September 2026
White Paper
Sponsored by GuidePoint Security

Managing Agentic AI Through the Identity Control Plane: What Organizations Should Look For

Grace Trinidad
Research Director, AI Security and Trust, IDC

Table of contents

Click any title to navigate directly to that page.

Executive summary

Identity has become an important security perimeter and a critical component of an organization's control plane. As workloads moved to the cloud and the network perimeter dissolved, the question that emerges as most important is no longer what sits inside the firewall. Instead, it is who, or what, is permitted to act and on what they are permitted to act.

AI agents are now being recognized as unique identities that require close, separate management in organizational systems. They are much like human users with credentials and permissions, yet unlike humans in their ability to work at machine speed and the risks that accompany that pace.

Identity teams are having to manage enormous changes to how identities are being created and managed. Every user, device, application, service, API key, bot, and AI agent requires verification, authentication, and authorization regardless of where or how it originates, including whether that's inside or outside the corporate network.

Securing agentic AI runs through three control planes: identity, which governs who or what may act; data, which governs what an agent may know and touch; and runtime, which governs what an agent may do in the moment. All three matter, but they are not peers. Data controls and runtime controls each depend on knowing, with confidence, which agent is acting. Identity is the control plane that allows the other two to work.

Governing AI agents is not the same as governing humans, and that difference is producing a consistent set of pain points across organizations.

- **You cannot see which agents you have or what they can access.** Most organizations cannot state an exact agent count, and the hardest agents to find are the ones running under a user's own credentials.
- **Traditional identity and privilege models do not fit.** An agent's risk is defined less by what it is than by what it can do, which pulls identity, privilege, monitoring, and incident response into a single problem.
- **The threat model has shifted to synthetic trust.** With multi-agent activity and AI-generated impersonation, human signals such as a voice on the phone or a face on a video call can be fabricated. Identity proofing that assumes a real person on the other end no longer holds.
- **Governance lags adoption.** Agents have entered an organization's environment under business pressures faster than security can write policy.
- **No single vendor covers the full framework.** Organizations are left assembling point tools while ownership of AI identity governance is contested internally.

Fortunately, among the security and identity leaders who participated in the research detailed here, there was strong consensus on the answer: every agent needs a named human sponsor and a managed lifecycle. Where organizations differ most is in how far they have matured in their approach to agent identity, not in the destination they are maturing toward: a purpose-built, distinctly governed identity class for every agent. The variation across organizations is a maturity gap, not a standing disagreement.

What organizations, enterprises, and government agencies need from their identity vendors is therefore practical:

- Visibility with automated discovery and an authoritative agent inventory
- An agent-first identity and privilege model
- An integrated governance framework
- Deepfake-resistant, adaptive verification
- Provable, no-lock-in solutions that meet them where they are on the maturity curve

None of these needs are speculative. They are already codified expectations, and each maps to a concrete capability that an identity program is supposed to deliver:

- **The OWASP Non-Human Identities Top 10 names the risks of unmanaged machine identities.** In plain terms, it expects a program to know every non-human identity it owns, retire each one cleanly when its work ends, keep privileges trimmed to the task, and rotate credentials before they age into liabilities. Its number one risk, improper offboarding, cannot even be attempted for an identity that never entered an inventory.
- **The OWASP Top 10 for Agentic Applications extends those expectations to agents.** Identity and Privilege Abuse targets the most common shortcut in agent deployments today: letting an agent borrow the identity of whatever is nearest, a user's live session, a cached administrative token, a shared service account. An agent that inherits a token from one workflow can reuse it in another, so a support agent finishing an IT task can carry the leftover administrative credential into a request that touches restricted HR data. Nothing was hacked; the agent simply held onto a privilege no one intended it to keep, and because it acted as someone or something else, the audit trail attributes the action to the wrong identity.

The expectation that follows is threefold:

1. Every agent acts under its own scoped identity with least privilege at the task level
2. Credentials are short-lived rather than inherited or cached
3. Runtime traceability records every action and decision under the agent's own name

This is also where the OWASP list and the panel converge from opposite directions, because agents running under a user's identity were exactly the population respondents said they could not see or count, and the inherited-credential pattern is the stage the study's least mature organizations are living in.

NIST and CSA name the controls. The NIST AI Risk Management Framework expects an organization to maintain an inventory of its AI systems as a basic governance practice, and its Govern, Map, Measure, and Manage functions extend to agents. The identity, logging, and data-protection control domains of the CSA AI Controls Matrix apply similarly. NIST's 2026 SP 800-53 control overlays and the NCCoE concept paper on agent identity make the extension explicit: agents are treated as distinct non-human identities (NHIs) with formal identification, managed lifecycles, and audited authorization.

OWASP Non-Human Identities Top 10
OWASP Top 10 for Agentic Applications
NIST AI Risk Management Framework
CSA AI Controls Matrix

- **The EU AI Act attaches legal force.** Article 12 requires high-risk AI systems to support automatic event logging across their lifetime, with general obligations in effect now, and high-risk application compliance required by December 2, 2027. An identity program must therefore be able to say which agents exist, who owns them, and what each one did, because no log can be produced for an agent no one can find.

