C1

2026

Future of Identity Report

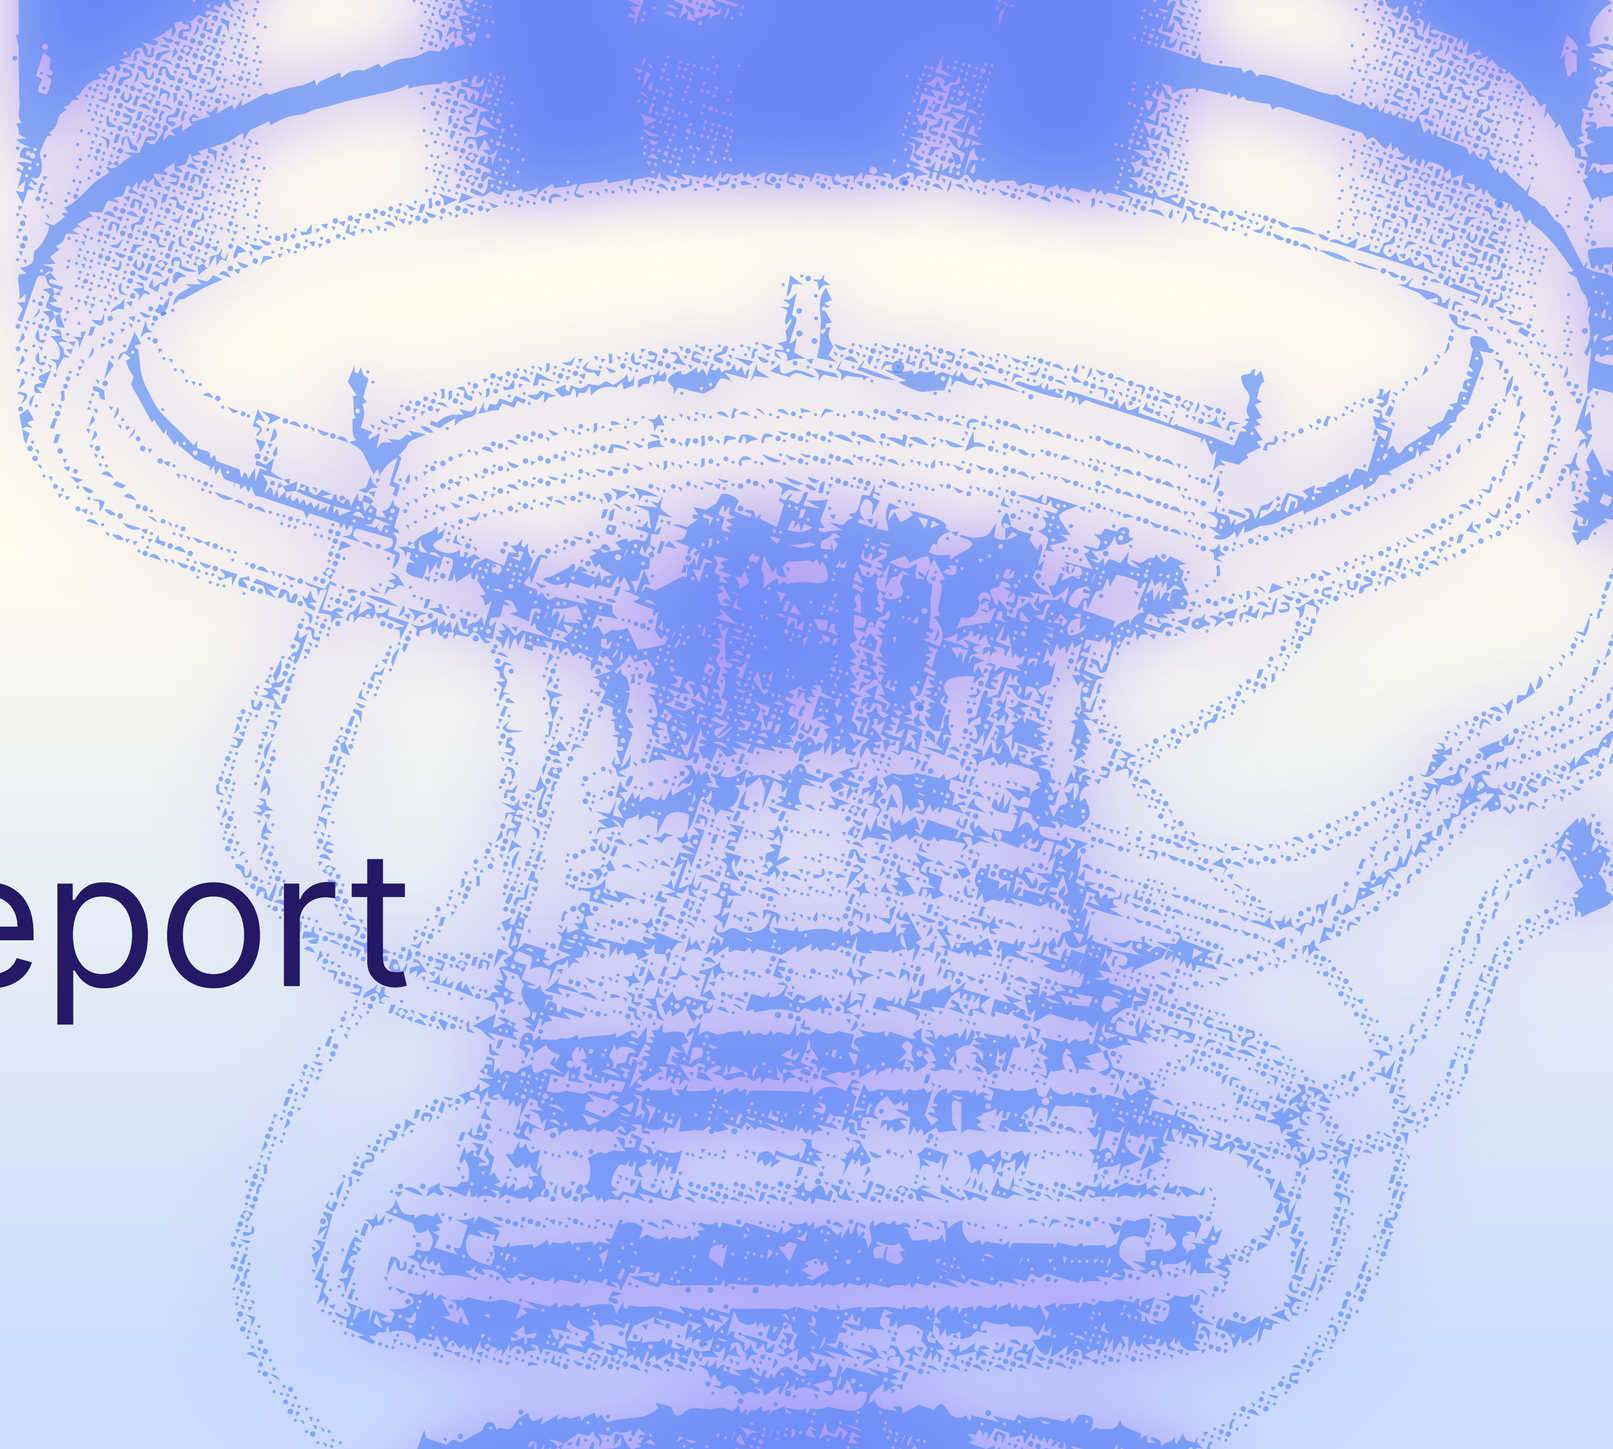

Identity is the unlock for the agentic enterprise.

AI agents are driving unprecedented productivity. Identity is the layer that enables them to operate safely, securely, and at scale.

Table of contents

Executive Summary

For the third annual Future of Identity report, ConductorOne surveyed 508 IT and security leaders at U.S. organizations with more than 1,000 employees. The results reveal a fundamental shift in how work happens across the enterprise. Autonomous software actors are no longer experimental tools. They are operational entities with real access, authority, and impact.

Last year, 96% of organizations said they planned to operationalize AI agents. Today, 95% report agents already performing autonomous IT or security tasks. The industry did not gradually evolve into autonomy. It crossed a threshold. More actors are operating inside enterprise environments, making decisions and executing workflows at machine speed.

We define the **agentic enterprise** as an environment where autonomous software entities perform work alongside humans, operating through identity and delegated access. In this model, identity becomes the system that enforces who and what can act, under what conditions, and for how long.

As autonomy accelerates, security teams face a growing imbalance between machine-speed execution and governance models built for human workflows. The agentic enterprise requires identity systems that can evaluate context continuously and enforce policy dynamically. Without that shift, organizations cannot safely scale AI-driven productivity.

What we found

AI agents are operational.

95%

95% of organizations report AI agents performing at least one IT or security task autonomously.
Identity governance is no longer about user logins alone. IGA now must govern autonomous execution.

Investment is accelerating.

91%

91% report increased IAM investment driven by AI adoption, signaling that identity is now central to how organizations enable and secure AI agents.

Security tooling is lagging behind autonomous adoption.

More than a quarter of respondents say existing identity tools were not designed for AI agents. As autonomous actors expand across environments, organizations are relying on governance models built for human workflows, creating new gaps between automation and security.

Access-related breaches remain pervasive.

80%

80% experienced at least one access-related attack or breach in the past year, resulting in increased investment in over 90% of organizations.

