NHIs are over-provisioned by default, and their access rarely reviewed once deployed, representing a potential entry point for attackers.

Surprisingly, despite respondents' claims of high visibility into NHIs, the second biggest challenge was reported to be "getting visibility into all NHIs across the business." Maybe they've already done the work to gain visibility into NHIs and know it's a challenge. Maybe they misread the question. Maybe both.

Here's what we do know:

- Those who experienced multiple identity compromises last year were more likely to say NHI governance was their top challenge.
- NHI credential rotation is a bigger challenge for larger companies.
- Financial services are more concerned about third-party compromises targeting NHIs.

What do you perceive as the biggest challenges in managing NHIs within your company?

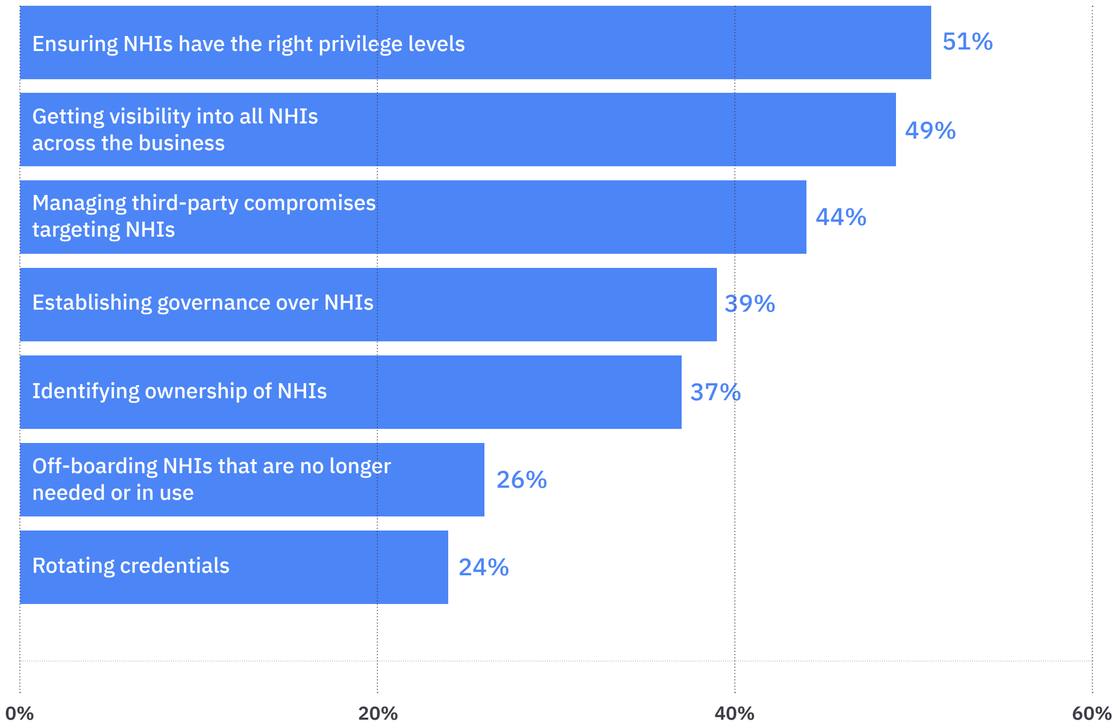

Pressure, Stress and Well-Being

Running on fumes

Cybersecurity is widely known to be a high-stakes, high-stress, profession. Now, it's a high-speed one, too. As attacks increase in severity and frequency, that stress also increases. One study from ISACA found that two-thirds of cybersecurity professionals said their role is more stressful now than it was five years ago.

We decided to ask our survey participants about their stress levels on a typical day. Overall, we found that about half say they face a "moderate" level of stress on a daily basis—a regular amount of stress that sometimes affects their productivity or mood. On the surface, this seems somewhat promising.

However, more than 1 in 4 respondents stated they experience "high" or "very high" levels of stress. Perhaps not unsurprisingly, those with higher levels of stress were more likely to indicate they had experienced multiple instances of identity compromise in the past year, showing that the job of a security leader naturally becomes more stressful when faced with increased attacks.

Which of the following best characterizes your stress level on a typical day?

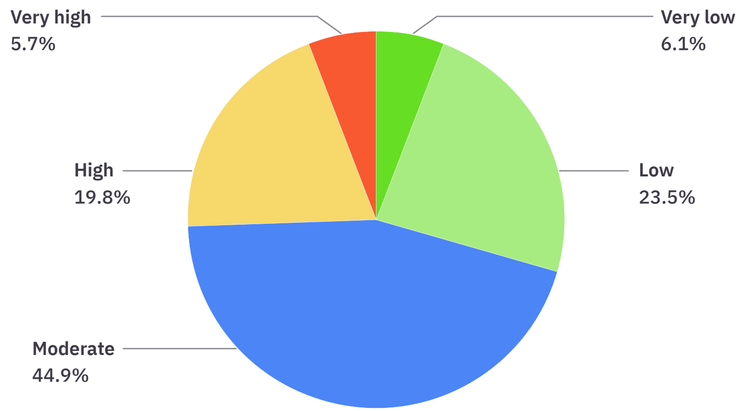

Pressure on the engine

For the second year in a row, ConductorOne has asked respondents about the biggest pressure they face in their role. Last year, the majority indicated their biggest pressure was keeping up with the rapid pace of technological change. This year, however, respondents stated their top pressure was preventing all cyber threats and ensuring no breaches occur.

