Omada

The State of **Identity Governance 2026**

Security, Scale, and Visibility in the Age of Non-Human Identities

Contents

Executive Summary

This report is based on primary research conducted with 577 identity and security leaders at U.S. enterprises.

Identity security has entered a period of rapid change. Organizations are investing heavily in identity governance, express strong confidence in their identity security capabilities, and are actively adopting Zero Trust and AI-driven automation. Identity is widely recognized as a critical security control, and concern about identity-related threats is high across both traditional and emerging attack vectors.

At the same time, the survey data reveals a growing gap between intent and execution . Executive reporting most consistently tracks provisioning, deprovisioning, and access lifecycle activity, while leading indicators of accumulated identity risk and control effectiveness are tracked less uniformly. Identity data remains fragmented across platforms, limiting unified visibility and oversight. Non-human identities now outnumber human identities, yet ownership and accountability are distributed across many teams. Agentic AI is rapidly being deployed into environments where governance models are still evolving.

Organizations are making real progress in identity governance and security, but identity environments are scaling faster than the security infrastructure management tools designed to govern them. Automation and AI are increasing the number, speed, and autonomy of access decisions, while visibility, integration, and ownership have not advanced at the same pace.

The findings in this report point to an inflection point. Identity security is no longer a periodic, human-centered function. It is becoming a continuous, machine-driven control layer that underpins Zero Trust, automation, and AI-enabled workflows. Organizations that succeed in this new environment will be those that treat identity security as a critical control surface, with clear ownership, consistent integration, and executive-level visibility.

Finding 1: IGA Investments Are Fueled by Risk Reduction and Compliance Efficiency

Survey Question

What are the key business challenges driving your IGA investment as part of an identity-first strategy? (Select up to three)

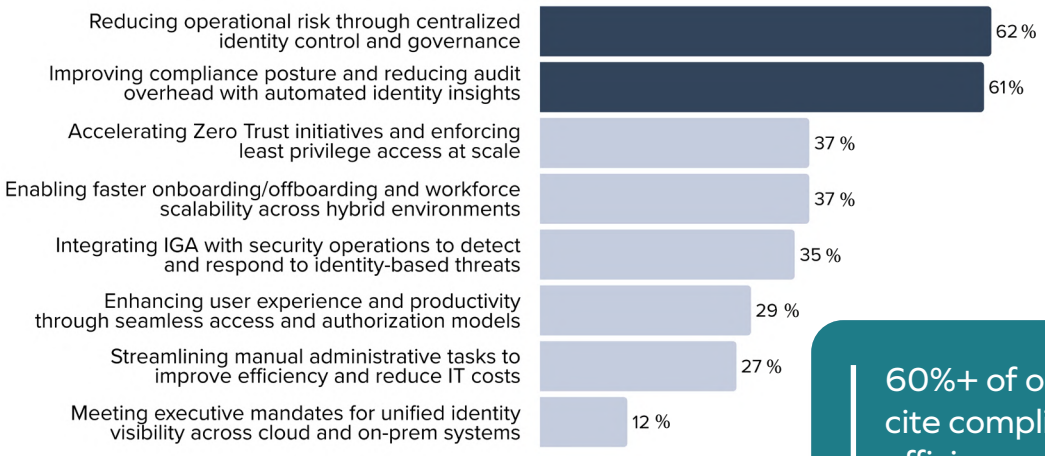

> 60%+ of organizations cite compliance and audit efficiency as primary drivers for IGA investment, while executive mandates for identity visibility rank far lower.

What the Data Shows

Organizations are investing in identity governance primarily to reduce risk and meet compliance demands. The most frequently cited drivers are improving compliance posture and reducing audit overhead, followed closely by reducing operational risk through centralized identity control. Accelerating Zero Trust initiatives is also a major driver, cited by roughly half of respondents.

By contrast, relatively few organizations identify executive mandates for unified identity visibility as a primary driver of IGA investment. This suggests that most identity programs are justified on operational and regulatory grounds rather than as an executive-level visibility or security initiative.

What This Tells Us

Organizations are investing in identity governance to address urgent, practical pressures. Compliance requirements and audit readiness remain major drivers, alongside the need to reduce operational risk through stronger, centralized control.

Zero Trust is also a prominent driver. Many organizations are using IGA to enforce least privilege at scale and to support more dynamic access decisions across cloud and hybrid environments.

At the same time, executive mandates for unified identity visibility rank much lower than risk and compliance drivers . This suggests many programs are still positioned as control and efficiency initiatives, not as a source of executive-level insight. That framing can shape which metrics are prioritized and what leaders see.

Why This Finding Matters

When investment is driven primarily by compliance and operational risk, reporting often prioritizes process performance and audit outputs. In fast-changing identity environments, leaders also need visibility into risk exposure and control effectiveness.

The next findings show how organizations express confidence in identity security, yet executive reporting and system integration do not consistently provide the risk visibility required for effective oversight.