The gap this study documents is between those expectations and current practice. Worldwide IDC research identifies the same gap at scale: abused NHIs are now effectively tied with phishing as the leading entry point in confirmed identity incidents. AI agents are the least covered identity type in NHI management programs, less than 1 out of 5 organizations run continuous identity discovery, and just over 1 out of 4 have fully automated identity governance for AI agents (IDC *Worldwide IAM Security Survey*, May 2026, n = 860).

This paper walks through the six problems security leaders described in their own words, the consensus they already share, a four-stage maturity model for locating your own organization, a definition-agnostic operating model you can adopt today, and the questions to put in front of any vendor before you buy. This paper also walks through four capability areas of need identified in the study **(Table 3, page 46)**, a subset of the eighteen capability areas in total that were identified by respondents **(Table 4, page 47)**. Each capability is ranked by demand strength and tiered Critical, High, Moderate, and Selective. The full mapping can be found in the Demand Map at the end of this document **(Table 4, page 47)**.

EU AI Act

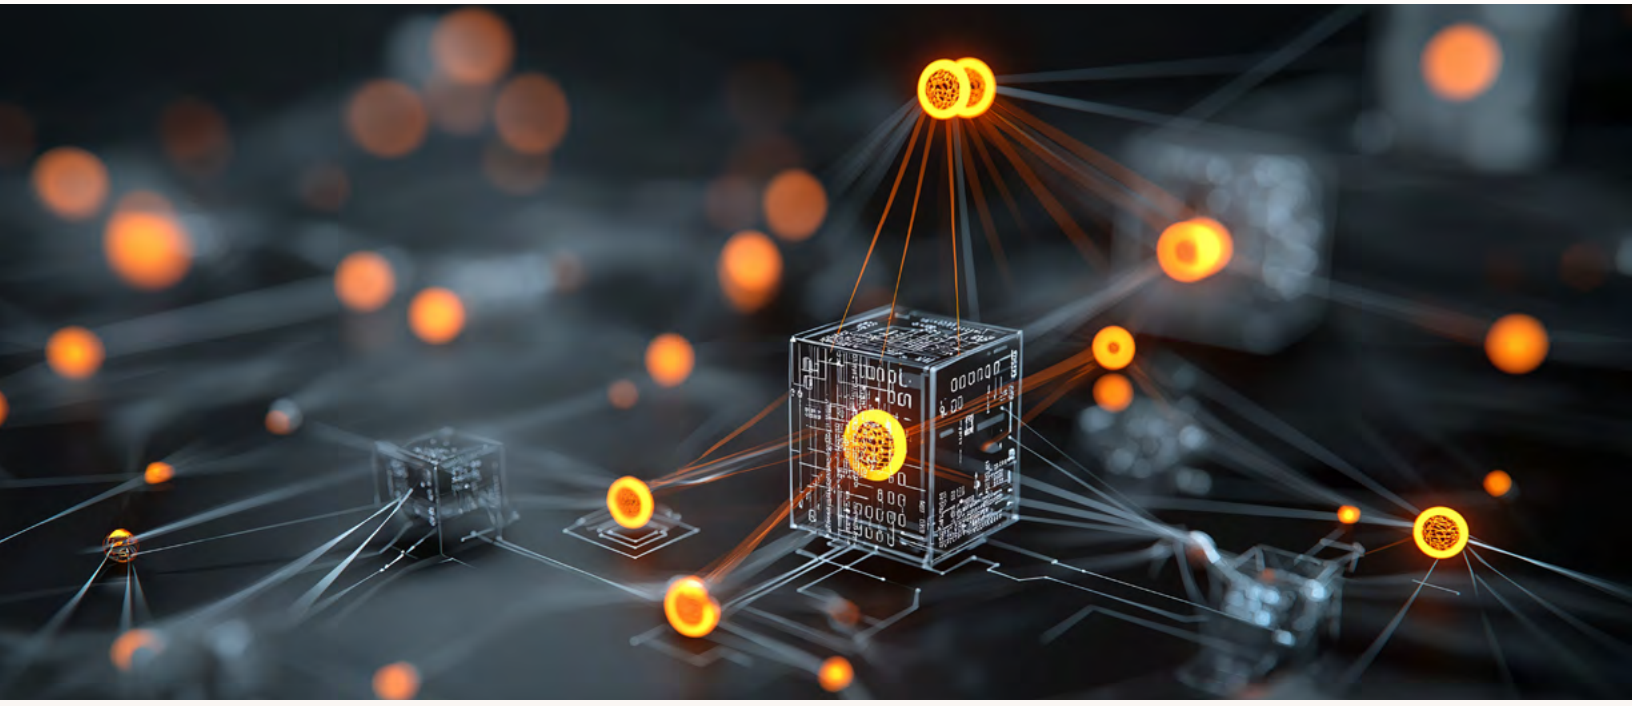

About this research

This paper draws on a three-session, asynchronous expert panel of 13 senior security and identity leaders across financial services, healthcare, technology and media, entertainment, hospitality, education, and regulated enterprise. Participants answered 31 open-ended questions in their own words. Responses were inductively coded, and findings were translated into a demand map of 18 vendor capability areas and 49 discrete needs, as well as a four-stage maturity model.