What's the biggest pressure you face in your role?	Rank
Preventing all cyber threats and ensuring no breaches occur	1
Keeping up with the rapid pace of technological change and new attack vectors	2
Managing with a reduced budget while maintaining security posture	3
Operating with reduced headcount while maintaining security posture	4
Being the person to take the blame in the event of a breach	5
Justifying the value of security investments to other business leaders and stakeholders	6

Surprisingly, one of the lowest ranked pressures was having to be the fall-person—the one who takes the blame in the event of a breach. It's not that security leaders don't feel the heat. They're just more focused on stopping threats than worrying about who takes the blame.

Yes, the stress is real, but so is the resilience. Security leaders don't fold under pressure. They operate in it.

Staying fueled

With such a demanding career, we wanted to know what security leaders are doing to support their mental health and well-being. Some of these methods may seem obvious, but we were encouraged to see the results. Not unsurprisingly, respondents with higher stress levels were less likely to report regular physical exercise, spending time with family and friends, or even getting adequate sleep.

There's a critical lesson here for security leaders: take care of yourself, and don't run out of fuel.

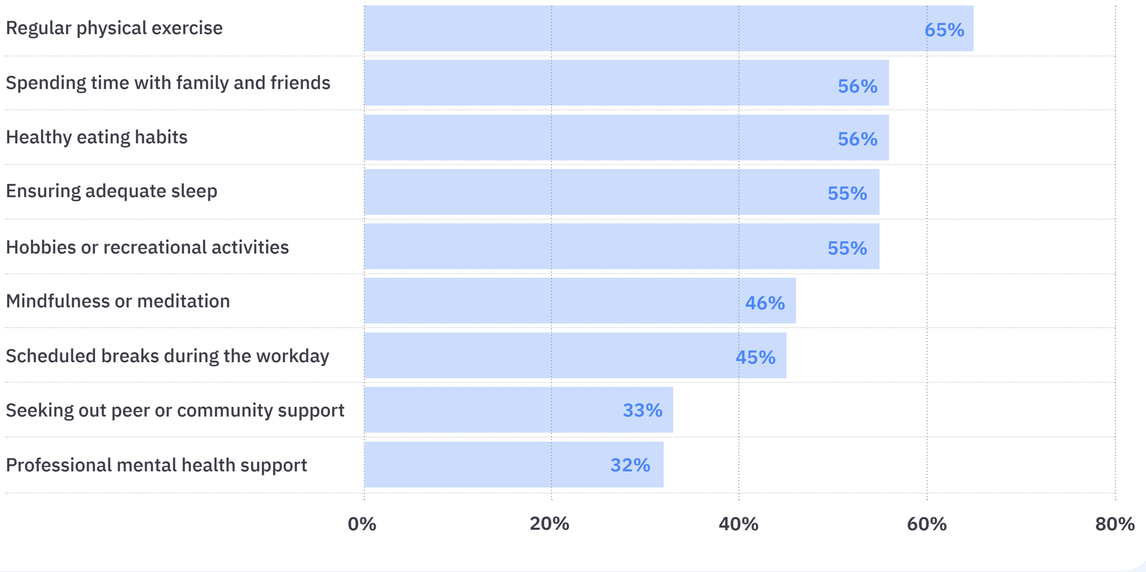

Conclusion

Agentic AI isn’t slowing down, and neither are the security leaders racing to meet it. Identity risks are growing, the pressure on security teams is mounting, and the margin for error is shrinking. Agentic AI is multiplying the complexity security teams face, and raising the bar for what it means to manage access and control.

But security leaders aren’t standing still. They’re embracing the shift, even as the ground moves beneath them: adapting their strategies, rethinking their priorities, and preparing for a future where identity isn’t just about users, but about agents acting on their behalf.

They’re not doing it blindly, but with intention. They’re figuring out best practices, and putting guardrails around AI agents so they can be deployed correctly.

As ConductorOne CISO Kevin Paige puts it:

> “We don’t put brakes on the train to make it go slow. We put brakes on so it can go fast. The same goes for agentic AI.”

Kevin Paige
Field CISO

The future of identity security isn’t about hitting the brakes, it’s about building the right tracks to accelerate velocity.