Finding 2: Confidence Is High, but Based on Perception, Not Measurement

Survey Question

Indicate your agreement with the following statements:

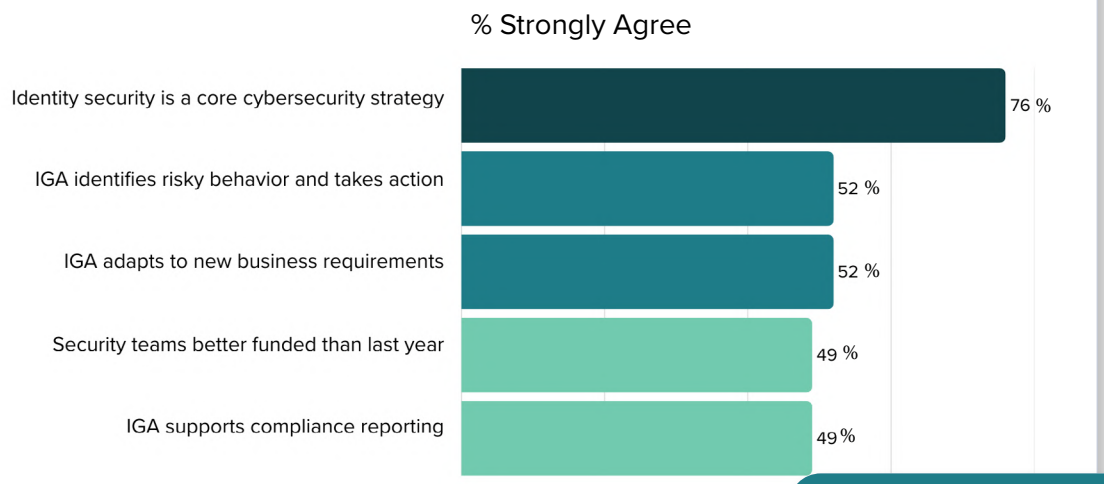

> 76% strongly agree identity security is core to cybersecurity strategy.

What the Data Shows

Respondents report high confidence across all statements. A strong majority agree that identity security is central to their cybersecurity strategy and that investment in security teams and resources has increased. Most also agree that their IGA solutions are adaptable, support compliance requirements, and can identify risky behavior.

Levels of disagreement are consistently low. Overall, the data reflects a broad sense that identity governance programs are capable and improving.

What This Tells Us

Organizations believe they are making progress in identity security, and that belief reflects real investment and focus. Identity has earned a strategic role, and many programs have matured.

However, confidence reflects perception rather than measured outcomes. Agreement with statements show belief in capability, not proof of consistent execution. They do not show how clearly identity risks reach leadership, how consistently controls are applied, or how quickly issues are detected and resolved.

As identity environments become more dynamic and automated, this difference matters. Feeling confident is not the same as having clear, continuous visibility into identity risk.

Why This Finding Matters

High confidence is an asset when it is paired with visibility. Without clear measurement and reporting, confidence can mask gaps that expand as scale increases.

The findings that follow show where executive reporting and system integration do not yet match the speed and complexity of modern identity environments.

Finding 3: Executives See Identity Activity Before Identity Risk

Survey Question

Which executive reporting metrics related to identity governance and access does your organization regularly track and present to senior leadership? (Select all that apply)

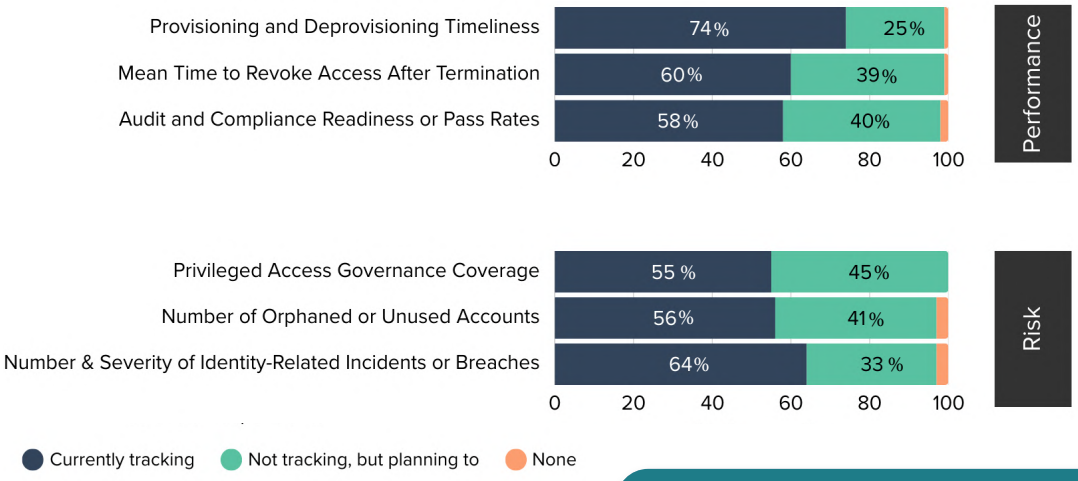

While 74% of organizations regularly report provisioning and deprovisioning timeliness to executives, leading indicators such as privileged access coverage and orphaned accounts are tracked less consistently.

Executive Callout: What Leaders Often See vs. What They Need to See

What leaders often see

- Provisioning and deprovisioning speed
- Number and severity of identity-related incidents or breaches
- Audit and compliance readiness or pass rates

What leaders need to see

- Accumulating privileged access and unused accounts
- Time to revoke access after role change or departure
- Where non-human and automated access is expanding without clear ownership

Closing this gap is essential as identity decisions become continuous, automated, and increasingly machine-driven.