Quotes are lightly copyedited for spelling, capitalization, and punctuation only; wording is unchanged, and any inserted word is marked in brackets.

Panel findings are corroborated where noted with results from three IDC surveys: the IDC *Worldwide IAM Security Survey (May 2026*, n = 860), the IDC *Unified AI Platforms and Governance Survey (March 2026*, n = 744), and the IDC *Agentic AI Functional Use Case Survey (February 2026*, n = 1,967).

The shift

Identity is a critical control plane, and non-human identities are exploding

Agentic identity is now a board-level and buying-center issue, and the reason is structural. The network perimeter has dissolved; identity, its permissions, and its context are the primary control surface for enterprise risk. Non-human identities, service accounts, application identities, API keys, machine workloads, and now AI agents already outnumber human identities in many environments, and agentic AI accelerates the imbalance. All respondents in this study cited non-human identities outpacing their identity and access management (IAM) programs, with ratios cited as high as 75 to 1.

Agents will push that figure higher.

It helps to be precise about where identity sits among the control planes that govern agentic AI, because three matter most: **data, runtime, and identity.**

- **Data controls** decide what an agent may access: classifications, entitlements, and boundaries around sensitive information
- **Runtime controls** decide what an agent may do as it executes: guardrails, approvals, monitoring, and containment

Both are indispensable, and both inherit their subject from the identity plane.

- A data permission must be scoped to some actor
- A runtime guardrail must bind to some actor
- An audit record must name some actor

When the identity beneath those controls is borrowed, shared, or unknown, data controls scope to the wrong subject, runtime controls cannot be enforced or tuned per agent, and the audit trail attributes actions to the wrong entity. **Therefore, Identity is not merely first among the three; it is the plane on which the other two operate.**

Agents are not ordinary service accounts. They reason over multi-step goals, call tools and APIs, act on behalf of users, and, in the course of completing complex tasks, can behave unpredictably. The broader market data confirms both the stakes and the priority. IDC research shows that abused non-human identities, meaning service accounts, API keys, and tokens, were the initial entry point in 19.0% of the most recent identity incidents among the 651 organizations reporting a confirmed incident, effectively tied with phished or stolen credentials at 19.5%. NHI compromise is no longer a theoretical exposure. Buyers have noticed: NHI and AI agent security is now a top two IAM program priority for 43.7% of organizations over the next 12 to 24 months, alongside identity governance modernization (50.7%) and privileged access management (PAM) modernization (50.6%) (IDC *Worldwide IAM Security Survey*, May 2026, n = 860).

Voice of the customer

"Identity has moved from being an IT administration function to a core security control."

Security Architect, Financial Services

"I would estimate a 75 to 1 ratio of NHI to human identities with expectations that AI agents will drive this higher."

Security Architect, Highly Regulated Enterprise

"With cloud, SaaS, APIs, and now AI agents, the perimeter is no longer the network and, fun fact, now it's the identity, the permissions attached to it, and the context around what that identity is doing."

Security Leader, Technology and Media

"That is where our IAM strategy has to evolve beyond human users and static access models into machine identity, agent identity, least privilege, policy enforcement, and continuous monitoring."

Security Leader, Technology and Media

"This surge in conversations stems from the fact that every AI initiative and project currently underway involves either creating or integrating with AI agents."

Security Architect, Enterprise IT

"AI agents are being deployed faster than normal security governance can keep up."

Security Architect, Financial Services

What this means for organizational leaders: Recognize that identity is a critical control plane for AI, and that agents are a new identity class to be secured in their own right, not an extension of existing accounts. Identity, data, and runtime controls will all be needed, and identity is the one that makes the other two work. Every downstream decision in this paper follows from that recognition.