Companies understand identity risk and are putting more budget towards reducing it.

Non-human identities are outpacing visibility.

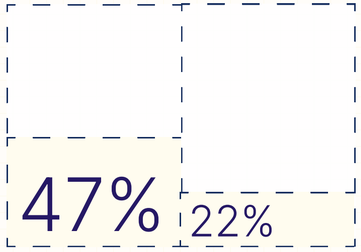

47% say NHIs outnumber human users, yet only 22% report full visibility. The attack surface is expanding faster than governance.

The defining challenge of identity security in 2026 is building governance models that can keep pace with autonomous systems.

The Agentic Enterprise Arrives

From pilot to production

91%

AI adoption has moved from experimentation to execution. Just one year ago, 89% said they planned to implement AI agents within their security teams, signaling that what was on the roadmap in 2025 has become operational reality in 2026.

The agentic enterprise is defined by the ability to orchestrate thousands of autonomous actors safely. Identity becomes the control plane, determining which agents can act, what they can access, and how their actions are governed in real time.

Agentic Maturity

Actively operating as an agentic enterprise: **31%**
Early workflows in production: **32%**
Piloting use cases: **25%**

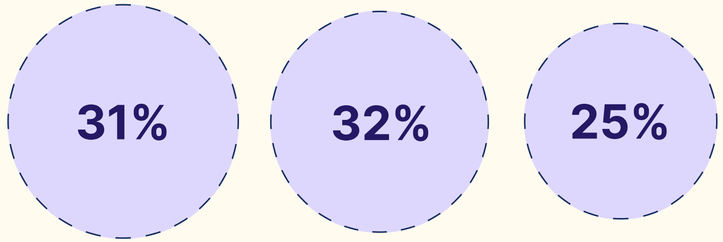

This marks a structural shift: identity systems need to be governing systems performing work, not just people.

Investment follows urgency

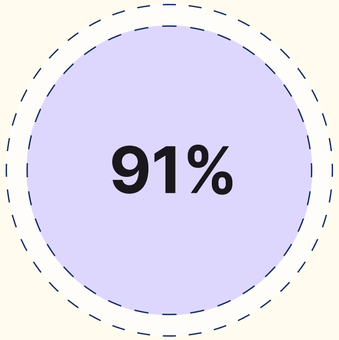

91% increased IAM investment due to AI adoption.

Organizations are increasing identity investment as governance gaps turn into real operational risk. Budget growth signals that identity is evolving into a real-time security layer for AI adoption.

The Identity Security Snapshot

Breaches aren’t slowing down

80% of organizations experienced at least one identity-related breach in the last year, with 82% the year prior.Despite rising awareness and investment, access risk remains persistent.

Automation increases speed, but governance models often lack real-time context. Access decisions still rely on periodic reviews, static roles, or manual approvals that were designed for human users. As AI agents and automated workflows act continuously, small gaps in policy or visibility can compound into large-scale exposure.

When identity governance loses intent, meaning who or what should have access, for what purpose, and for how long, autonomous systems can amplify mistakes at machine speed.