What the Data Shows

Executive reporting is most mature around provisioning deprovisioning and access lifecycle activity, with nearly three-quarters of organizations tracking provisioning and deprovisioning timeliness. A majority of organizations also report tracking lagging risk indicators, such as the number and severity of identity-related incidents or breaches, as well as audit and compliance readiness.

By contrast, leading indicators of identity risk are tracked less consistently. Metrics such as privileged access governance coverage, the number of orphaned or unused accounts, and mean time to revoke access after termination are reported by a narrower share of organizations, with many indicating these metrics are planned rather than currently tracked.

What This Tells Us

Leaders are more likely to see how quickly access is provisioned and what incidents have already occurred than to see where identity risk is actively building.

Operational metrics and incident counts are familiar, easier to standardize, and align well with traditional service management and security reporting. Leading indicators of identity risk, however, require broader visibility across identity sources, tighter integration between IAM, HR, security, and audit systems, and analytics that can normalize identity data across environments.

When those foundations are incomplete, executive dashboards skew toward activity and outcomes rather than exposure and accumulation of risk. As a result, organizations may respond effectively to incidents while lacking early warning signals that could prevent them.

Why This Finding Matters

As identity environments grow more dynamic, automated, and non-human, this imbalance limits effective oversight. Lagging indicators help organizations understand what has already gone wrong; leading indicators help them understand what is likely to go wrong next.

Without consistent visibility into privileged access, orphaned accounts, and access revocation delays, organizations risk managing identity security reactively rather than proactively. This gap becomes more consequential as automation and AI increase the speed and scale of access decisions across the enterprise.

Finding 4: Zero Trust Is Widely Adopted, but Interoperability Gaps Limit Visibility

Survey Questions

To what extent is your IGA program integrated with Zero Trust security principles beyond basic least privilege?

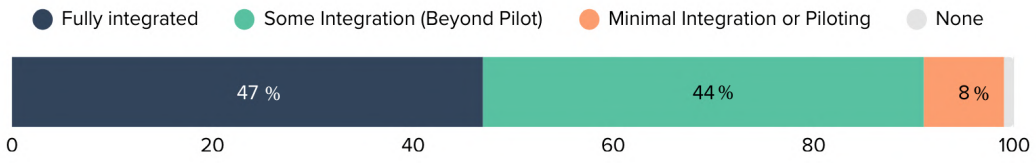

Survey Questions

Which factors present the biggest obstacles to achieving seamless interoperability across security solutions from different providers?

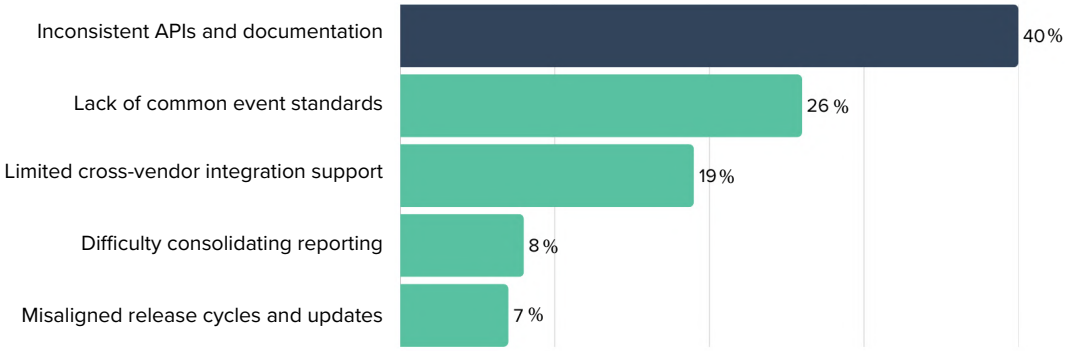

> 95% of organizations report some level of Zero Trust adoption, yet interoperability gaps remain the top barrier to unified identity visibility and reporting.

What the Data Shows

Most organizations report that their IGA programs are aligned with Zero Trust principles. Nearly all respondents indicate some level of integration, and close to half describe their programs as fully integrated.

At the same time, organizations report significant technical barriers to interoperability across security platforms. The most frequently cited obstacles include inconsistent API support, lack of common standards for event signaling, and difficulty consolidating identity data for unified reporting and analytics.

What This Tells Us

Organizations have broadly adopted Zero Trust as a guiding security model, and identity governance sits at the center of that approach. However, Zero Trust only works when identity information can move reliably between systems.

When identity data is spread across tools with different APIs, inconsistent standards, and limited integration, policies cannot be applied or reported consistently. Individual systems may enforce Zero Trust decisions correctly, but those decisions remain isolated.

This explains why organizations can pursue Zero Trust while still lacking clear, unified visibility . Identity decisions are being made across many platforms, but the data needed to understand and explain those decisions does not come together in one place.

Why This Finding Matters

Zero Trust increases the importance of identity governance rather than reducing it. When access decisions are continuous and context-aware, gaps in integration become gaps in control. Without interoperability across identity, security, and governance systems, Zero Trust initiatives risk becoming a collection of local controls rather than a coherent operating model.