Demographics and Methodology

Industry

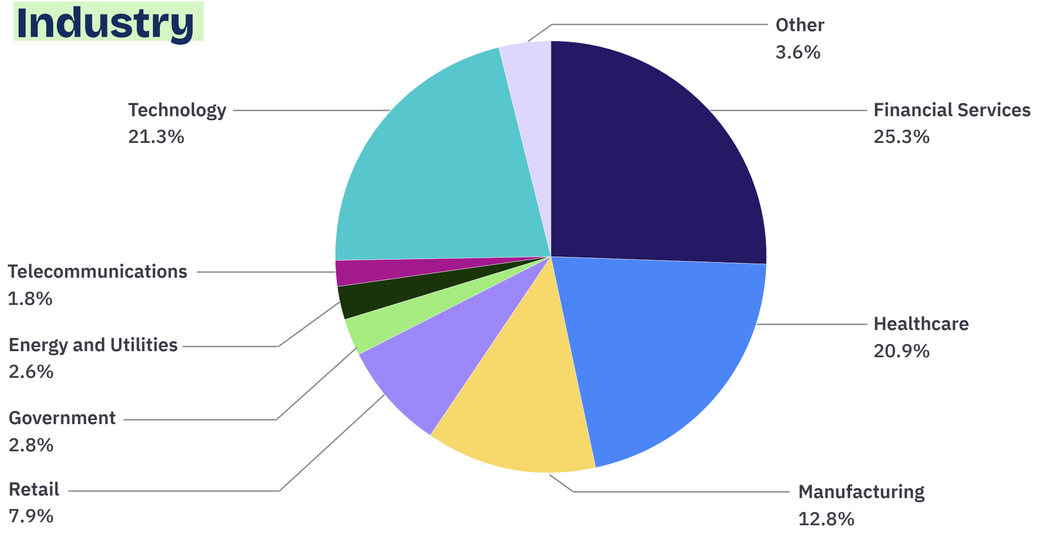

Company Size (by Employees)

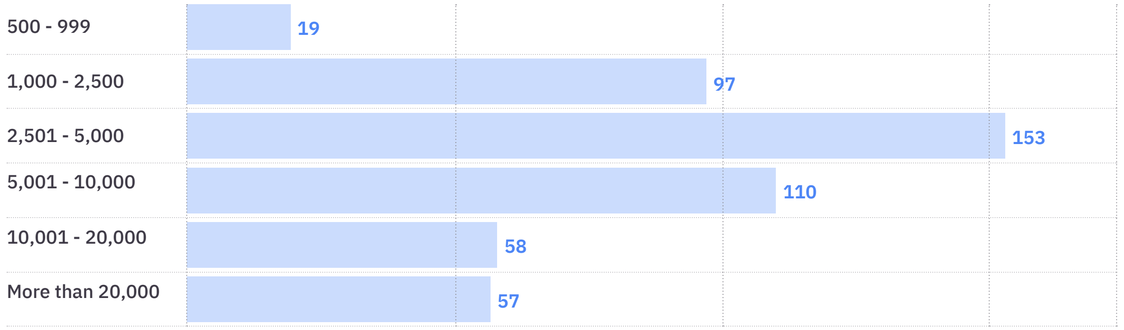

Company Structure

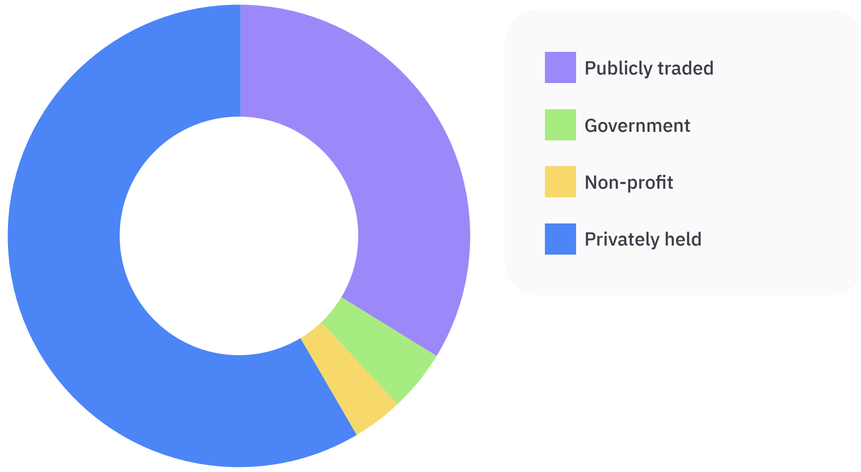

Methodology

The *2025 Future of Identity Security Report* findings are based on the results of an online survey conducted in April 2025 that examined the opinions of 494 US-based IT professionals, manager level and higher at companies of 500+ employees, whose roles involve cybersecurity.

ConductorOne

2025 Future of Identity Security Report

Talk to us to learn more