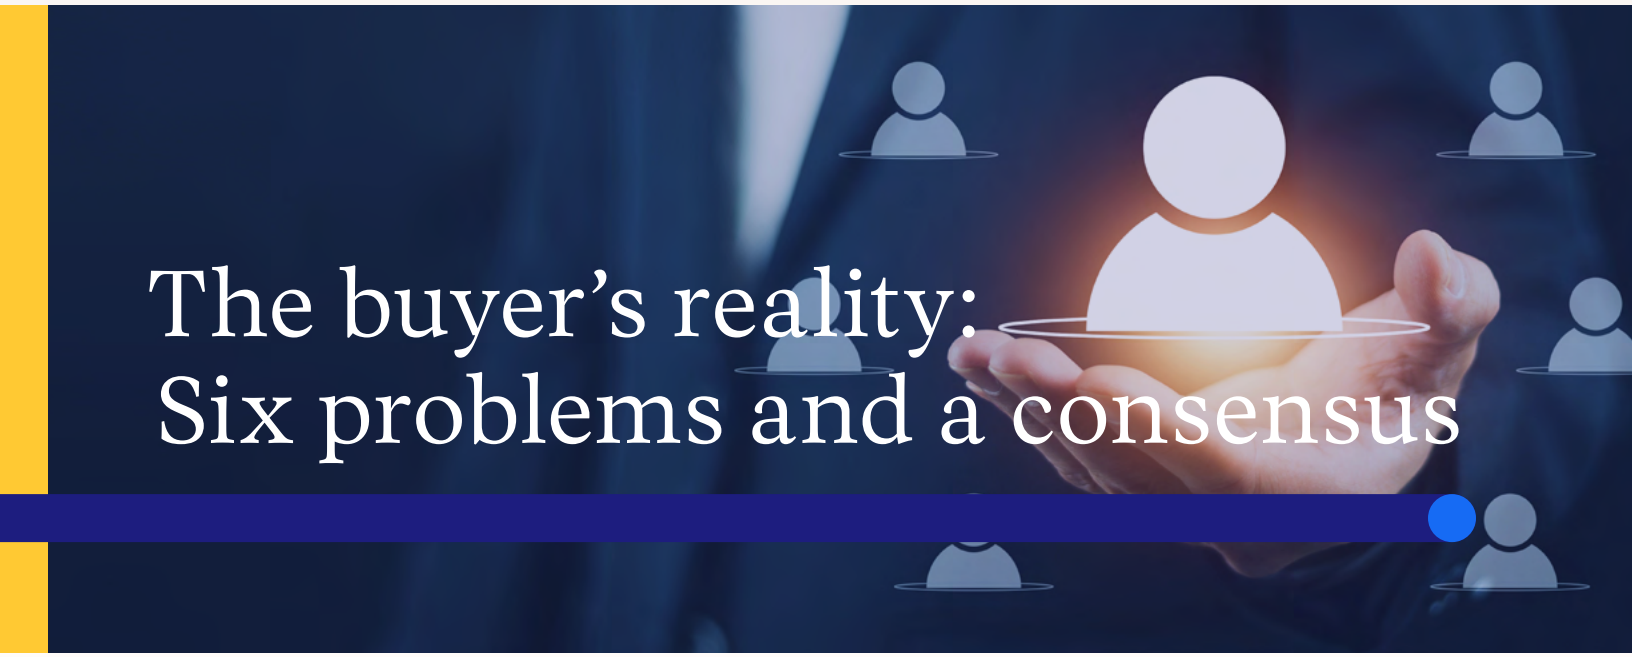

The buyer's reality: Six problems and a consensus

"I cannot see or count my agents"

The need: Automated discovery and an authoritative inventory.

Most organizations cannot state an exact agent count. The sprawl is a product of ease and decentralization. Business units and developers stand up agents faster than any central team can track them, and the hardest agents to see are those running under a user's identity, with no service account attached and no separate credential trail. Security and compliance now require what few programs have: a real-time, inventory recording each agent's owner, business purpose, data accessed, permissions, credentials, approval status, expiration, and monitoring state. The evidence on this point is unanimous. Every respondent flagged the issue of the inability to count the total number of agents operating in their environment. The Demand Map **(Table 4, page 47)**, says the same thing from the other direction. Agent and NHI discovery and visibility drew all respondents, one of the broadest areas of need in the study.

Worldwide research suggests the visibility gap is even more stark. A substantial 77.3% of organizations report high or very high confidence that they can see all human and non-human identities across on-premises, cloud, and SaaS. Yet the same survey shows that only 18.5% run identity discovery continuously or in near real time, inventory or ownership gaps rank among the most cited non-human identity challenges (42.0%), and bots, RPA, and AI agents are the least covered non-human identity type in current management

programs, secured by only 41.5% of organizations. High confidence built on quarterly snapshots of the categories that exclude agents is exactly the false assurance the study panel's discovery stories puncture.

The OWASP Non-Human Identities Top 10 makes improper offboarding its number one risk (NHI1:2025), and says an identity that never entered an inventory cannot be offboarded at all. An authoritative inventory is the implicit precondition for nearly every control on that list.

NIST's AI Risk Management Framework calls for organizations to maintain an inventory of their AI systems as a core Govern-function practice, and NIST's 2026 work on AI agents treats formal identification and enterprise-grade lifecycle management of agents as the baseline, which presumes you can see and count them.

The EU AI Act points in the same direction: Article 12 requires high-risk AI systems to support automatic event logging across their lifetime, an obligation that applies from December 2, 2027. But no organization can log, monitor, or produce records for agents it has not discovered.

In short, study participants described the gap between their own experience and the controls that regulators and standards bodies already consider table stakes.

Voice of the customer

It is true for us. We do not have an authoritative agent inventory."