This pattern is consistent with broader industry research. The Identity Defined Security Alliance (IDSA) reports near-universal adoption of Zero Trust principles, while also noting that most organizations implement Zero Trust selectively rather than as a fully integrated operating model. This reinforces the challenge of achieving unified visibility across fragmented identity and security systems.

Finding 5: Concern About Identity Threats Is High Across All Attack Types

Survey Question

How concerned is your organization about the following identity-related security threats?

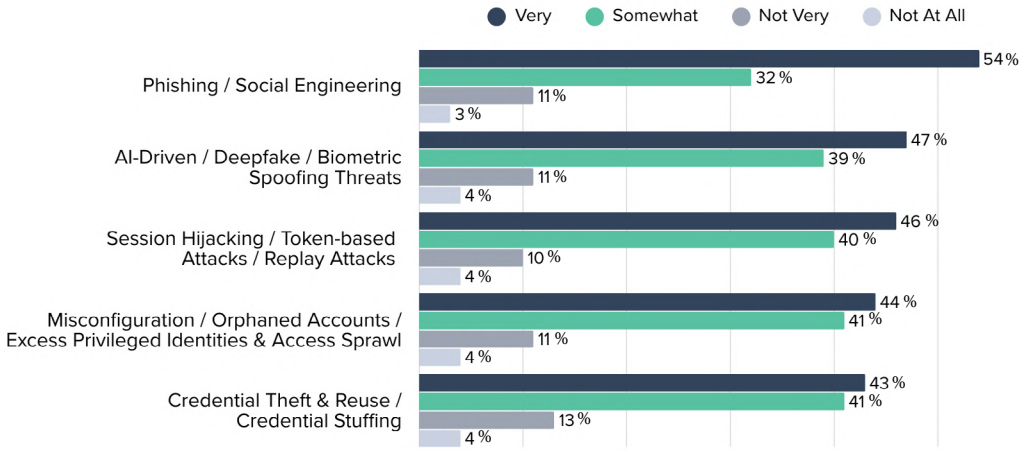

> More than 80% of respondents report concern about each major identity threat category, with AI-driven risks viewed at levels comparable to traditional attack vectors.

Threats assessed include phishing and social engineering, credential theft and reuse, session hijacking and token-based attacks, misconfiguration and orphaned accounts, excess privileged access, and AI-driven or deepfake-enabled threats.

What the Data Shows

Concern levels are high across all identity-related threat categories. A strong majority of respondents report being either very concerned or somewhat concerned about phishing and social engineering, credential theft and reuse, session hijacking, and misconfiguration or orphaned accounts.

Notably, concern about AI-driven and deepfake-enabled threats is comparable to concern about more established identity attack vectors. No single threat category stands out as marginal or low priority.

What This Tells Us

Organizations understand that identity is a primary attack surface. The breadth of concern across both traditional and emerging threats reflects a clear awareness that identity-related risk is not limited to one vector or technology.

This awareness exists alongside the visibility and integration gaps highlighted in earlier findings. Organizations recognize the risks, but many lack a unified view of where identity exposure is increasing or which controls are most effective at reducing it.

As identity environments grow more complex, concern alone is not enough. Without consistent visibility across users, non-human identities, privileges, and automated access decisions, organizations struggle to translate awareness into targeted action.

Why This Finding Matters

High concern signals urgency, but it does not guarantee readiness. When risk awareness outpaces governance visibility, organizations tend to respond reactively rather than proactively.

This disconnect becomes more pronounced as AI-driven automation and non-human identities increase the speed and scale at which identity-related decisions need to be made. The findings that follow show how agentic AI adoption is accelerating identity operations, often faster than governance practices can adapt.

Finding 6: Organizations Use GenAI to Scale Identity Security

Survey Questions

What are your organization's primary business goals for using GenAI in security or identity governance?

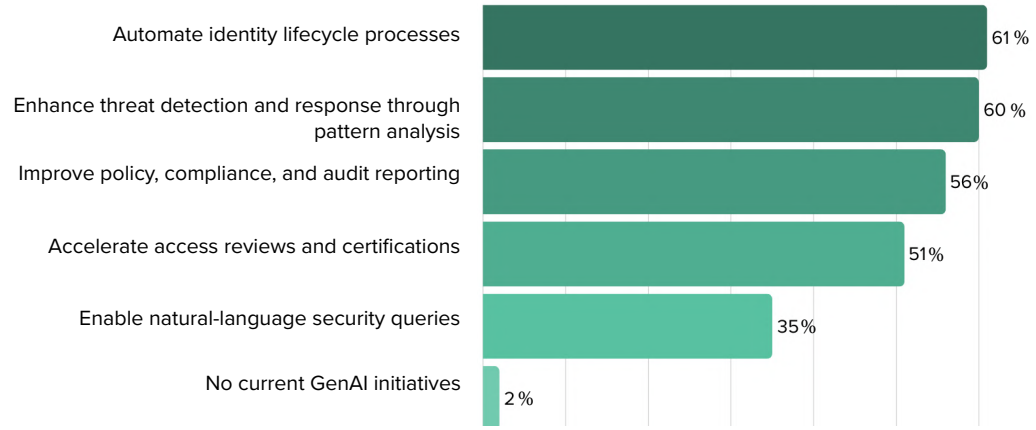

Survey Questions

Which GenAI capabilities do you consider essential for modern identity governance?

