CYBERARK®
THE IDENTITY SECURITY COMPANY®

Report

2025 Identity Security Landscape

Perspectives on risk and readiness from security leaders.

A Security Matters Research Report by CyberArk

Table of Contents

Executive Overview

CYBERARK®

Executive Overview

Welcome to the 2025 Identity Security Landscape! This study wouldn't be possible without the generous insights from our 2,600 security decision-makers across 20 countries around the globe — a big thank you to our contributors and researchers.

If you scrolled here for the first time, welcome. This report specifically examines cyberattack trends impacting identity across modern IT ecosystems and shares insights on how security professionals can and should prepare.

Our returning readers are well aware that AI is arming both sides of the security battle, helping attackers and defenders alike. But what's more interesting this year is how the race to adopt AI has inadvertently expanded the attack surface with a surge of machine identities. Welcome to the third dimension of AI: attackers use it to create new threats; defenders use it to defend against them — and businesses incur new identity-centric risks as they embed agentic AI across the enterprise.

On one hand, we're seeing the most relentless and sophisticated cyberattacks of the modern age, with 9 out of 10 organizations reporting a successful identity-centric breach. Over half (51%) fell victim to phishing and vishing attacks multiple times. At the same time, respondents tell us sanctioned and unsanctioned adoption of AI is adding to cybersecurity risks. Organizations now report that 72% of employees regularly use AI tools on the job — yet 68% of organizations still lack identity security controls for these technologies. Machine identities now outnumber human identities by more than 80 to 1. Some would call this "unprecedented" — we prefer *overachiever in the field of firsts*.

Machine identities now outnumber human identities by more than 80 to 1.

The outlook on the geopolitical front is not much brighter. Last year, the Election Cyber Interference Threat Research Report warned that state-sponsored attackers would step up the use of AI in their disruptive operations against the U.S. and its allies. Nation states aren't just sponsoring these attacks; they're joining forces with cybercriminal organizations to ramp up cyber espionage and disinformation. They're hitting businesses, critical infrastructure and even the financial world, including a recent $1.5B crypto heist from ByBit. In December, the U.S. confirmed that Chinese government hackers gained remote access to the Treasury in what it described as a "major cybersecurity incident."

Executive Overview

AI has captured the world's imagination. But, as philosopher Paul Virilio once said, "When you invent the ship, you also invent the shipwreck." The same AI that can protect can also attack. It can detect vulnerabilities — and exploit them.

In the race to adopt AI, organizations are also inadvertently creating a surge of unmanaged and unsecured machine identities that overburdened teams don't have the visibility to manage. The privileged access of AI agents represents an entirely new threat vector that existing security models aren't built to handle. To stay resilient in this "overachieving" identity threat landscape, we can't wait for someone else to take the wheel. We must own our identity risk strategy and modernize our approach so we can adapt, respond and recover.

If you were already buckled up, maybe also bite down. What a time to be alive.

Here's what you'll find in this year's report:

1. AI's potential to be an identity-centric threat trifecta.
2. The shocking surge of machine identities, the scope of human identities with unsecured privileged access and the unique challenges both present for the enterprise.
3. The emergence of identity silos and how they undermine business resiliency.

Protecting sensitive and confidential data from breaches or leaks is paramount to maintaining trust and operational resiliency. As always, we'll dig into the data to highlight what's evolving — and share the steps you can take now to help your organization make the right kind of cybersecurity history.

Sincerely,

Clarence Hinton
Chief Strategy Officer, CyberArk

Clarence Hinton
Chief Strategy Officer

AI has captured the world's imagination. But, as philosopher Paul Virilio once said, "When you invent the ship, you also invent the shipwreck."