TOP ATTACK VECTORS

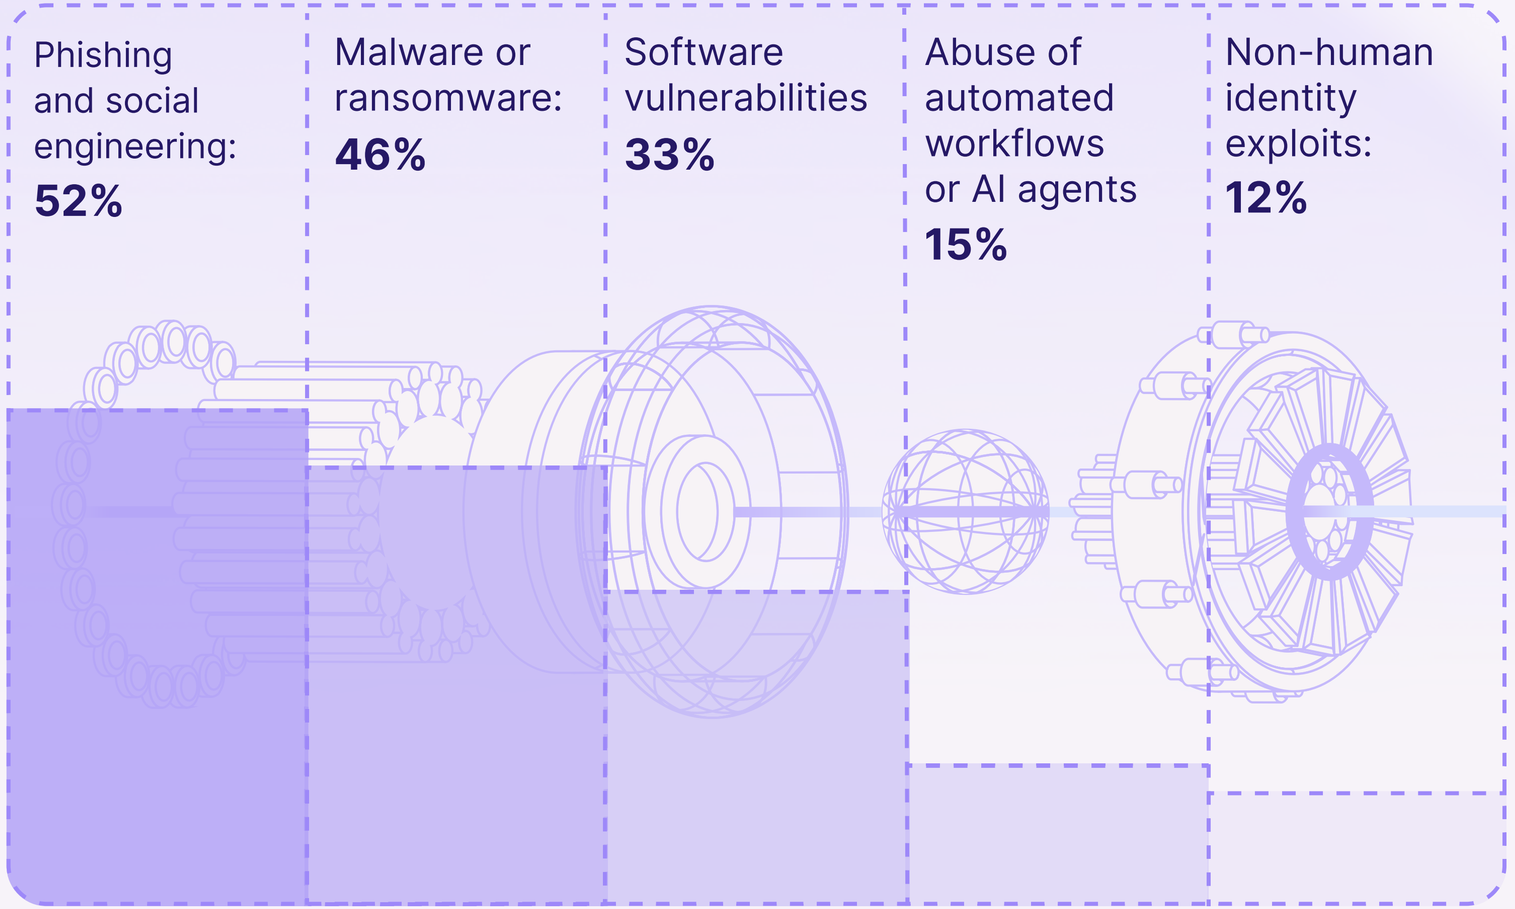

The Identity Security Snapshot

Complexity remains the biggest barrier

System complexity leads all IAM challenges (55%), followed by tooling limitations (41%). Complexity has now ranked as the top identity challenge for multiple years in a row, signaling that traditional tooling struggles to keep pace with expanding automation and identity sprawl. More than a quarter of respondents say existing tools were not designed for AI agents.

Every new SaaS application, automation, or AI workflow increases identity sprawl. Security teams are struggling because legacy identity platforms cannot scale to autonomous environments.

IAM CHALLENGES

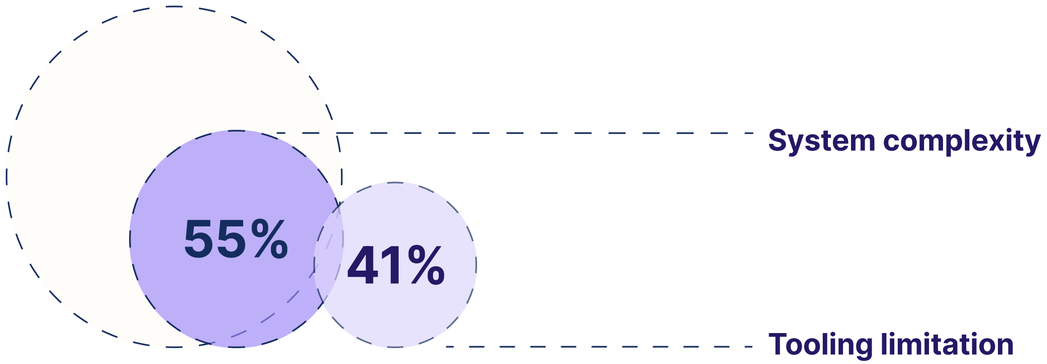

Security is the primary driver

Improving security (81%) and reducing risk (66%) dominate IAM priorities. Improving security has ranked as the top identity priority for two consecutive years, rising from 77% in 2025 as organizations shift identity from a productivity tool to a core security control. Identity has decisively changed from an IT productivity item to a foundational security discipline.