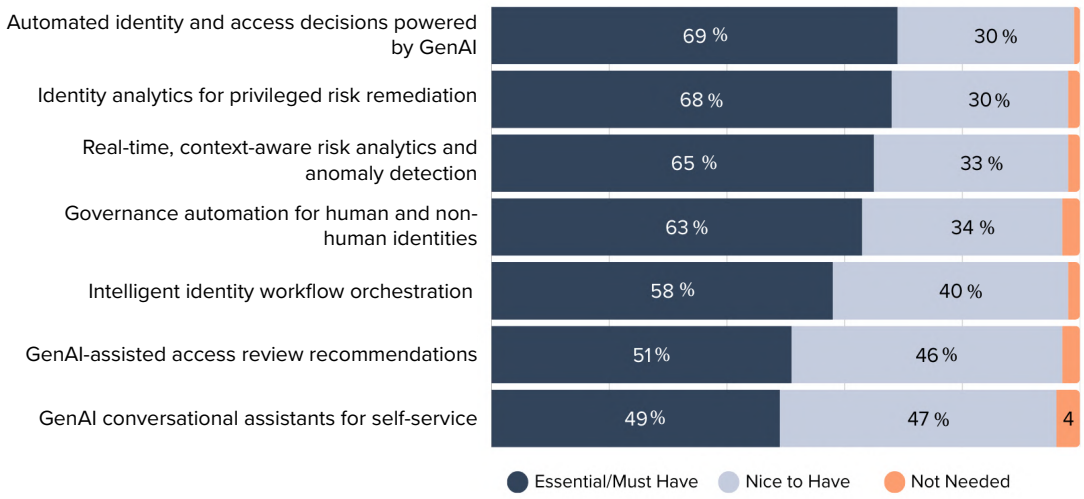

> Over 60% of organizations cite automating identity lifecycle processes and scaling identity operations as their primary GenAI use cases.

What the Data Shows

Organizations are adopting GenAI in identity programs to handle scale, speed, and complexity. The most frequently cited goals focus on automating identity lifecycle processes, improving operational efficiency, and enhancing the ability to analyze identity data.

Respondents also identify AI-driven analytics and automation as top priorities when evaluating identity governance solutions. Capabilities such as AI-powered insights, automated decision support, and adaptive governance rank higher than traditional feature enhancements.

Taken together, the data shows that GenAI is viewed primarily as a way to manage the growth in identities, access changes, and policy decisions that identity teams are responsible for. Organizations are not adopting AI for experimentation alone; they are adopting it to keep identity operations manageable as environments become more dynamic and automated.

What This Tells Us

Identity teams are under pressure to manage a rapidly growing number of identities, access changes, and automated decisions with the same or fewer resources. Cloud adoption, remote work, non-human identities, and automation have increased the number of access decisions dramatically. GenAI is being positioned as a response to that pressure.

This is a practical choice. Manual processes and periodic reviews cannot keep pace with modern identity environments. Organizations are turning to GenAI to automate decisions, surface risk signals, and reduce operational load.

At the same time, the data shows that GenAI adoption is focused on scaling existing processes, not rethinking governance models. AI is being used to accelerate what organizations already do, rather than to redefine how identity risk is measured, owned, and governed.

Why This Finding Matters

GenAI is becoming a core component of identity operations. As it takes on a greater role in provisioning, access decisions, and analysis, it also becomes part of the control surface.

The findings that follow show that governance practices have not yet caught up with this shift. Organizations are moving quickly to apply AI to identity, but many have not yet adapted ownership models, credential practices, and oversight mechanisms to match the new scale and autonomy AI introduces.

Finding 7: Agentic AI Governance Expectations Outpace Current Controls

Survey Question

How does your organization handle credential management for agentic AI agents?

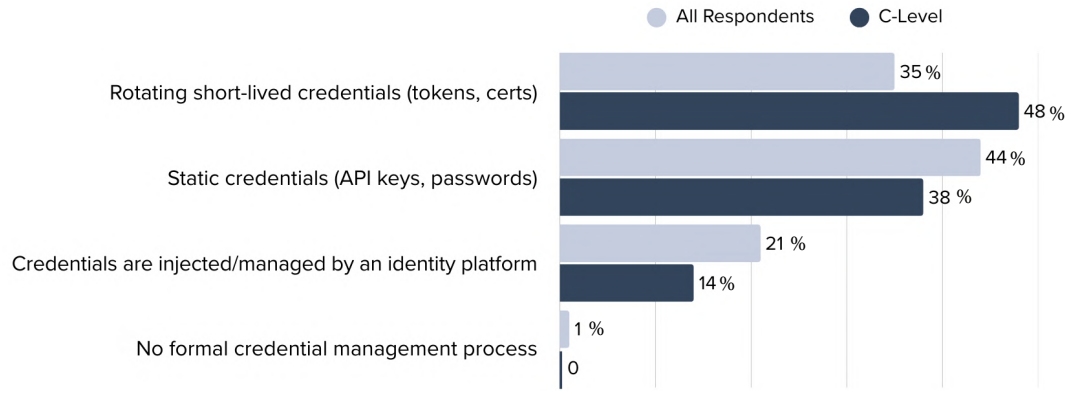

Survey Question

Does your organization assign a unique identity to each agentic AI agent?