At a Glance

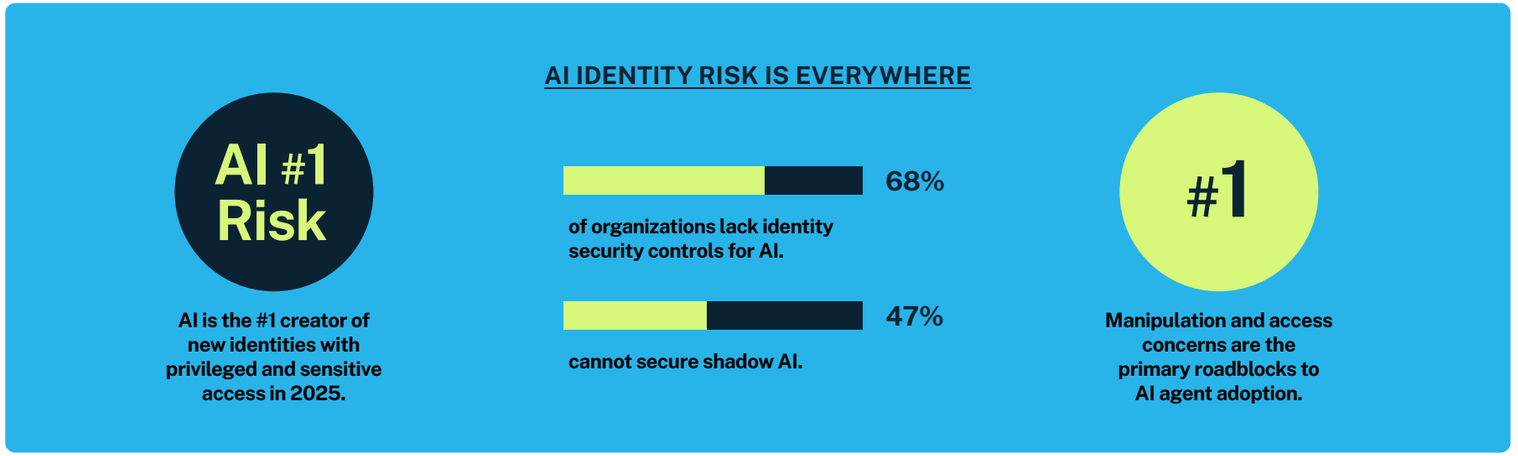

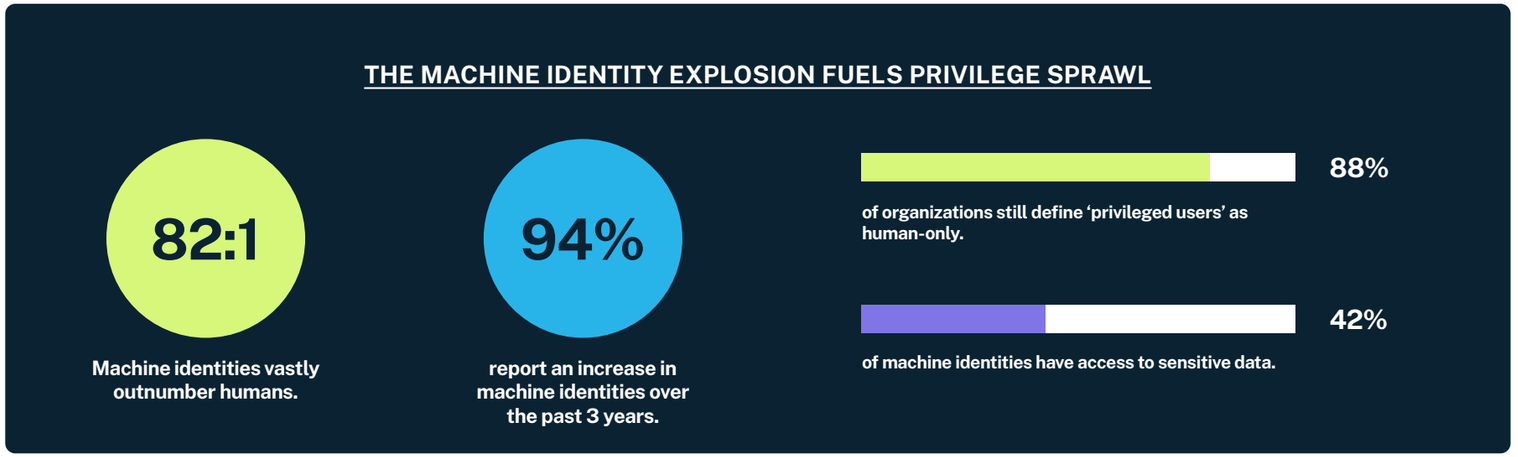

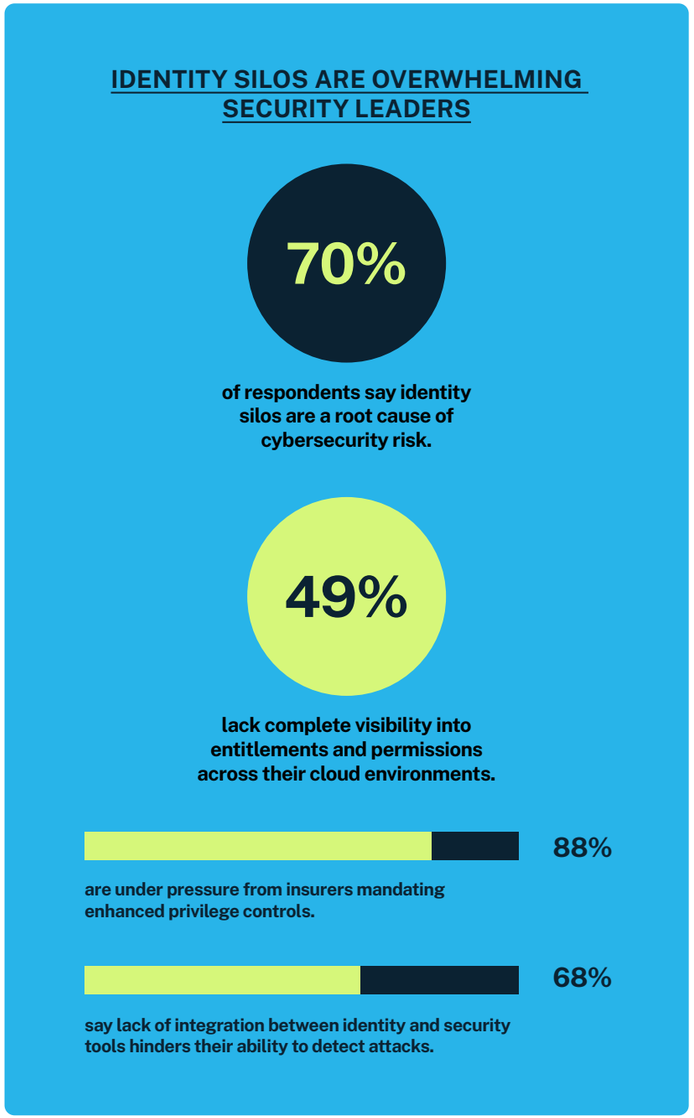

Figure 1. Key trends highlighting the impact of AI, machine identity and silos on identity risk (n=2,600).

The AI Trifecta: Attacker, Defender and Identity Risk

CYBERARK®

The AI Trifecta: Attacker, Defender and Identity Risk

In 2025, we'd be hard-pressed to find a place where AI has not relieved humans of manual and repetitive processes. It now regulates our grid, monitors our crops, directs traffic, and strengthens our cybersecurity arsenal. Our survey found that 94% of respondents (Figure 2) use AI and LLM processes to enhance their overall identity security strategies. Figure 3 shows that 61% are considering using AI to secure both human and machine identities in the next 12 months. Unfortunately, bad actors have had a head start using AI to make their attacks faster, smarter and harder to stop.

In addition, our report found that AI and LLMs are expected to drive the creation of the most new identities with privileged and sensitive access in 2025. This means that organizations must now secure the AI systems they deploy — and the new identities those systems create. Essentially, we must now manage AI as a weapon that can break into our systems; AI as a defender that secures our systems; and now, AI itself as a system we must secure.

We kick off this year's report by taking a closer look at how these three dimensions of AI triangulate the pressure on security teams.

Clunky, error-filled spam — and other things we miss

In the last 12 months, phishing has remained the leading cause of identity-related breaches. What's changed is the AI-driven scale, sophistication and success rate of these attacks.

Attackers can send AI-generated phishing emails that are highly personalized, context-aware and nearly indistinguishable from legitimate senders. They can use AI to analyze public data, mimic tone and formatting and adapt messaging in real time — making it easier to deceive even security-savvy users. And because AI can automate and coordinate outreach across email, chat and voice channels, social engineering campaigns are more convincing than ever before.

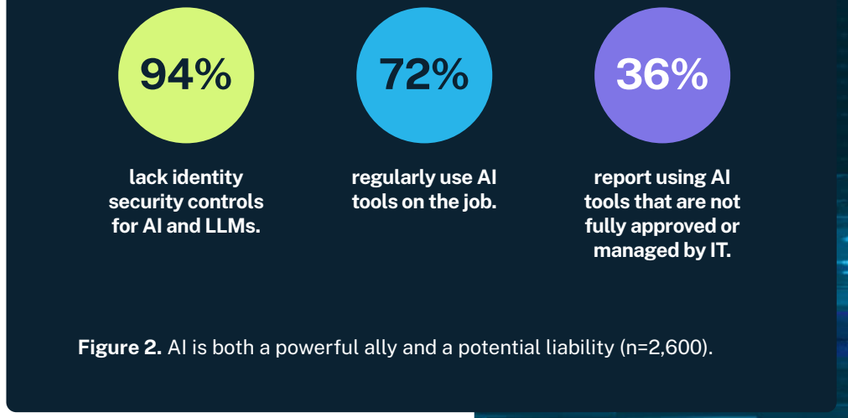

Figure 2. AI is both a powerful ally and a potential liability (n=2,600).

Figure 3. Top processes organizations are considering for securing identities with AI (n=2,600).

WHAT WE ASKED

What are your organization's **primary use cases** for AI and LLM applications? (Multi-select question)

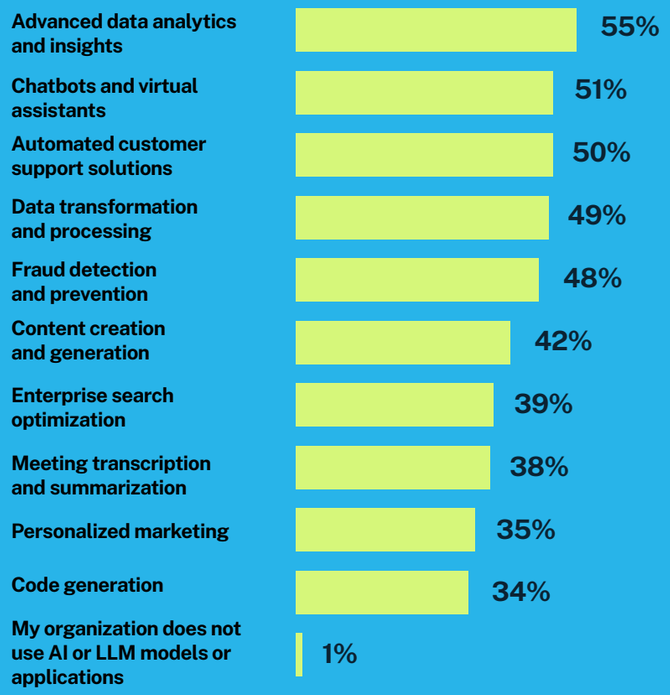

Figure 4. Use cases and LLM applications (n=2,600).

The AI Trifecta: Attacker, Defender and Identity Risk

AI-generated phishing then becomes an ultra-effective entry point for attackers who want to harvest credentials, escalate privileges and fast-track the exploitation of vulnerable applications, compromised privileged access and credential-based attacks. Nine out of 10 organizations reported experiencing a successful breach of this nature. Over three-quarters of respondents reported falling prey to successful phishing attacks (including AI-driven deepfake scams) within their organizations — and more than half of these fell victim multiple times.

Case in point: In February, scammers targeted prominent Italian business figures, including Giorgio Armani, using AI to mimic the voice of Guido Crosetto, Italy's Defense Minister. The fraudsters requested financial assistance under the guise of freeing kidnapped journalists, leading at least one victim to transfer €1 million to a Hong Kong bank account.

Identity security's new clutch player

For security teams, AI can reduce response times from hours to seconds. As it has no pesky human needs, it can ceaselessly analyze historical attack patterns, predict what's next, prioritize vulnerabilities and automatically shut down threats. Security operations centers (SOCs) can use AI to sift through mountains of identity-related threat data in real-time — not to replace human analysts, but to augment them.

AI also handles time-consuming, repetitive tasks and surfaces useful insights, allowing security teams to focus on bigger threats and make smarter, more strategic decisions. When paired with security orchestrations, automation and response (SOAR) systems, this human-AI collaboration can make incident responses more efficient and adaptive. In Figure 4, 55% of organizations say they use AI for advanced analytics and anomaly detection. Respondents cite AI as one of the most impactful tools for reducing identity-related threats in 2025.

AI handles time-consuming, repetitive tasks and surfaces useful insights, allowing security teams to focus on bigger threats and make smarter, more strategic decisions.

Meet your new sidekick/supervillain

But as AI-driven cybersecurity becomes a frontline defense strategy, securing the AI systems — including their machine identities — becomes just as critical. AI's reliance on vast amounts of data increases the risk of breaches, misuse and unauthorized access. Figure 5 shows that 82% of organizations know that using AI models opens access to sensitive data and creates cyber risks.