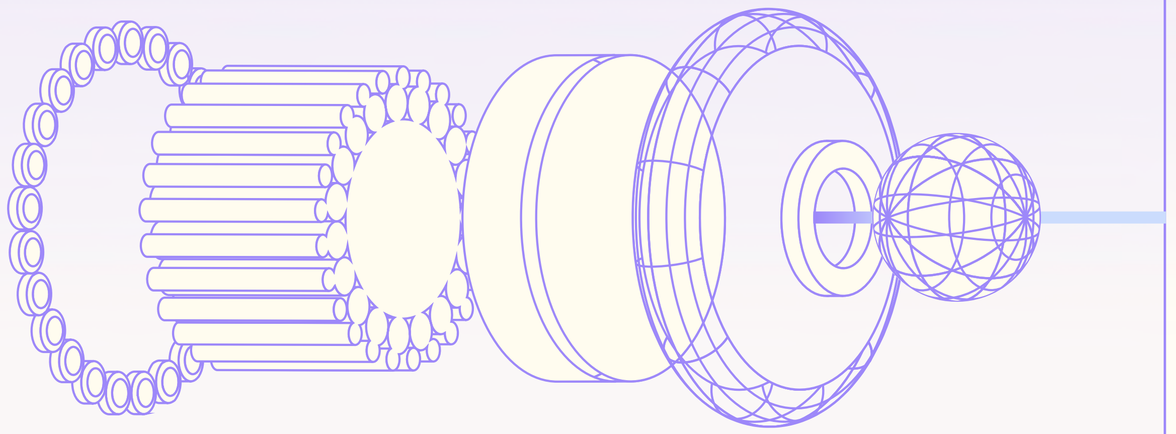

Non-Human Identities: The Governance Gap

NHIs are multiplying faster than governance

Nearly half of organizations report more non-human identities than human users. Service accounts, API keys, automation workflows, and AI agents are driving rapid growth across enterprise environments. As this expansion accelerates, identity shifts from a directory problem to a governance and orchestration challenge.

Non-human identities often operate with persistent credentials, elevated permissions, or limited oversight, making them attractive targets for attackers. Without clear ownership and visibility, compromised tokens or service accounts can be abused long before they are detected.

Visibility is declining year over year. Only 22% of organizations report full insight into their non-human identities, down from 30% the previous year. This suggests organizations are discovering identity sprawl faster than they can secure it.

The implication is twofold: security teams face increasing exposure to traditional attacks like credential abuse and privilege escalation, while also preparing for a future where automated systems hold real operational power. Without continuous visibility and policy enforcement, non-human identities expand the attack surface faster than governance models can adapt.

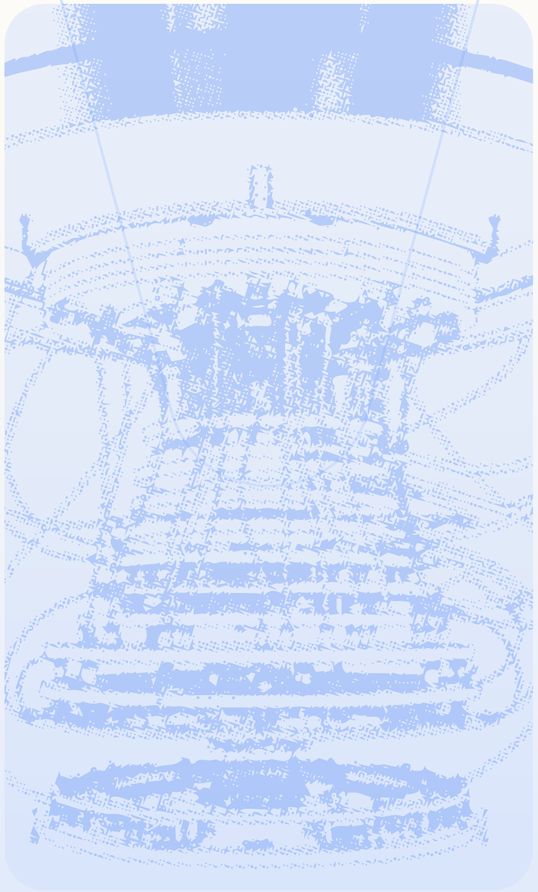

Non-Human Identities: The Governance Gap

Urgency is high, confidence is mixed

87% rate NHI risk as moderately to extremely urgent.
The top challenges include:

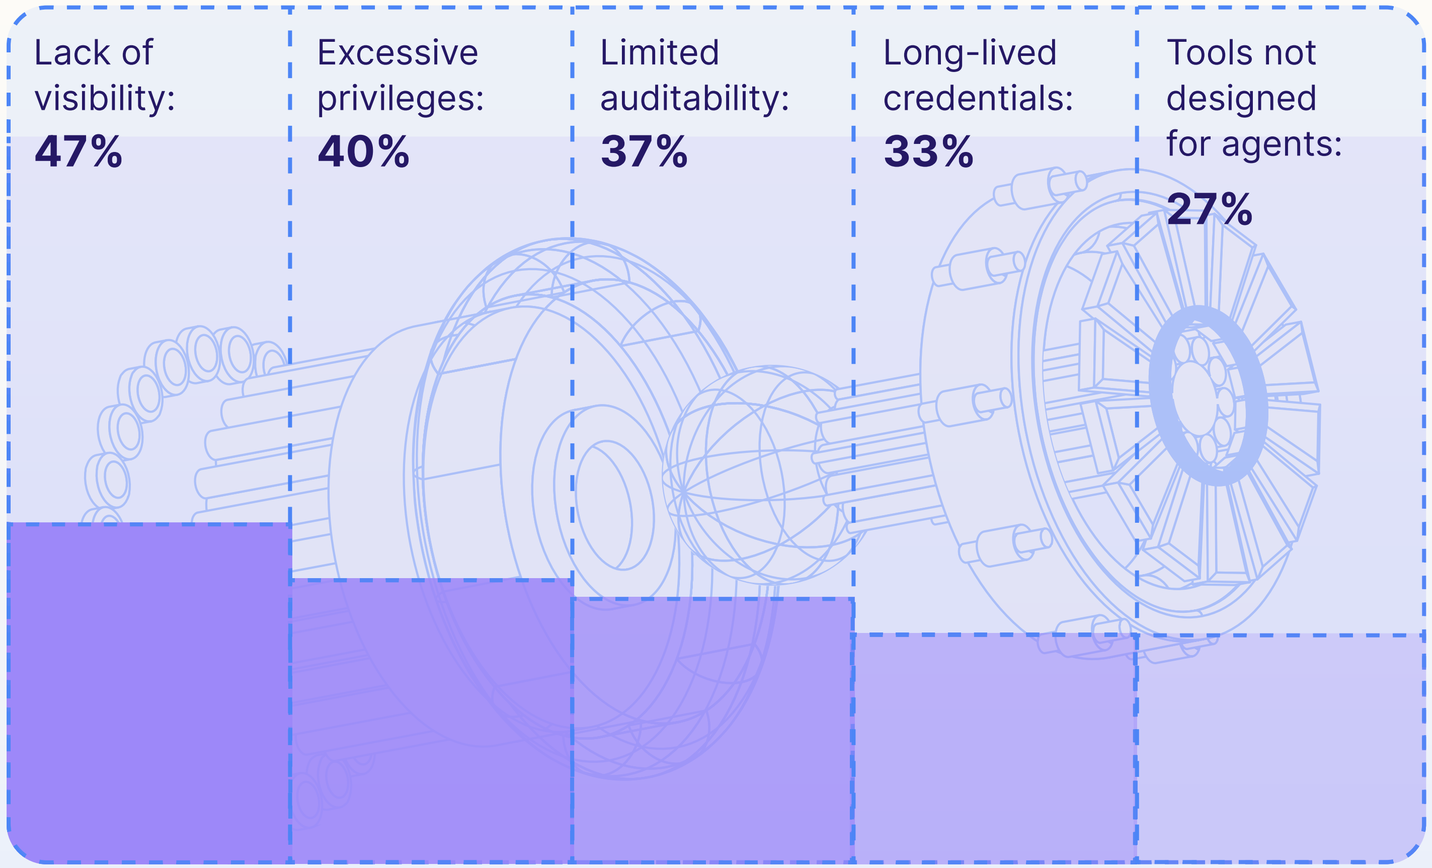

This gap between urgency and confidence defines the current identity market. Organizations know the risk exists. Many lack platforms built for NHI management.

IAM for NHIs moves from roadmap to reality

45% already use IAM tools for NHI governance, and another 45% plan to within 12 months. Human-only identity governance is ending. Unified governance across human and non-human identities is becoming a baseline expectation.

AI Agents at Work

Autonomous execution is already here

The fastest acceleration in enterprise autonomy is already underway.

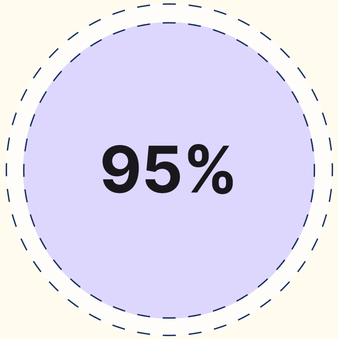

95% of organizations report AI agents performing tasks such as helpdesk workflows, security monitoring, access reviews, and privileged actions.