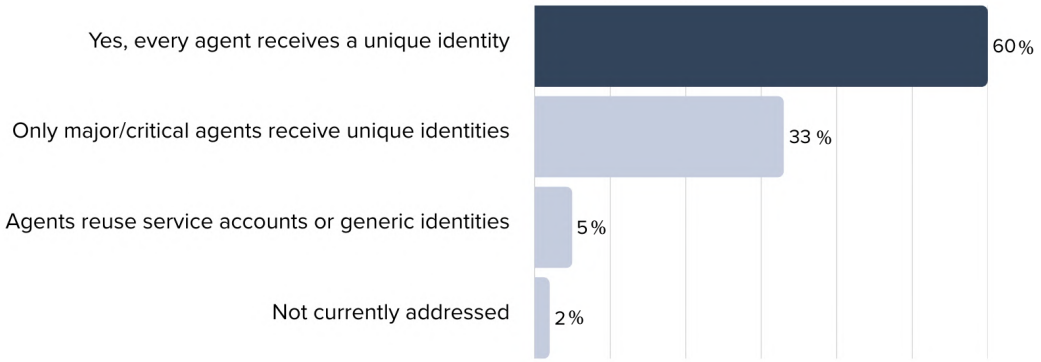

> While 48% of executives report the use of rotating credentials for AI agents, overall responses indicate that these practices are not yet applied consistently across organizations.

What the Data Shows

Credential management practices for agentic AI vary widely across organizations. Many respondents report using stronger practices such as rotating, short-lived credentials and assigning unique identities to AI agents. At the same time, a significant portion rely on static credentials or shared accounts, indicating inconsistent application of governance controls.

Differences also appear between executive and non-executive responses. C-level respondents are more likely to report the use of stronger practices, such as rotating credentials and unique identities, than the overall respondent population.

What This Tells Us

Organizations recognize that AI agents require stronger controls than traditional automation. Expectations for governance are high, especially at the leadership level.

Differences between executive and non-executive responses suggest that governance practices may not be consistently understood or visible across teams, especially in environments where AI agents are managed by multiple functions. In some cases, this may reflect uneven implementation; in others, it may indicate gaps in visibility or communication about how controls are deployed.

As AI agents operate continuously and at scale, reliance on static or shared credentials can increase exposure over time, particularly when ownership and oversight are unclear.

Why This Finding Matters

As AI agents take on more responsibility, they must be governed as identities, not tools. Unique identities, strong credential hygiene, and clear ownership are foundational controls, not optional enhancements.

When governance expectations outpace actual practices, organizations increase their exposure without realizing it. This mismatch becomes more difficult to correct as AI-driven workflows expand, creating conditions for the challenges explored in the next findings.

Finding 8: Non-Human Identities Dominate, but Ownership Is Fragmented

Survey Question

What is the approximate ratio of non-human identities to human identities in your organization?

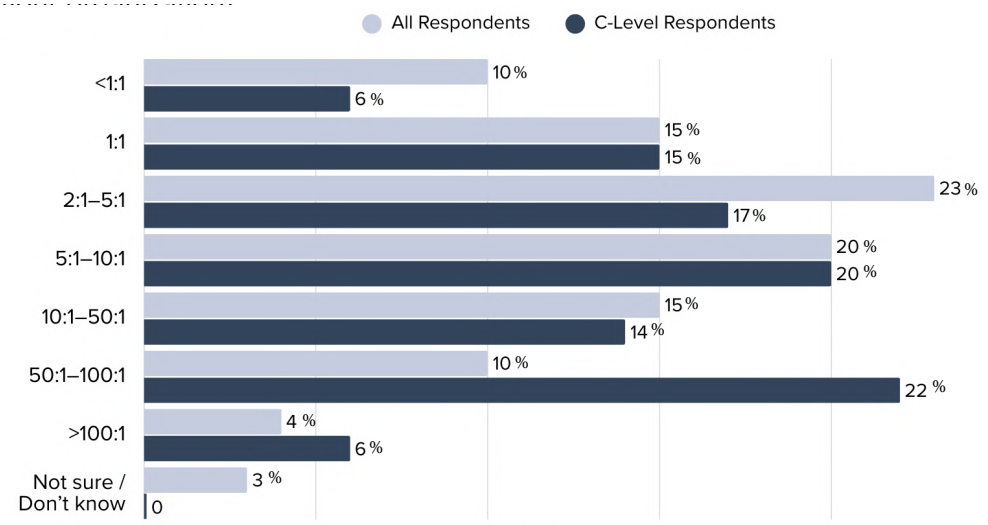

Survey Question

In your organization, who is primarily responsible for identifying and managing the security of non-human identities?