In the wrong hands, AI models can be manipulated into executing database queries, running external API calls or even accessing networked machines. Studies show that attackers are finding new ways to "jailbreak" (manipulate LLMs into secretly extracting and sending users' personal information, such as names, IDs, email addresses, payment details, etc.) with nearly 100% success rates on various models.

Jailbreaking AI models isn't just a theoretical exercise — it's a growing security concern as organizations rush to deploy AI without fully understanding its ramifications. Incidentally, that's why CyberArk's new FuzzyAI tool is making waves — it has successfully jailbroken every model it has tested. As an open-source project, now available on GitHub, it can help organizations and researchers systematically identify and fix AI security gaps before attackers exploit them.

Shadow AI: No one approved it. Everyone's using it.

Enterprises are using multiple approaches when hosting their AI tools, often adopting leading global LLM AI models (such as OpenAI, Google, Amazon Bedrock and Meta AI), coupling public training datasets with proprietary enterprise data to train the AI to solve problems. While 64% say that all of their organization's AI tools are approved and managed by IT, there are knowledge gaps. Almost half (47%) tell us that their organization is unable to secure and manage all of the "shadow AI" tools that are in use (Figure 5).

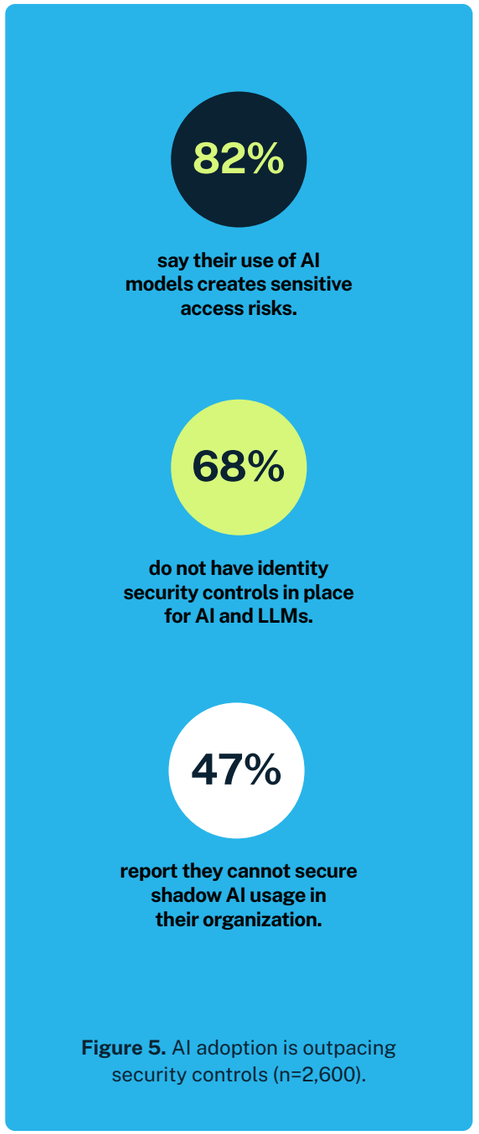

Figure 5. AI adoption is outpacing security controls (n=2,600).

The AI Trifecta: Attacker, Defender and Identity Risk

In many companies, the use of AI has drifted outside the purview of IT or security teams. Shadow AI, or employees or departments using AI applications, models or AI-powered features without official approval, is on the rise. Our report found that 36% of respondents report using AI tools that are not fully approved or managed by IT, leading to shadow AI risks.

Unlike shadow IT, shadow AI can be even harder to detect; AI capabilities are often embedded invisibly into approved software, meaning that organizations may not know which AI tools are processing company data. This is a problem. Let's say an employee inadvertently submits proprietary or personal data to an AI service — it could be stored or logged outside the company. Or a finance team might unknowingly expose an API key or confidential records by including them in an AI prompt, which then gets logged by the AI provider. Without the right controls in place to protect AI inputs, decision-making or training data, attackers can corrupt any one of these processes using injection attacks, model poisoning or any number of attacks du jour to bias AI behavior.

Compounding this risk is the diverse landscape where AI models live (Figure 6). As AI deployments expand and oversight thins out, organizations may be innovating beyond what they can secure. Whether hosted on-premises or in the cloud, companies must now decide how they'll secure AI training, rollout and operationalization. Without policies and monitoring, shadow AI piles on the security and regulatory pain, exposing companies to compliance violations, data leaks and other no good, very bad times.

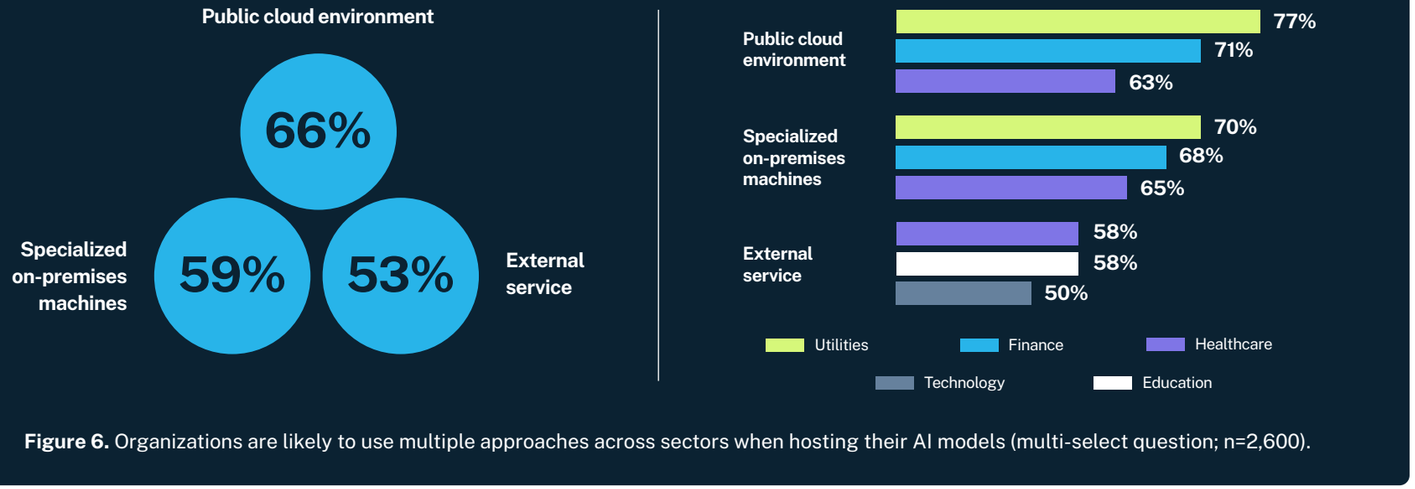

Figure 6. Organizations are likely to use multiple approaches across sectors when hosting their AI models (multi-select question; n=2,600).