C-Suite, Education and Care Services

"We do not have a good handle on agents that are running in our environment, particularly ones that run as the user so they don't have a service account tied to them."

Identity Architect, Healthcare

We are handling AI agents by first trying to improve visibility: where they exist, what systems they connect to, what data they access, and whether they are tied to a clear owner."

Security Architect, Financial Services

Simply highlighting the fact that we cannot truly tell which/how many agents exist is motivating for internal leaders."

Identity Leader, Hospitality

IAM became a board-level priority for me when my discovery report revealed that we have over 500 AI agents operating within our network and systems, far more than the handful I initially estimated."

Security Architect, Enterprise IT

My organization is prioritizing visibility over control - you cannot control what you cannot see/know about."

Identity Leader, Hospitality

If you asked me to swear to an exact count including everything devs and business units have spun up, I couldn't do it with full confidence."

Security Leader, Technology and Media

What organizational leaders should be asking for: Automated discovery and a continuously updated, authoritative inventory spanning SaaS, cloud, and your vendor stack. If a discovery assessment surprises you the way it surprised the architect quoted above, treat that as the start of the program, not an embarrassment.

"Agents break my identity and privileged-access tooling, built for humans"

The need: Agents as a first-class identity, and PAM and IAM reinvented for machine speed.

Identity governance and privileged access tools were designed around the rhythm of human employment. A person is hired, changes roles, and eventually leaves, and at each of those moments the tools grant, adjust, or revoke access. Agents have none of those moments. No one hires an agent, it has no manager and no job title, and it never resigns, so the tools have no natural point at which to act, and no obvious person to hold accountable. The tools also assume that software behaves predictably, running the same task the same way on a consistent schedule. However, autonomous AI agents decide what to do next, so the permissions that fit yesterday's behavior do not fit today's. Study respondents were clear. Existing controls were not built for this kind of actor, and the data confirms how widely that problem is felt. Agent and NHI lifecycle and ownership was a critical demand area, raised by all research participants.

What security leaders want next is equally plain. PAM today mostly means keeping powerful credentials locked in a vault and checking them out when someone needs them. That model worked when that someone was a human administrator at a keyboard. For agents operating at machine speed, the vault has to become more like a dispatcher, issuing short-lived, tightly scoped permissions the moment a task begins and taking them back the moment it ends.

Voice of the customer

"It needs a new architecture and a new design."

Security Architect, Enterprise IT

"PAM vendors need to shift from plain old secrets management to being part of real-time process as part of the AI control plane."

Security Architect, Highly Regulated Enterprise

"We manage AI agents as a separate identity type, and have specific conventions for each agent instance too."

Security Architect, Technology

"Our identity and access management approach with AI agents has changed because these are now non[-]human entities and not just system accounts that they utilize."

Security Leader

"IGA or PAM tools break down in several areas, including accountability (ownership), lifecycle management, context, and finally, behavior."

C-Suite, Financial Services

"Today they are mostly managed through existing service account or application identity models, but we are moving toward treating them as a distinct identity type because they can access data, call APIs, and act on behalf of users."

Security Architect, Financial Services

"PAM has to change in several areas to address the power of agents and their ability to take action exponentially faster than humans."

C-Suite, Financial Services

Survey data quantifies how far the market is from closing the category gap. Worldwide IDC research shows only 27.3% of organizations have fully automated identity governance and access controls for AI agents; 37.8% have policy-driven automation, 27.1% are still on role-based controls, and 7.8% are at basic manual controls or have not started. Nearly three quarters of organizations are governing a machine-speed identity class with something less than machine-speed controls.

Figure 1

Identity governance and access-control maturity for AI agents

G3. How mature is your organization in applying identity governance and access controls to AI agents?