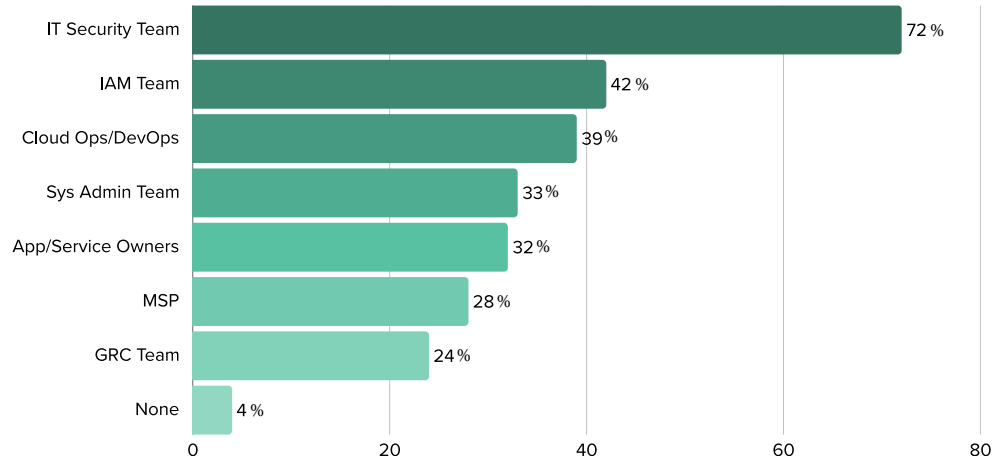

Values represent % of ownership. Respondents could select multiple teams if ownership was shared.

> Executive respondents are far more likely to report NHI-to-human ratios of 50:1 or higher, suggesting the scale of non-human identities is often underestimated across organizations.

What the Data Shows

Non-human identities now outnumber human identities in most organizations. Executive respondents are more likely to report very high ratios of non-human to human identities, often far exceeding estimates provided by non-executive respondents.

This aligns with broader industry research. IDSA reports that non-human identities already outnumber human identities by a significant margin, underscoring how quickly machine and service identities have become the dominant identity type in modern environments.

Responsibility for non-human identities is widely distributed. Security teams are most frequently cited, followed by IAM teams, DevOps, cloud operations, application teams, and infrastructure teams. A small percentage of respondents report having no clearly defined owner at all.

Overall, the data shows both scale and fragmentation. Non-human identities are numerous, and accountability for them is shared across many functions.

What This Tells Us

The growth of non-human identities has outpaced traditional ownership models. Service accounts, API keys, automation, and AI agents are created across platforms and teams, often outside the workflows originally designed for human identity governance.

Differences between executive and non-executive responses suggest that the scale of non-human identities may not be consistently understood across organizations. Executives are more likely to report very high NHI-to-human ratios, while practitioners often report lower estimates based on the identities they directly manage.

This discrepancy suggests that many organizations may be underestimating the true number of non-human identities in their environments. If so, the magnitude of the AI- and automation-related governance challenge is likely larger than many teams currently assume.

When ownership is spread across teams, governance becomes harder as scale increases. Without clear accountability, organizations struggle to answer basic questions about purpose, privilege, and lifecycle. These gaps can create persistent exposure, especially when non-human identities operate continuously and with elevated access.

Why This Finding Matters

Non-human identities are no longer an edge case. They represent a large and growing share of the identity population and play a central role in automation and AI-driven workflows.

When ownership is unclear and scale is underestimated, controls weaken as environments expand. This challenge underpins many of the governance gaps described earlier, including those related to AI agents, credential hygiene, and executive visibility.

Agentic AI represents the most advanced expression of this trend, combining non-human identity scale with autonomy and continuous decision-making.

Finding 9: Agentic AI Is Moving Faster Than Governance Models

The previous findings highlighted gaps in visibility, ownership, and credential practices. This section examines what happens when those gaps intersect with the scale and autonomy of agentic AI.

Survey Questions

Is your organization currently using, piloting, or planning to deploy agentic AI?

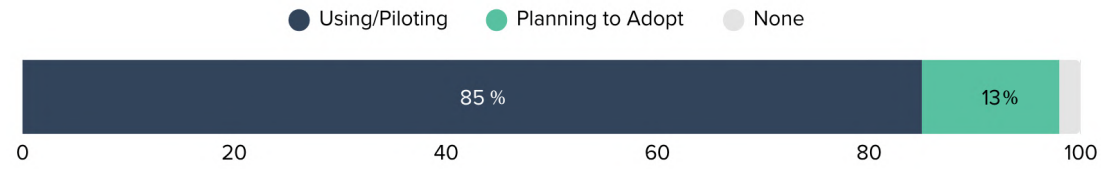

Survey Questions

Is your organization currently using, piloting, or planning to deploy agentic AI?