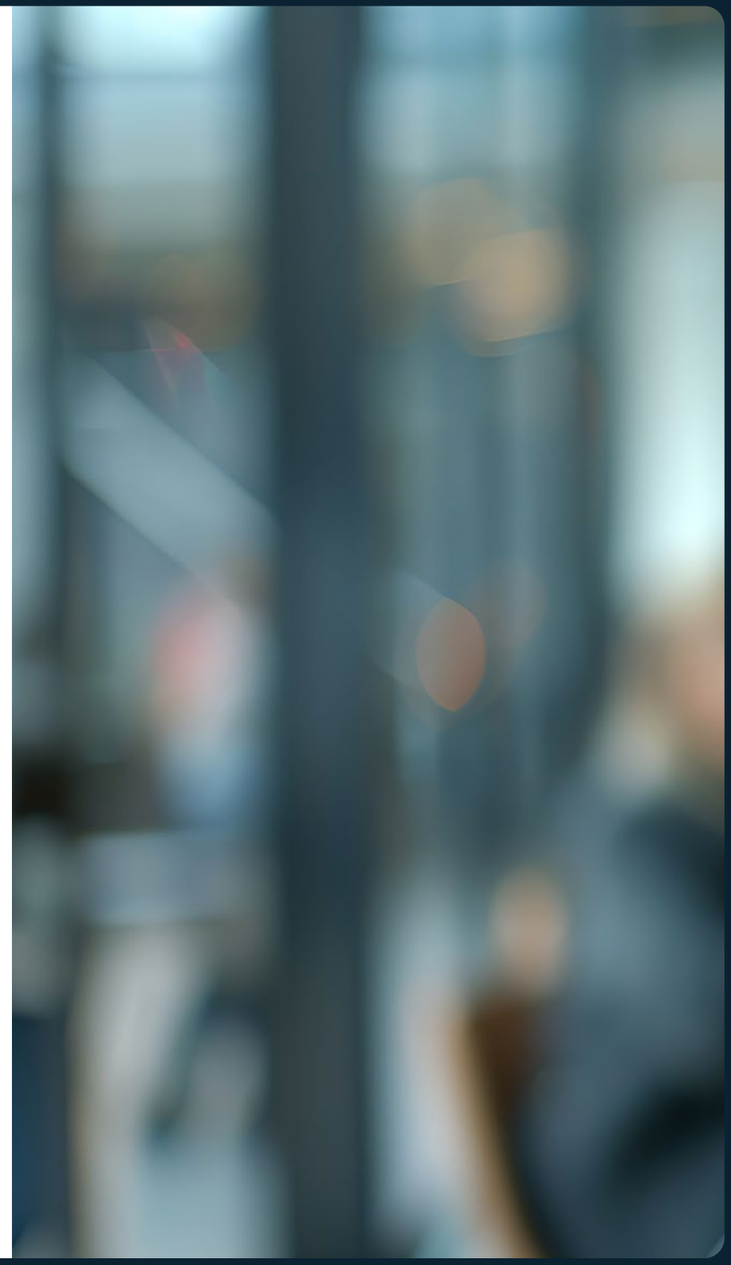

Future shock: The emergence of AI agents

In case securing AI was not enough of a challenge, AI agents can be your new endurance sport. AI agents introduce an entirely new layer of complexity — as dynamic, machine identities with human-like autonomy. Rather than just an information-processing content tool, AI agents are machine identities that perceive, reason and act based on defined goals. Now imagine securing thousands or even millions of these entities: ensuring proper authentication with systems (and other agents), regulating their privileged access to sensitive data and maintaining strict lifecycle control to avoid rogue agents with lingering permissions across diverse systems and geographies — you get the idea. If not properly controlled and monitored at scale, a lot can go wrong.

The attack surface of an AI agent spans three critical layers:

1. **Infrastructure layer:** Credentials on the system where the agent resides.
2. **Access layer:** Privileges or entitlements associated with the agent.
3. **Model layer:** The AI itself, which can be tricked or hijacked.

While the first two reflect familiar challenges in securing machine identities, the third introduces unique risks tied to the AI's non-deterministic behavior and ability to reason — which lends itself to, well, misbehavior. Without guardrails, AI agents at the model layer can be manipulated into executing malicious commands, leaking data, escalating privileges or granting unauthorized access faster than a human ever could. Traditional IAM systems aren't equipped to handle the authentication, authorization and monitoring protocols required for thousands (or millions) of these intelligent entities.

While not yet widely deployed, experts predict that by 2028, AI agents will be making at least 15% of day-to-day work decisions. The benefits are undeniable — but without preparation, organizations risk racking up hefty security debt (Figure 7). Duct tape fixes will not fly here. Organizations will need rock-solid backend infrastructure. Best practices include:

- **Privileged access controls** that ensure AI identities aren't exploited for unauthorized access.
- **Governance that allows for continuous visibility** into the activities of AI-driven machine identities.
- **Codes of conduct that align AI use with responsible deployment** and regulatory compliance. CyberArk supports model context protocols (MCPs) — early AI design standards that ensure context-aware, interoperable and secure AI agent workflow across the enterprise.

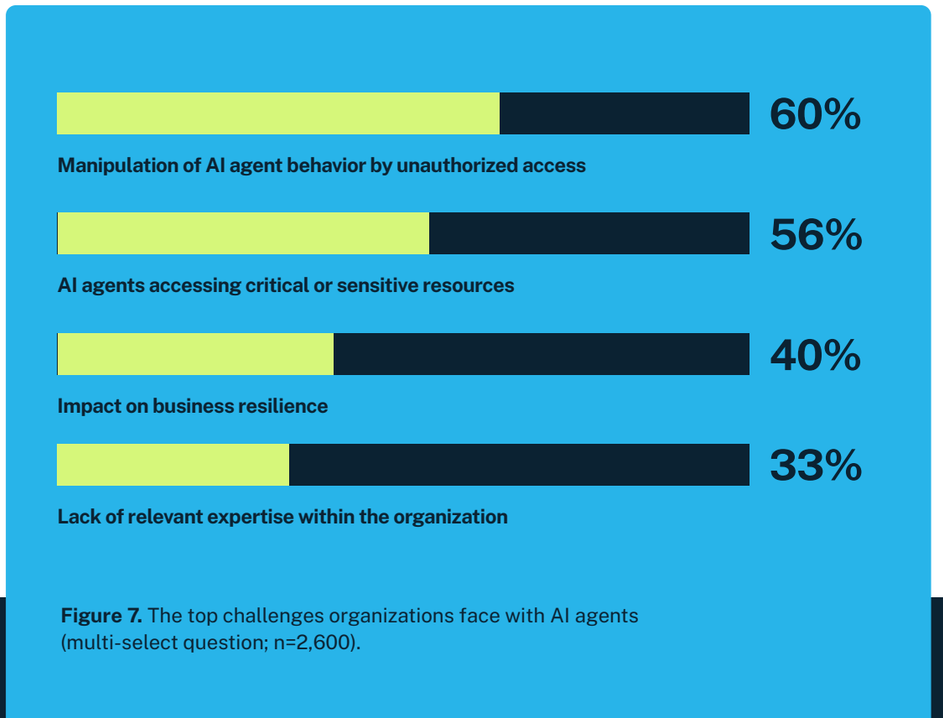

Figure 7. The top challenges organizations face with AI agents (multi-select question; n=2,600).

CyberArk Insight

AI is now an integral part of how we do business. The path forward must include proactive measures around how these AI-driven solutions and services are developed, deployed and used.

We recommend a three-tiered approach:

- **Secure Development:** Developers who write code and create models that help AI systems must follow strong security practices that ensure training data is clean and representative.
- **Secure Deployment:** When an AI system is moved from the testing phase to an operational environment where it interacts with users or other systems, the operational environment must adhere to strict identity security measures to protect it from tampering, unauthorized access and manipulation.
- **Secure Use:** To ensure attackers can't leverage user access, we must integrate AI into identity security models — not as an afterthought but as part of a holistic strategy.

Secure Identity = Secure AI Agents

Machines that behave like humans require both human and machine security controls. Each agent must be uniquely identified, authenticated and governed, just like a human user — but also with the added rigor required for machine-scale operations. Without these dual-layered protections in place, we risk repeating the identity chaos of early RPA implementations where impersonation, over-privileged access and lack of governance left the door wide open to exploitation.

Machine Identities: The Sprawl Awakens

CYBERARK®

Machine Identities: The Sprawl Awakens

Though invisible to the human eye (or audit log), machine identities quietly keep digital infrastructures humming. Without them, our devices, clouds, servers, applications, containers and software processes would be as secure as a lock with the key taped next to it. However, every day, new cloud workloads, AI/ML services, automated processes and interconnected systems come online, requiring new machine identities for authentication. Not surprising that the volume, variety and velocity of machine identity growth show no signs of slowing down: 94% of organizations report an increase in machine identities over the past three years. Enterprises are operating amidst a staggering proliferation of machine identities: **more than 80 machine identities for every human identity** — nearly doubling since this data was first reported in 2022 (Figure 8). This ratio grows as high as 96:1 for the finance sector and 100:1 in the U.K. — a hair-raising challenge for human security teams.

Over half of survey respondents (54%) predict AI and LLM tool adoption will continue to drive the creation of machine identities with privileged and sensitive access. As the machine identities often have a direct channel to privileged resources, the attack surface isn't widening — it's exploding.

Dust off your vision boards, folks. Organizations have not improved their understanding of 'privileged user' — **88% still define 'privileged user' as human-only**, up from 61% in 2024. In Figure 8, we found that 42% of machine identities — and 68% of bots and machine accounts — have access to sensitive data (compared with 37% of human users). Only 12% (Figure 9) consider machine identities to be 'privileged users.' To fix this gap, we need to move beyond the human-centric definition of 'user' and redefine 'privilege' for machines. Privileged access for non-human identities may look different — but it must be just as visible, managed and governed.

As a rule of thumb, **machine identity security (MIS) secures all non-human identities that matter** — from bots and service accounts to scripts, cloud workloads and AI agents. As these entities gain autonomy and access, MIS is no longer just an IT hygiene issue. It's a core pillar of enterprise security.