Just one year ago, 96% expected this future to arrive. The industry moved from prediction to execution in less than twelve months.

This surge introduces a structural challenge.
More autonomous actors are entering enterprise environments faster than governance models can evolve. Organizations now face an imbalance between machine-speed execution and human-designed security controls.

Autonomy is increasing

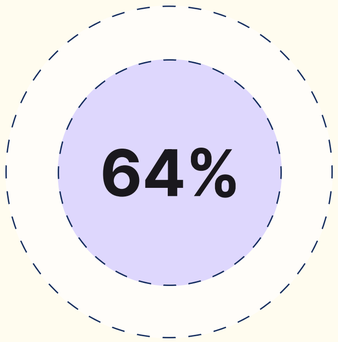

64% allow AI agents to act autonomously within policies or with post-action review.

Security teams cannot manually review machine-speed decisions. Governance must shift toward precise policy enforcement and continuous auditing.

The access model gap

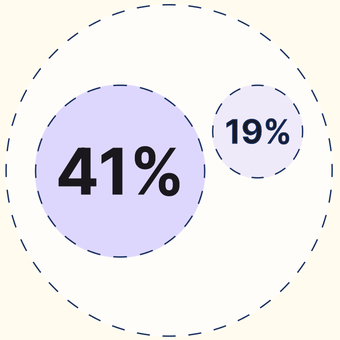

41% still rely on standing permissions, while only **19%** use continuous policy-based enforcement. Static access models were built for human workflows.

Without identity systems that adapt decisions in real time based on context, organizations cannot safely scale AI-driven productivity. This mismatch between static access models and dynamic AI behavior is one of the largest structural risks revealed by the survey.

Pressure Points

Security teams are shifting toward oversight and policy engineering

AI adoption is not simply automating tasks. It is changing how security teams operate.

37%
say AI replaces manual execution

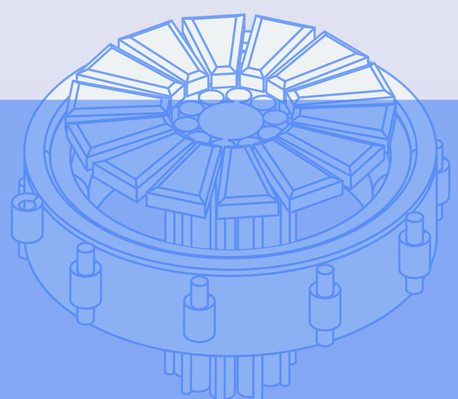

33%
say it augments existing roles

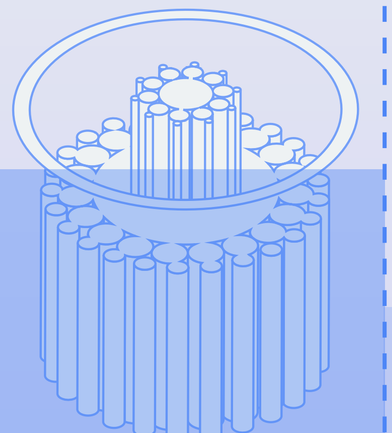

17%
report redesigning roles around policy definition and oversight

The data suggests a transition from operational execution toward governance engineering.
As agents take on provisioning, monitoring, and remediation tasks, human roles increasingly focus on defining policies, validating outcomes, and maintaining auditability.

The role of the practitioner is shifting from executing access changes to designing the systems that govern autonomous actors.

Conclusion

The enterprise has entered its first truly autonomous phase.

AI agents are executing work at scale, expanding productivity while simultaneously widening security gaps. Identity has become the system that determines whether autonomous environments remain controlled or drift into unmanaged risk.

The organizations that succeed will not be those that deploy the most AI. They will be the ones that build identity systems capable of governing autonomous actors continuously, dynamically, and at machine speed.