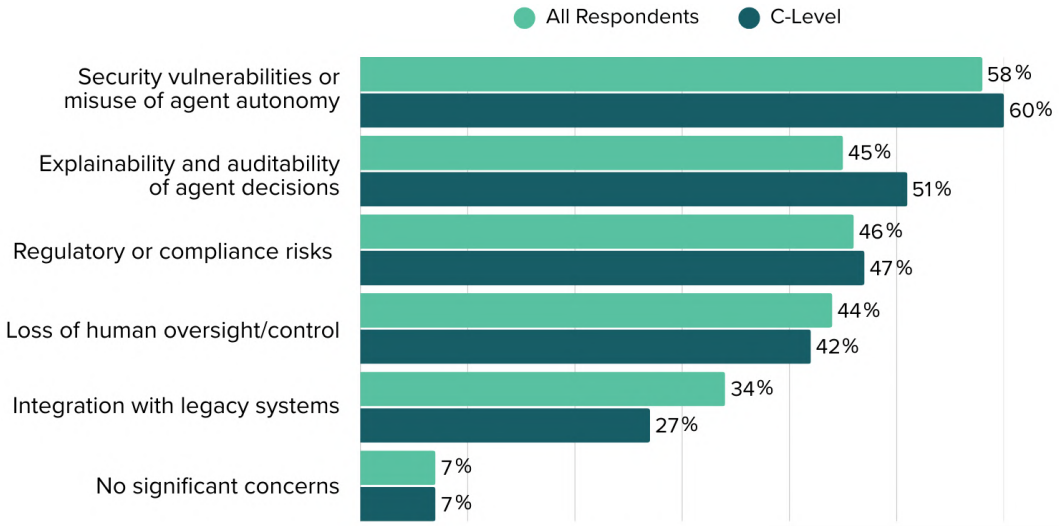

> 85% of organizations are already using or piloting agentic AI, while security is the most frequently cited concern associated with its deployment.

What the Data Shows

Agentic AI adoption is widespread. A large majority of organizations report that they are already deploying agentic AI, piloting it, or planning near-term implementation.

Security concerns rank highest among the risks associated with agentic AI. Respondents consistently identify security vulnerabilities as a top concern, ahead of regulatory, operational, or ethical considerations.

At the same time, earlier findings show uneven governance practices for AI agents, including variation in credential management, identity assignment, and ownership.

What This Tells Us

Organizations have made a clear decision to move forward with agentic AI. Adoption is no longer theoretical. Autonomous agents are being introduced into production environments where they can act continuously and at scale.

Governance, however, is still catching up. While leaders recognize the risks and expect strong controls, many organizations are still adapting identity models, ownership structures, and oversight mechanisms to account for autonomous behavior.

This creates a familiar pattern. Technology adoption accelerates first, driven by efficiency and competitive pressure. Governance follows more slowly, constrained by legacy processes and fragmented accountability.

Why This Finding Matters

Agentic AI changes the risk profile of identity security. Autonomous agents can create, modify, and act on access without human intervention. When governance maturity lags adoption, small gaps in control can persist and expand unnoticed.

Other research reflects similar concerns. The Identity Defined Security Alliance reports that nearly one-third of organizations have already experienced an AI-generated identity incident, highlighting that AI-driven identity risk is no longer theoretica l.

This finding reinforces the central message of the report. Identity governance must evolve to match the speed, scale, and autonomy of modern environments. Without that evolution, organizations risk automating access faster than they can govern it.

What the Data Makes Clear

Across the findings, a consistent picture emerges. Organizations are investing heavily in identity governance, express confidence in their capabilities, and are actively adopting Zero Trust and AI-driven automation. Identity security is widely recognized as critical, and concern about identity-related threats is high.

At the same time, visibility and governance have not kept pace with scale. Executive reporting continues to emphasize operational activity over risk exposure. Identity data remains fragmented across platforms, limiting unified reporting and oversight. Non-human identities now outnumber humans, yet ownership is dispersed across teams. Agentic AI is being introduced into environments where governance models are still evolving.

These conditions coexist. Strong intent, growing investment, and real progress are happening alongside structural gaps in visibility, integration, and accountability. This is not a failure of awareness or effort. It reflects how quickly identity environments have changed, and how much faster automation and AI are increasing the number and speed of access decisions.

The data points to a clear inflection point. Identity governance is no longer a periodic, human-centered control function. It is becoming a continuous, machine-driven operating layer that underpins Zero Trust, automation, and AI-enabled workflows.

Organizations that treat identity governance as a strategic control surface, with clear ownership, consistent integration, and executive-level visibility, will be better positioned to manage this shift. Those that rely on fragmented tooling, manual processes, or incomplete reporting will find it increasingly difficult to explain, control, and trust the access decisions being made on their behalf.

The next phase will be defined not by whether identity is important, but by whether it is governed at the same speed and scale at which it now operates. For identity and security leaders, this data offers both a benchmark and a roadmap: visibility into where peers are today, and clarity on where governance must go next.

About the Research

This report is based on primary research commissioned by Omada and conducted by independent research agency BCKL Group as part of the State of IGA for 2026 study.

The research is based on a survey of 577 identity, access management, and cybersecurity professionals who play a primary role in their organization's identity security management. Respondents represent U.S.-based enterprises with 1,000 to 10,000 employees, including organizations with 1,000–5,000 employees (62%) and 5,000–10,000 employees (38%).

Over half of respondents hold leadership roles, including C-level executives (19%) and senior management (37%), alongside IAM and IGA practitioners (36%), HR (6%), and Help Desk roles (2%). This mix reflects both executive and operational perspectives on identity governance.

The survey explored current practices, priorities, and perceptions related to identity governance, Zero Trust, non-human identities, and the adoption of GenAI and agentic AI. Results are self-reported and are intended to highlight patterns, gaps, and emerging trends rather than assess individual organizational maturity.

Omada

Omada reduces identity risk by securing access for every identity, human and non-human. Our AI-powered platform uses intelligent automation to enforce just-in-time, least-privilege access across hybrid and cloud environments. With unified visibility and continuous risk evaluation, security teams can detect threats faster, eliminate standing privilege, and scale identity security with confidence. Learn more at **omadaidentity.com**

V011326