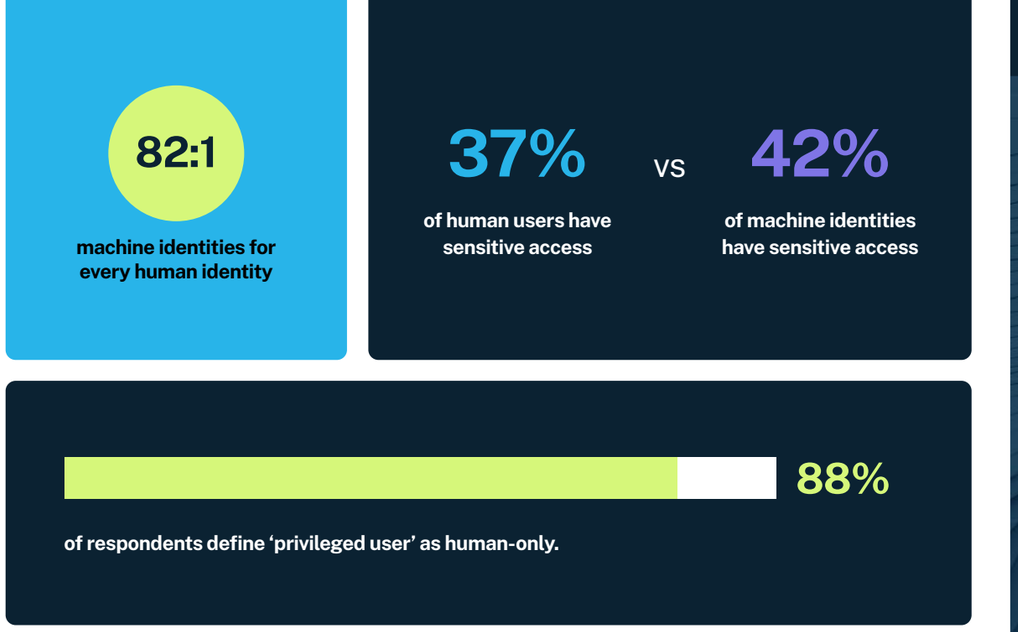

Figure 8. Machine identities outpace humans in volume and access risk (n=2,600).

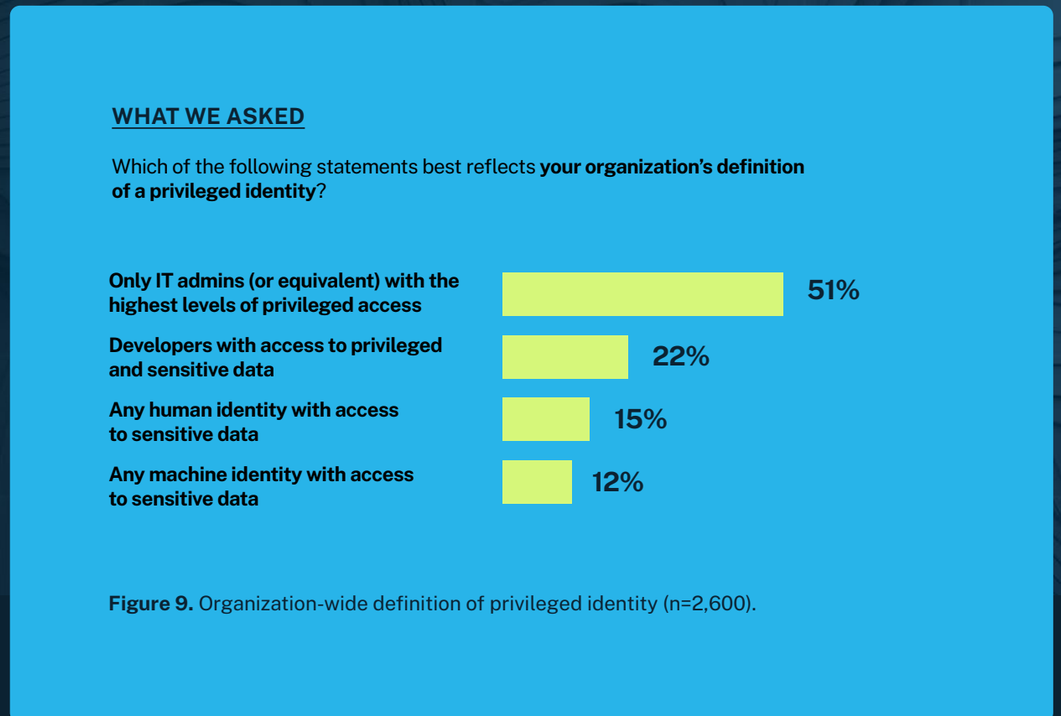

Figure 9. Organization-wide definition of privileged identity (n=2,600).

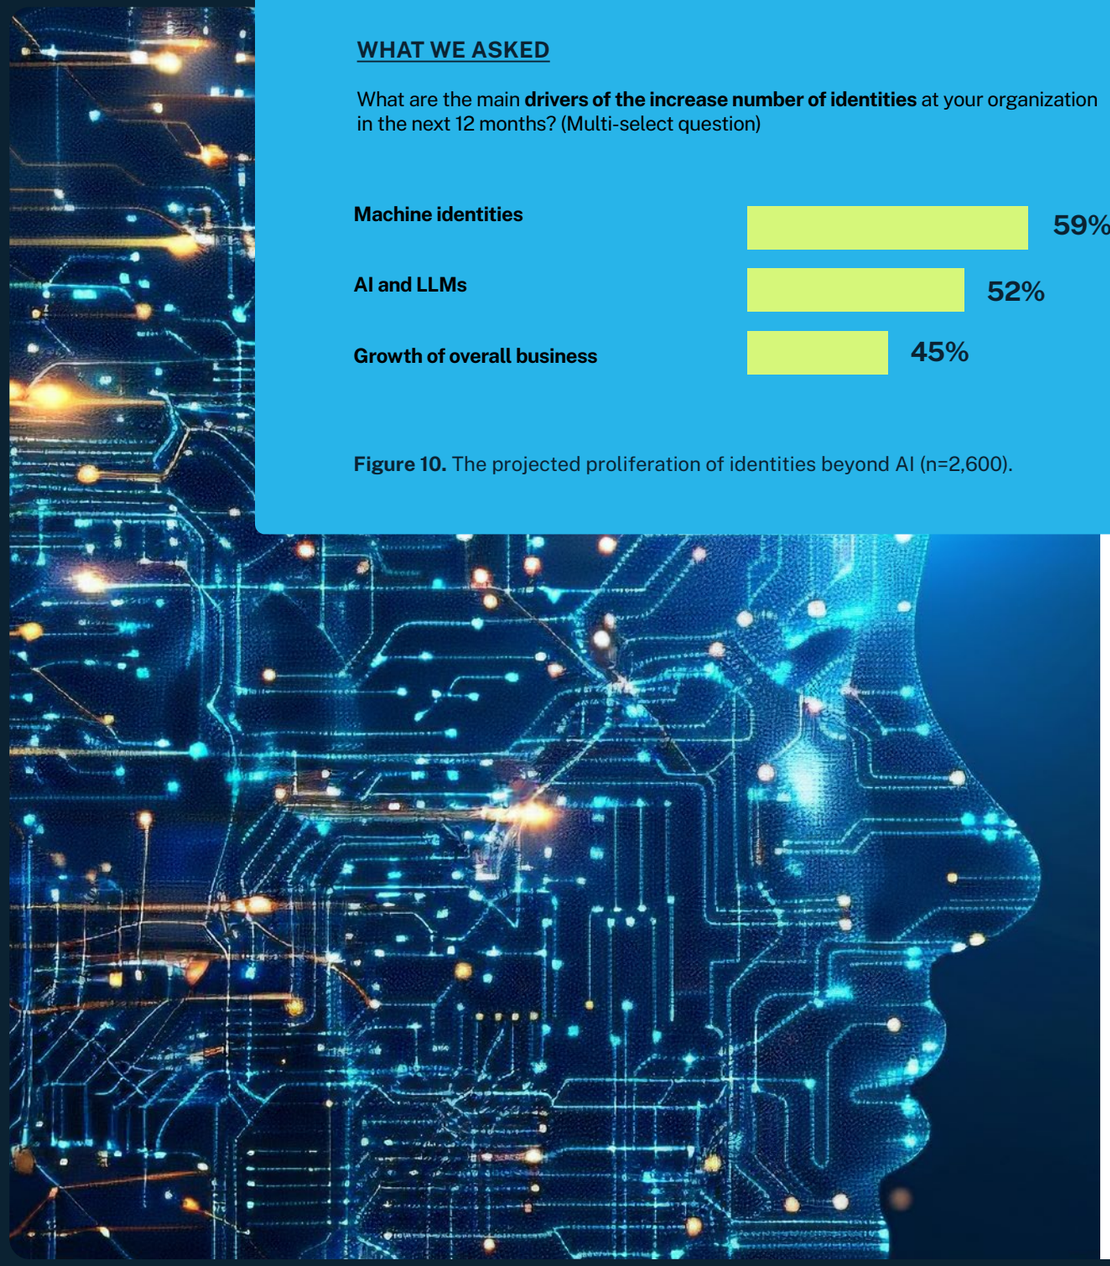

WHAT WE ASKED

What are the main **drivers of the increase number of identities** at your organization in the next 12 months? (Multi-select question)

Figure 10. The projected proliferation of identities beyond AI (n=2,600).