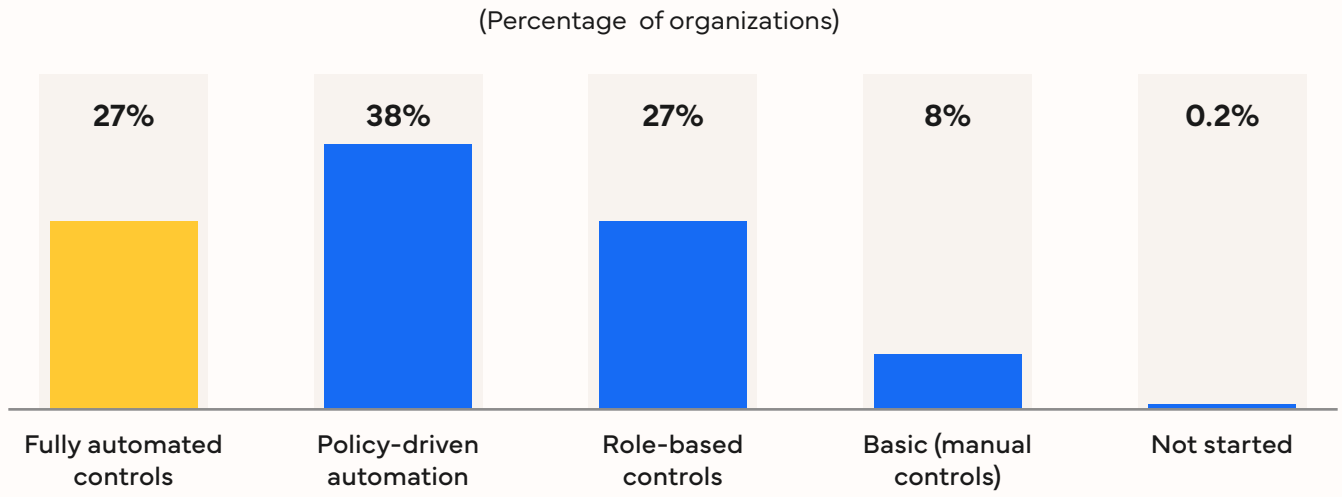

n = 860 (all organizations); Source: IDC *Worldwide IAM Security Survey*, May 2026 (EU25120009-S), unweighted

Here again, the frameworks have already moved. NIST's NCCoE 2026 concept paper on software and AI agent identity proposes treating agents as distinct NHIs requiring enterprise-grade lifecycle management, with authorization built on OAuth 2.0, OpenID Connect, SPIFFE and SPIRE, and attribute-based access control. And the OWASP Top 10 for Agentic Applications 2026 names identity and privilege abuse (ASI03) as a top risk precisely because agents inherit sessions and secrets from human-era tooling, producing privilege escalation and attribution gaps. The category gap respondents describe is the same gap these bodies are now writing guidance to close.

NCCoE 2026 concept paper on software and AI agent identity

What organizational leaders should be asking for: Solutions that model agents as a distinct, governable identity, and privileged access re-engineered for dynamic, machine-speed non-human identities. If a vendor's answer to agents is "put them in the vault," that is a human-era answer to a machine-era problem.

"Privilege is sprawling at machine speed"

The need: Automated right-sizing, just-in-time access, and zero standing privilege.

Agents accumulate more privilege than the people who built them, and AI-assisted provisioning can multiply that privilege at machine speed. The target state respondents described is task-scoped, time-bound, least-privilege access, with a clean return to zero standing privilege when the task ends. Just-in-time and ephemeral credentials work well for cloud-native agents but break against legacy systems that only understand standing service accounts. So the realistic program is a mix: ephemeral where the platform allows it, tightly governed standing access where it does not. Privilege right-sizing and just-in-time (JIT) access is a Critical demand area.

Voice of the customer

"AI just provisions access faster and creates privilege sprawl at machine speed."

Security Architect, Financial Services

"They need to be task-scoped, time-bound and have least privilege at the action/tool level."

C-Suite, Education and Care Services

"The new strategy needs stronger discovery, ownership, least privilege in real time, expiration, behavioral monitoring, and rapid revocation for every identity."

C-Suite, Financial Services

"PAM for AI really just in [a] nutshell means moving away from the classic vault-and-checkout model built for a human admin at a keyboard, toward brokering short-lived, identity-bound credentials for workloads moving at machine speed."

Security Leader, Technology and Media

"If I defer this decision, I may end up over-assigning permissions for these dynamic identities, which would increase our risk of unauthorized access and lateral movement within our systems."

Security Architect, Enterprise IT

The target state respondents describe maps directly onto the OWASP Non-Human Identities Top 10, which flags overprivileged NHIs (NHI5:2025) and long-lived secrets (NHI7:2025) as core risks. What the panel calls privilege sprawl at machine speed is those two risks compounding as the agent population grows.

The survey data shows how few organizations have reached the target state, and how fast the pressure is building. Only 18.0% of organizations enforce least privilege through just-in-time access with automated remediation; 34.4% use automated entitlement right-sizing, and nearly half still rely on policy-driven (31.4%) or manual (16.2%) access

reviews (IDC Worldwide IAM Security Survey, May 2026, n = 829). Meanwhile the autonomy those privileges attach to is ratcheting upward: 29.1% of organizations already allow agents to make decisions in most areas with human oversight only for critical issues or to operate with full autonomy on some critical issues, and 97.5% are discussing granting agents more autonomy, 67.9% actively (IDC *Agentic AI Functional Use Case Study*, February 2026, n = 1,967). Autonomy is expanding faster than privilege discipline is maturing, which is precisely the sprawl dynamic the panel describes.

Figure 2

How organizations enforce least privilege today

C1. Which best describes how your organization enforces least privilege today?