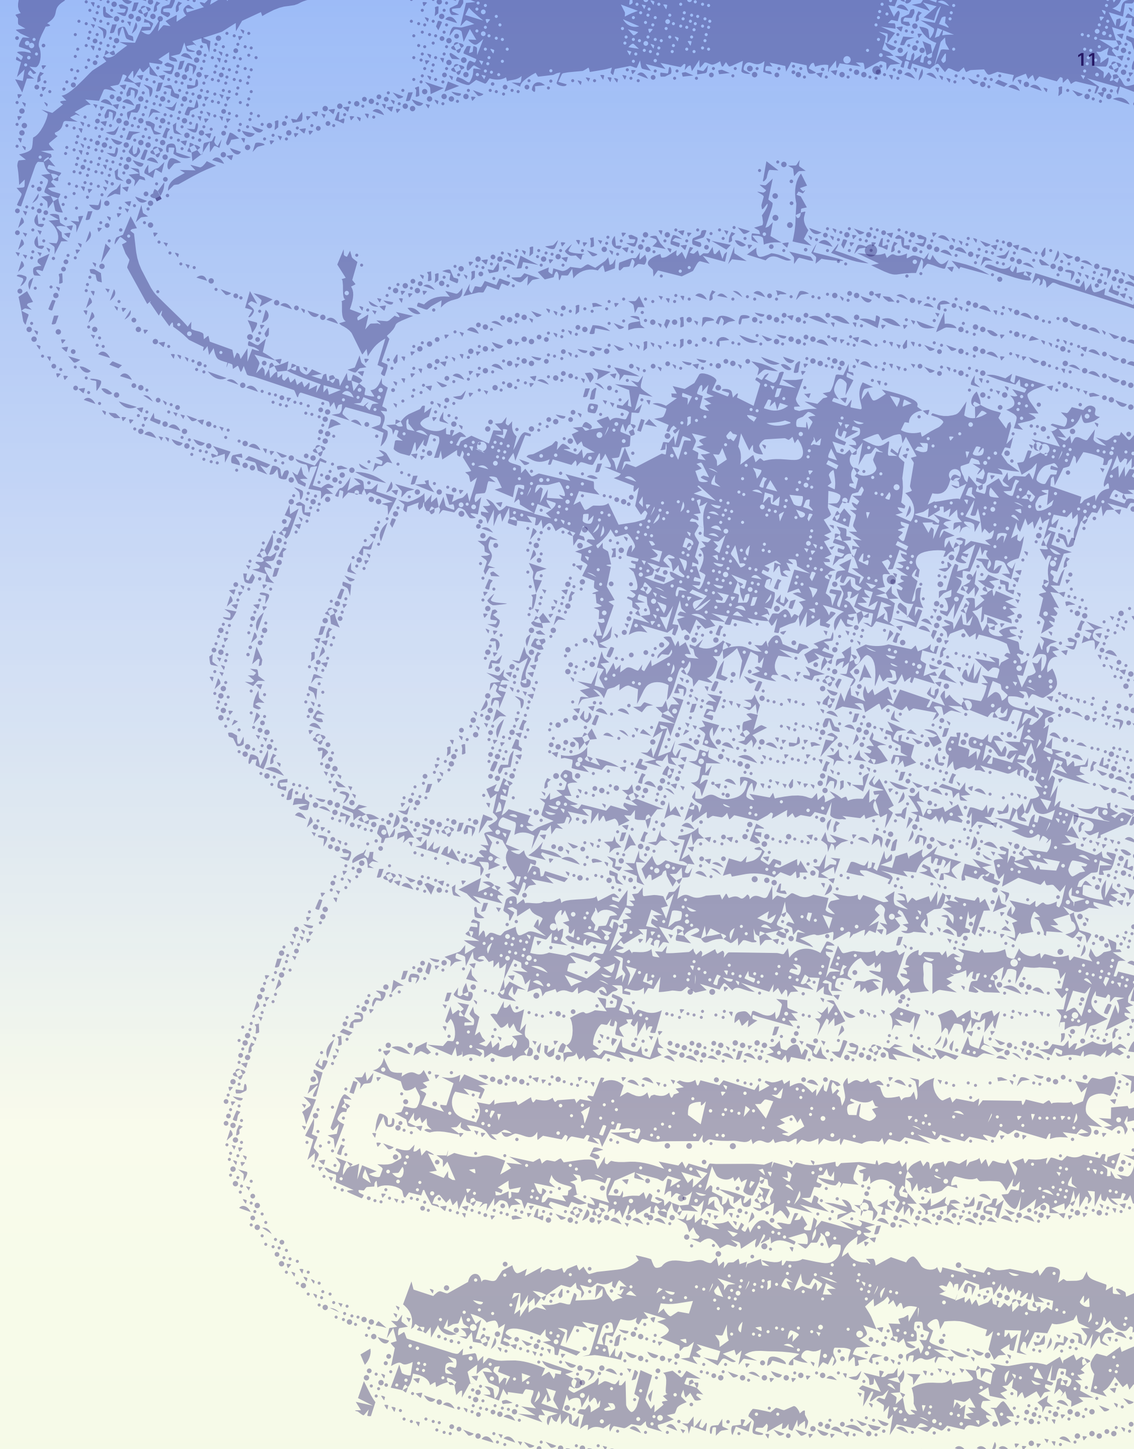

C1

2026 Future of Identity Report

Learn more at
C1.ai