Any (unsecured) port in a storm

Cybercriminals aren't picky eaters — any identity is fair game. And machine identities, the quietest users in your environment, are often the most vulnerable. If a machine identity is left with excessive or standing privileges, an attacker can find a pathway to assume the identity of that machine account. They can register their own devices or apps on corporate identity systems to persist their access without detection. They can extract API keys, certificates, or secrets from code repositories, logs, or configuration files. They can hijack abandoned service accounts. They have lots of options.

In line with these findings, machine identities emerged as this year's top perceived identity risk in terms of the most unmanaged, unknown identities across the IT environment, with 33% of respondents admittedly not controlling the risk by only applying 'privileged' to human and not machine identities.

As we noted in our intro, security teams are fielding mixed signals from regulators. In the U.S., recent shifts suggest a move toward a more hands-off approach to AI oversight, raising questions about whether deregulation will fuel innovation — or simply widen the blast radius. While some federal guidance still recognizes machine identities as critical gatekeepers, much of the momentum now lies outside the U.S.

In the Asia-Pacific region, for example, Australia's Cyber Security Act 2024 ushered in the country's first broadly applicable cybersecurity-specific law. The Act reflects a growing trend: governments moving to tighten identity controls and codify how AI systems are secured and governed. Meanwhile, the European Union is pressing forward with its AI Act, which means that companies operating in the EU must closely monitor and document their AI models to meet rigorous new standards or face substantial fines.

The proliferation of identities isn't just an AI problem — it's a broader, systemic shift across modern infrastructure that will require careful planning. Over the next 12 months, 59% of respondents predict that machine identities (from cloud workloads to app credentials and automated services) will be a leading driver behind identity growth, outpacing even AI and LLMs (Figure 10).

Ignorance is risk: The human side of identity sprawl

Pining for the good ol' days? Rest assured, human identities remain a familiar headache, with identity pains being less about proliferation and more about privilege. Across the three major cloud platforms alone, there are over 40,000 distinct privileges. Users in the broader workforce need access to dozens of SaaS applications, multiple cloud platforms and AI tools across the enterprise — and they're logging in from everywhere. Unmanaged endpoints represent significant security blind spots that are challenging to monitor and protect effectively. With little to no visibility, IT and security teams aren't just unaware of the potential risks — they're are also unable to enforce them.

CyberArk Insight

Organizations should put their attention on centralizing solutions and adopting an identity security strategy that recognizes all identities — workforce, IT, developer and machine — pose security risks at every stage in the lifecycle, from creation to consumption. Unfortunately, there are no magic bullets, just a lot of security homework. Some food for thought:

- **Secure every identity with controls that can monitor, analyze and audit user and admin sessions to detect threats.** Privileged access management (PAM) and least privilege controls are critical to ensure that every identity has only the necessary access rights required for their role.
- **Reassess your definition of privileged user** to include every machine, service account and workload.
- **Make sure you know what you're managing.** Security teams need to discover secrets in cloud service providers' built-in (native) secrets stores.
- **Automate the certificate lifecycle** across all application and workload types from the initial request to installation. This results in fewer errors and avoids squandering precious security team resources.
- **Reduce unmanaged endpoint risks** by adopting secure browsing solutions.
- **Leverage different approaches,** like just-in-time or dynamic secrets rotation, strong authentication and authorization mechanisms, and role-based access controls (RBAC).

THE PEOPLE BEHIND THE PRIVILEGED ACCESS

The biggest risk to an organization's security isn't just AI. The CyberArk 2024 Employee Risk Survey gathered insights from over 14,000 employees worldwide and shed light on just how risky our human work behaviors can be:

- 60% have used a personal device to access work-related apps, emails, or systems in the last 12 months.
- 36% use the same password for both personal and work accounts.
- 65% admit to bypassing security policies in the name of productivity.
- 40% habitually download customer data.
- 1 in 3 can alter sensitive or critical data.

CyberArk, CyberArk 2024 Employee Risk Survey, Dec. 2024.

Breaking Silos, Taking Names

CYBERARK®

Breaking Silos, Taking Names

For most enterprises, identity security didn't start out as part of the grand strategy. It was assembled brick by brick as organizations built out their technology stack. During the normal course of business — a merger here, a legacy system there — multiple groups ended up using independent systems and different technologies to achieve slightly diverse versions of the same-ish goal. So yeah, silos: great for farmers, deadly for business resiliency. Our survey found that 70% of respondents identify silos as a root cause of organizational risk. Factor in hybrid infrastructure and some unsupervised AI app usage (Figure 11), and it starts to feel less like a strategy and more like a trust fall.

Privileged access: More management, less mystery

This fragmentation has profound implications for tracking entitlements and permissions. While 94% of organizations use tools that automatically protect and monitor all cloud sessions, 68% say that the lack of integration of their identity and security tools hinders their efforts to detect attacks (as high as 84% in government organizations). Meanwhile, attackers aren't dealing with any of these roadblocks — they can be light on their feet and move seamlessly across environments.

Silos also make compliance more difficult and drive up premiums. Since their last cyber insurance renewal, 88% of respondents report that insurers are demanding stricter privilege controls, and 89% noted that cyber insurers are requiring stricter adherence to the principle of least privilege.

WHAT WE ASKED

Which of the following, if any, have created **identity silos** in your organization? (Multi-select question)

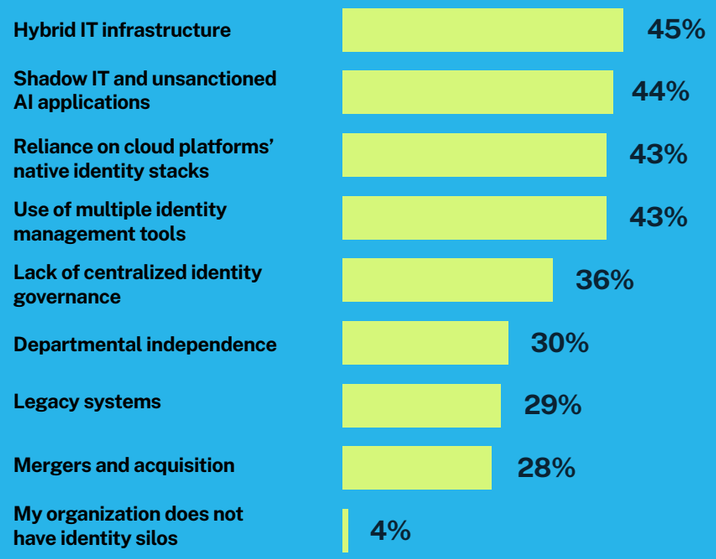

Figure 11. The causes of identity silos in organizations (n=2,600).

88% say they face stricter requirements from cyber insurance providers to implement privilege controls.

You can't secure what you can't see

Almost half (49%) of survey respondents report that their organization lacks full visibility into entitlements and permissions across their entire cloud environment (Figure 12).

Even where identity controls do exist, they're unevenly applied. Fewer than 40% report coverage for cloud infrastructure and workloads. Controls drop further for DevOps environments (35%), AI and LLMs (32%), and service accounts (23%) — despite these being some of the fastest-growing areas of risk (Figure 13).

IGA TIES THE WHOLE ROOM TOGETHER

Think of IGA and PAM as the bread and butter — the Iron Man and Jarvis — of your security infrastructure: separate; they're capable; combined, they're complete.

Identity governance and administration is critical to opening up visibility across the enterprise. Working in tandem with identity and access management (IAM) and privileged access management systems, IGA centralizes and automates identity management across on-premises, cloud, and hybrid environments. It helps ensure consistent identity policies, access controls, and least privileged access based on real-time risk assessment. This is critical for both Zero Trust alignment and for managing and securing those machine identities at scale.

IGA automation also improves oversight and eliminates human latency while helping organizations comply with a variety of government and industry regulations. This, in turn, can save your people from burnout.

WHAT WE ASKED

To what extent do you agree or disagree with the following? *My organization does not have full visibility of entitlements and permissions across my entire cloud environment.*

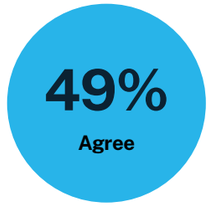

Figure 12. Respondents' level of visibility across their cloud environment (n=2,600).