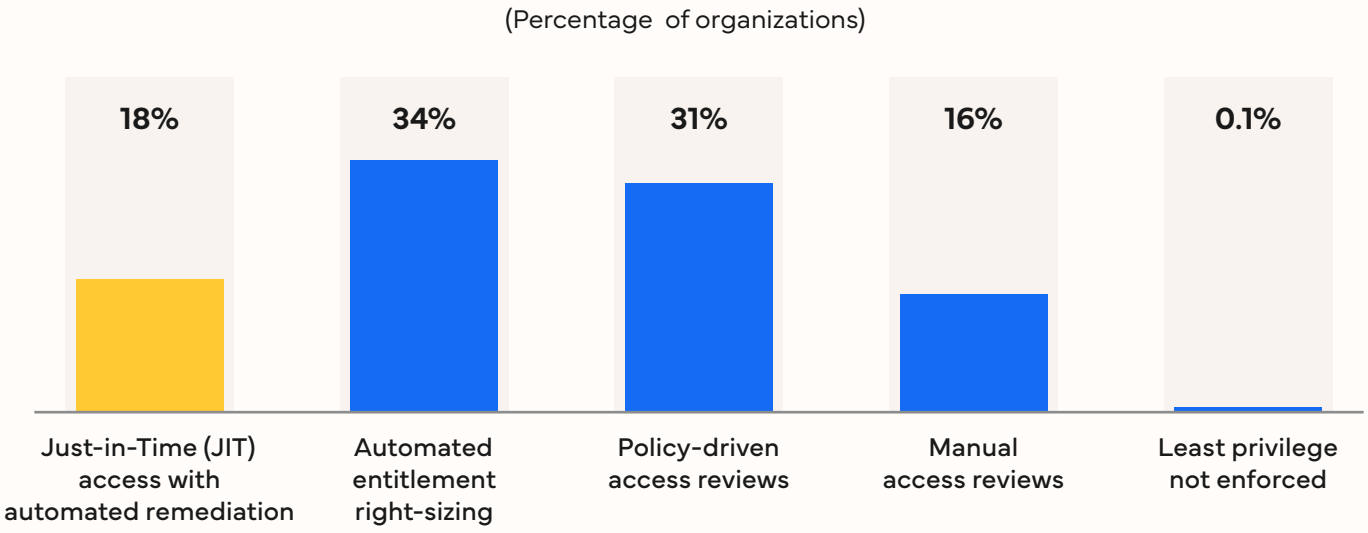

n = 829 (all organizations); Source: IDC *Worldwide IAM Security Survey*, May 2026 (EU25120009-S), unweighted

What organizational leaders should be asking for: Automated least-privilege and right-sizing, ephemeral and just-in-time credentials, and scoped secret and token management, with an honest account from the vendor of how and where the model degrades against your legacy estate.

“I can no longer trust the human signal”

The need: identity proofing and verification that survive deepfakes.

AI makes impersonation cheap and convincing. The highest-ranked concern among respondents is deepfake attack against the help desk to gain access, and against executives and accounts payable for fraud. Several organizations in the respondent panel have already experienced attempted AI-augmented deepfake attacks.

The defense that emerged is to stop trusting the human identity signal itself and to continuously triangulate independent data points to assemble confidence: liveness checks, step-up authentication, out-of-band approvals, and behavioral or continuous signals that no single fabricated artifact can satisfy. Respondents also expect vendor detection models to be tested against current adversarial techniques, not just known samples. AI-driven threat detection and identity threat response is a critical demand area; deepfake and synthetic-identity defense is a Moderate area of demand. A concern nearly all respondents named but few dwelled on, and one that over-indexes among organizations still extending existing identity constructs.

Voice of the customer

Highest concern: Deepfake attack against helpdesk to gain access, or against senior executive[s] and accounts payable for fraud."

C-Suite, Education and Care Services

We've had a couple of instances of deepfake audio impersonations against several senior executives and are looking into products that can help detect this."

Security Leader, Entertainment and Media

Controls like stronger identity proofing, liveness checks, step-up authentication, and out-of-band approvals become more important."

Security Architect, Financial Services

AI has changed our identity threat model by making impersonation faster, more believable, and harder to detect with legacy controls."

Security Architect, Financial Services

A deepfake voice or legit-looking campaign targeting specific helpdesk admins puts us at the mercy of the weakest link in the user population."

C-Suite, Financial Services

The rise of deepfakes has compelled us to revise and update our verification processes for both incoming and outgoing video and voice calls."

Security Architect, Enterprise IT

Phantom identities designed to mimic service accounts are a top concern."

Security Architect, Technology

One example is DPRK (North Korean) IT worker infiltration of companies which has been an on-going concern."

Security Architect, Highly Regulated Enterprise

The market is already voting with its roadmaps. Identity threat detection and response is the single most selected identity security priority for the next 12 to 24 months, chosen by 66.5% of organizations, ahead of identity resilience and recovery (52.7%) and phishing-resistant or passwordless authentication (39.9%) (IDC *Worldwide IAM Security Survey*, May 2026, n = 860).