WHAT WE ASKED

For which of the following environments and devices **does your organization have identity security controls** in place? (Multi-select question)

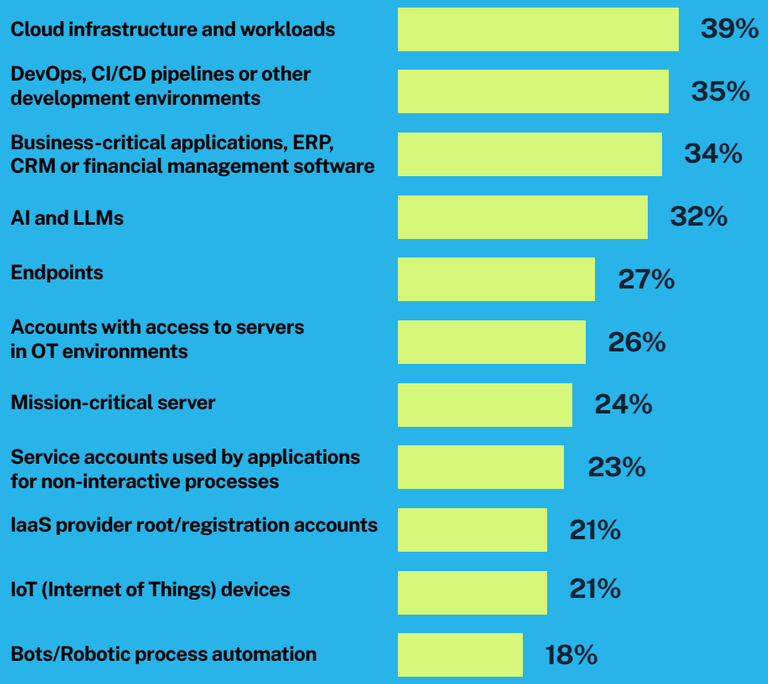

Figure 13. Identity security controls across different environments (n=2,600).

WHAT WE ASKED

What are your organization's **top strategic security priorities** in 2025? (Multi-select question)

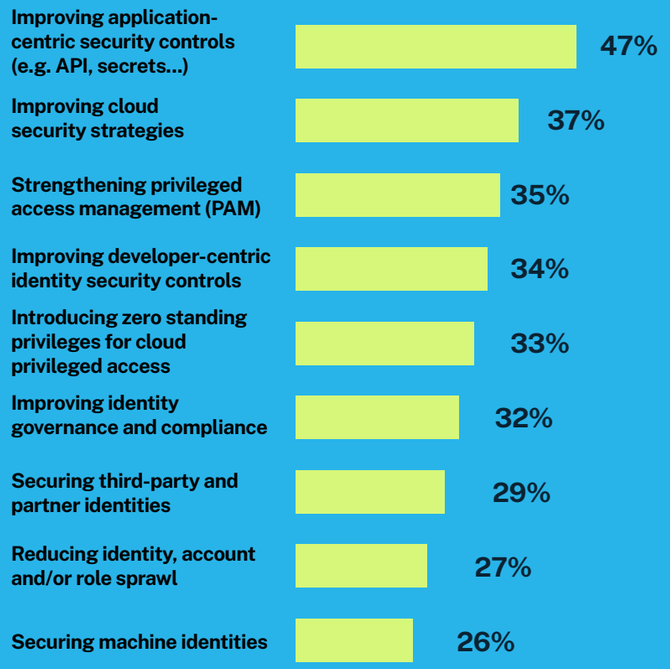

Figure 14. Identity security priorities for 2025 (n=2,600).

The top 6 strategic identity security investments for 2025

A majority — 87% of organizations — say they experienced at least two successful identity-centric breaches in the past 12 months, ranging from supply chain attacks and compromised privileged access to identity and credential theft. However, 75% of security professionals agree that business efficiencies are prioritized over strong cybersecurity in their organization.

Some good news: organizations recognize these challenges and are prioritizing critical priorities in the year ahead.

As shown in Figure 14, nearly half (47%) of organizations want better application-based security controls to protect unique environments, while 35% acknowledge that stronger privileged access management (PAM) controls are the way to go. And, given the inconsistent oversight across increasingly complex ecosystems, 32% plan to invest in identity governance and compliance (IGA).

75% of security professionals agree that business efficiencies are prioritized over strong cybersecurity in their organization.

Breaking Silos, Taking Names

CyberArk Insight

Addressing fragmented legacy solutions is the best way to strengthen your organization's overall posture and ultimately build a resilient enterprise.

- **Think like the attacker.** Keep up to date on modern threats so you stay clear-eyed about gaps in your controls.
- **Adopt a "build to protect mindset"** where every identity, resource and account is guarded with automation and the right level of intelligent privilege control from the moment it's created.
- **Streamline and automate IAM and identity security processes.** Manual processes cause delays and gaps in security that bad actors exploit.
- **Consolidate and centralize identity security tools** to improve operational efficiency, increase timeliness of security and simplify the disjointed identity processes.
- **Modernize with a faster, more adaptive IGA solution** built to tackle complex, multi-cloud use cases.

70% of respondents identify silos as a root cause of organizational risk.

Parting Thoughts

Parting Thoughts

Today's cybersecurity threats and AI buzz have become so pervasive that they often fade into background noise, but we cannot tune out this blaring siren.

Ransomware has become a funding mechanism for nation-states — and cybercrime is the most lucrative and scalable business model available.

Adversaries are embedding AI into their tactics, making attacks more scalable and efficient, while organizations are embedding AI into their workflows, creating new security blind spots. Soon, AI agents will be making decisions like humans and scaling like machines.

The sheer volume, variety and velocity of machine identities is already forcing organizations to rethink how they manage and secure them. Meanwhile, security teams are drowning under the weight of too many tools, too many alerts and too few resources to keep up.

At the center of all of this is identity. No matter the method, the goal of every attacker is to compromise an identity. That singular focus should simplify and reshape how you respond to threats this year, and beyond.

In 2025, the collective risks we face often feel existential — but we have powerful resources to help us meet the moment. AI may be rewriting the rules, but identity security contains the risks. Organizations need to secure every identity, human and machine, and deliver a cohesive end-to-end experience that matches real-world identity threats. With the right preparation, we can build defenses that overachieve, outthink, outpace and outlast whatever comes next.

To build business resilience, organizations need a practical, risk-based approach — grounded in identity security. We must:

- **Authenticate and secure** AI agents at scale.
- **Manage and limit** access to sensitive data.
- **Control** AI identity lifecycles to prevent rogue access.
- **Consolidate** security tools with experienced, trusted partners.

As AI agents take on more responsibilities and the boundaries of privileged access expand, this strategy allows us to effectively anticipate, withstand and recover from cybersecurity incidents without disrupting operations.

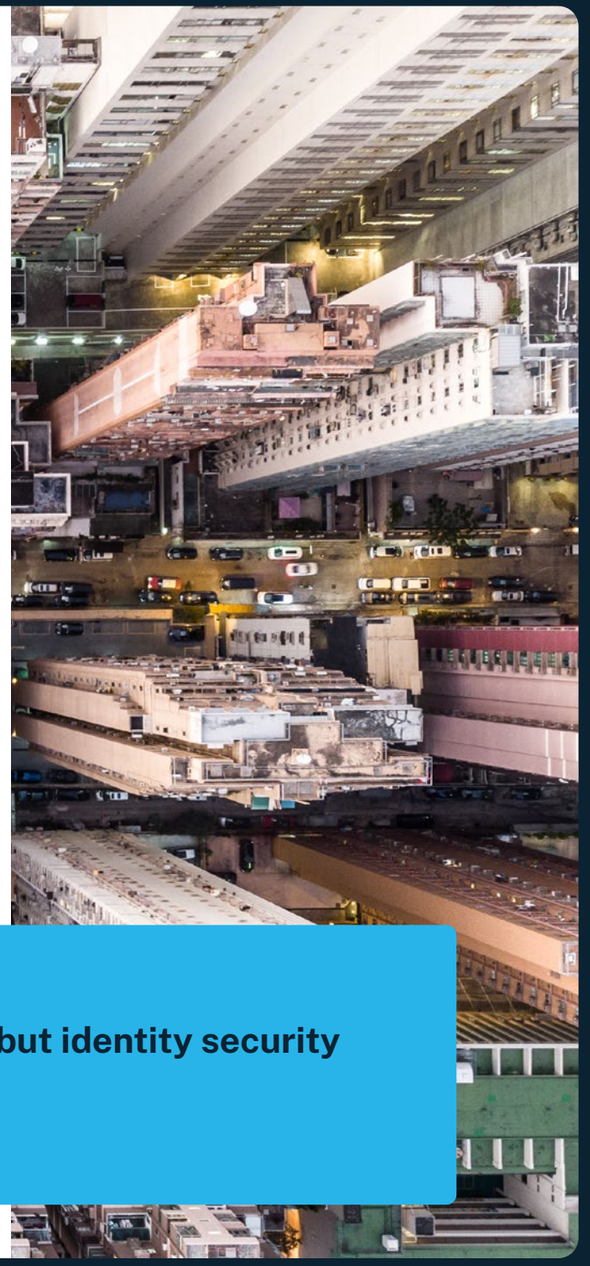

AI may be rewriting the rules, but identity security controls the risks.

Appendix

Appendix

Methodology and Demographics

Results from the CyberArk 2025 Identity Security Threat Landscape Report was fielded across private and public sector organizations between January and February 2025. It was conducted by B2B technology research partner Vanson Bourne amongst 2,600 cybersecurity decision makers based in Brazil, Canada, Mexico, the U.S., France, Germany, Italy, the Netherlands, South Africa, Spain, the U.K., UAE, Australia, India, Hong Kong, Israel, Japan, Saudi Arabia, Singapore and Taiwan.

RESPONDENTS BY GEOGRAPHY

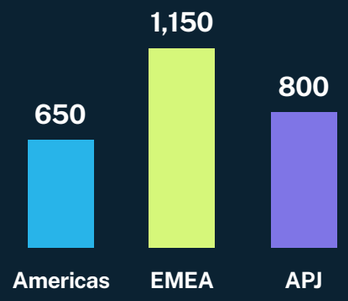

Figure 15. Breakdown of respondents by geography (n=2,600).

RESPONDENTS BY SECTOR

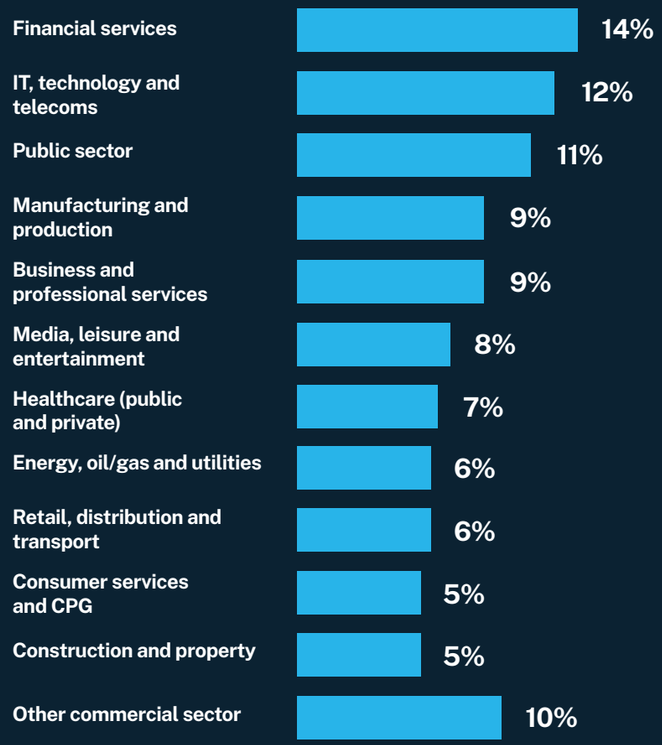

Figure 16. Breakdown of respondents by sector (n=2,600).

RESPONDENTS BY COMPANY REVENUE

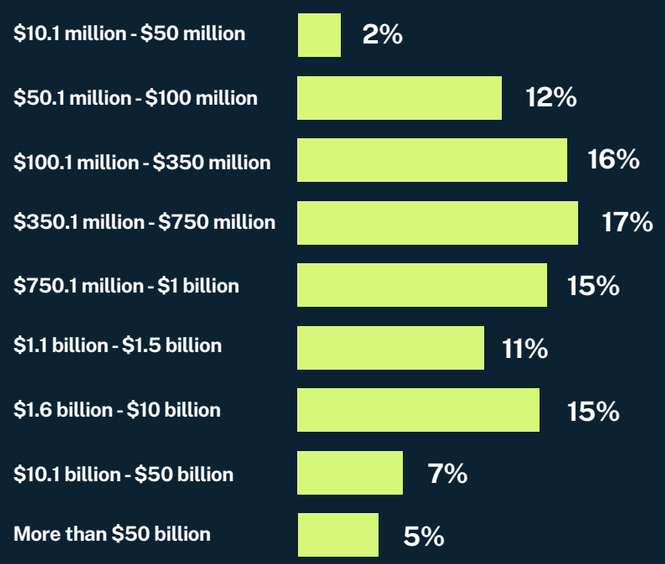

Figure 17. The 2024 global annual revenue (USD) reported by respondent organizations (n=2,600).

RESPONDENTS BY DEPARTMENT

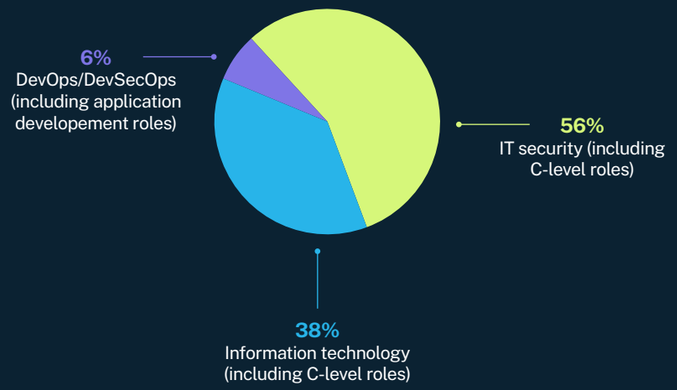

Figure 18. Breakdown of respondents by department (n=2,600).

RESPONDENTS BY TITLE

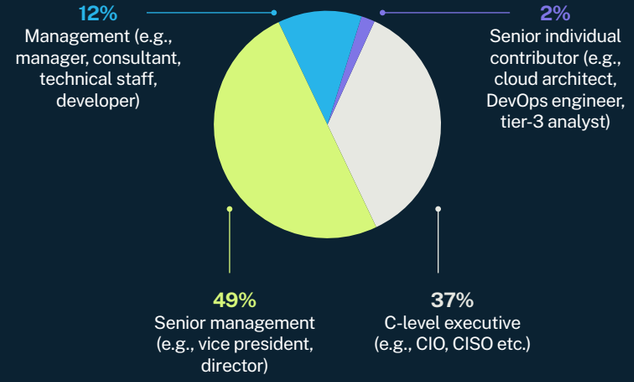

Figure 19. Breakdown of respondents by job title (n=2,600).

RESPONDENTS RESPONSIBLE FOR IDENTITY SECURITY

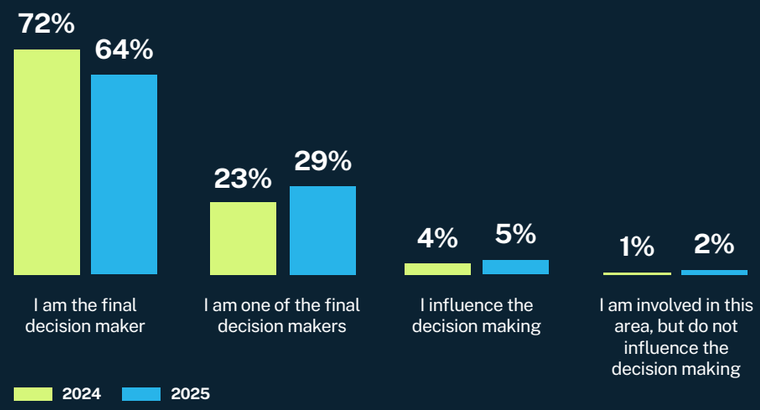

Figure 20. Breakdown of respondents by identity security responsibility (n=2,600).

SECURITYMATTERS

CYBERARK® Research

Global investments in AI is massive. The identity security risks? Even bigger. Get best practices to guide your identity security journey.

Learn More

About CyberArk

CyberArk (NASDAQ: CYBR) is the global leader in identity security, trusted by organizations around the world to secure human and machine identities in the modern enterprise. CyberArk's AI-powered Identity Security Platform applies intelligent privilege controls to every identity with continuous threat prevention, detection and response across the identity lifecycle. With CyberArk, organizations can reduce operational and security risks by enabling zero trust and least privilege with complete visibility, empowering all users and identities, including workforce, IT, developers and machines, to securely access any resource, located anywhere, from everywhere. Learn more at www.cyberark.com.

CYBERARK®