Figure 3

Identity security program priorities, next 12 to 24 months

N2. What are the highest priorities for your identity security program over the next 12 to 24 months? [Select up to 2]

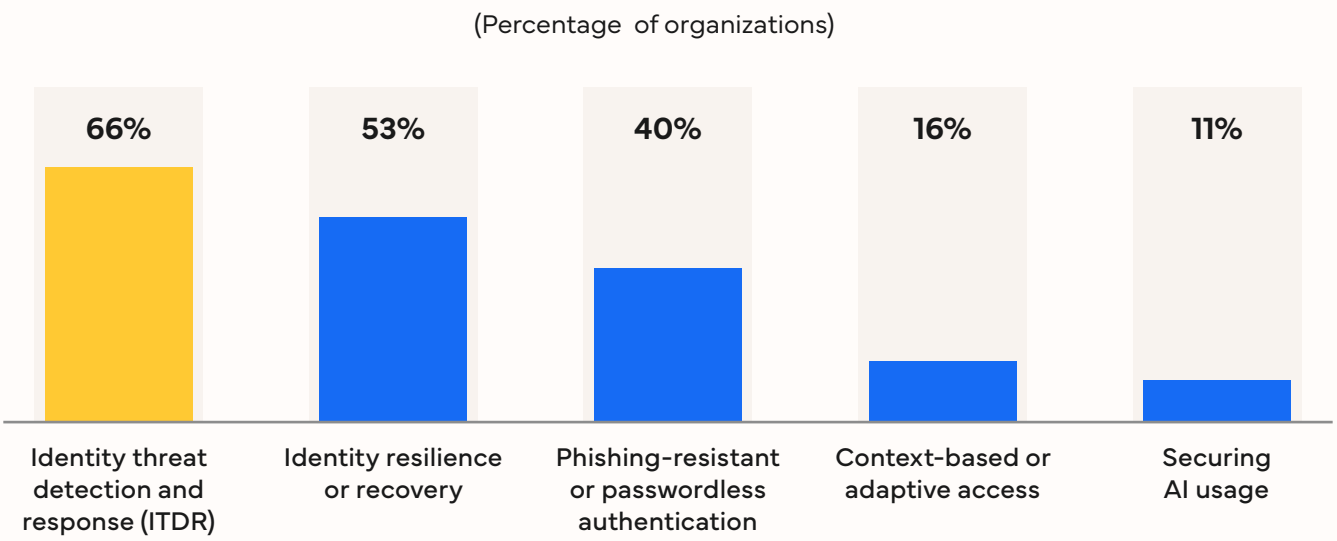

n = 860 (all organizations); multiple responses allowed, totals exceed 100%. Source: IDC *Worldwide IAM Security Survey,* May 2026 (EU25120009-S), unweighted

What organizational leaders should be asking for: Deepfake and synthetic-identity detection, hardened identity proofing and help desk verification procedures, and identity threat detection and response that operates at machine speed.

"Governance and ownership lag adoption, and no single vendor covers it"

The need: an end-to-end governance framework and an integrated platform, not point tools.

Agents entered the environment under business pressure faster than security could write policy. Ownership of AI identity governance is contested across the Chief Information Security Officer's Office (CISO), Identity and Access Management (IAM), Governance, Risk, and Compliance (GRC) and the business, and no single vendor covers the full framework. Respondents were emphatic that governance cannot be a separate program; it has to connect identity, data, cloud, and third-party risk, groups that historically worked within their own siloes.

The evidence here is the strongest in the study. Governance, risk, and policy framework for agentic AI is the single strongest demand area, and identity platform consolidation is a high-tier area.

Voice of the customer

"The gap is common because the AI tools slipped into the environment unnoticed until they were."

Identity Architect, Healthcare

"Most companies are in a race to deploy agents with little regard for security, governance, auditability, [and] explainability."

Security Architect, Technology

"AI governance cannot be treated as a separate program from identity, data, cloud, and third-party risk."

Security Architect, Financial Services

"The gap can be immense, but the worst part of it is if we learn too late, and the maturity of these shadow AI projects now has production data, the reversal of these poor practices is costly."

C-Suite, Financial Services

"The gap between visibility and control is the cost of compliance. Because you cannot comply if you do not have control."

Security Architect, Technology

"[O]rganizations have embraced AI tools quickly, and very haphazardly, and with IT and security resources stretched thin, governance has taken a backseat to business priorities."

Security Leader, Entertainment and Media

The contested-ownership finding replicates almost perfectly in the quantitative data. Asked who is primarily responsible for AI governance, no single role or body reaches 15%: a dedicated AI governance team leads at 14.5%, followed by an AI center of excellence at 13.4%, the CIO/CTO at 11.7%, a chief AI officer at 11.4%, AI/ML engineers at 11.2%, and the CEO/CFO at 10.9%, with 6.0% saying no one person is responsible (IDC *Unified AI Platforms and Governance Survey*, March 2026, n = 744).

Identity strategy, by contrast, has a clearer home. The CISO organization owns it in 45.8% of organizations, central IT in 29.1% (IDC *Worldwide IAM Security Survey*, May 2026, n = 860).

Figure 4

Identity strategy has a clear home; AI governance does not

Panel A: P1. Who has primary ownership for enterprise identity strategy in your organization?
Panel B: D02. Who is primarily responsible at your organization for AI governance?

Panel A: Enterprise identity strategy owner

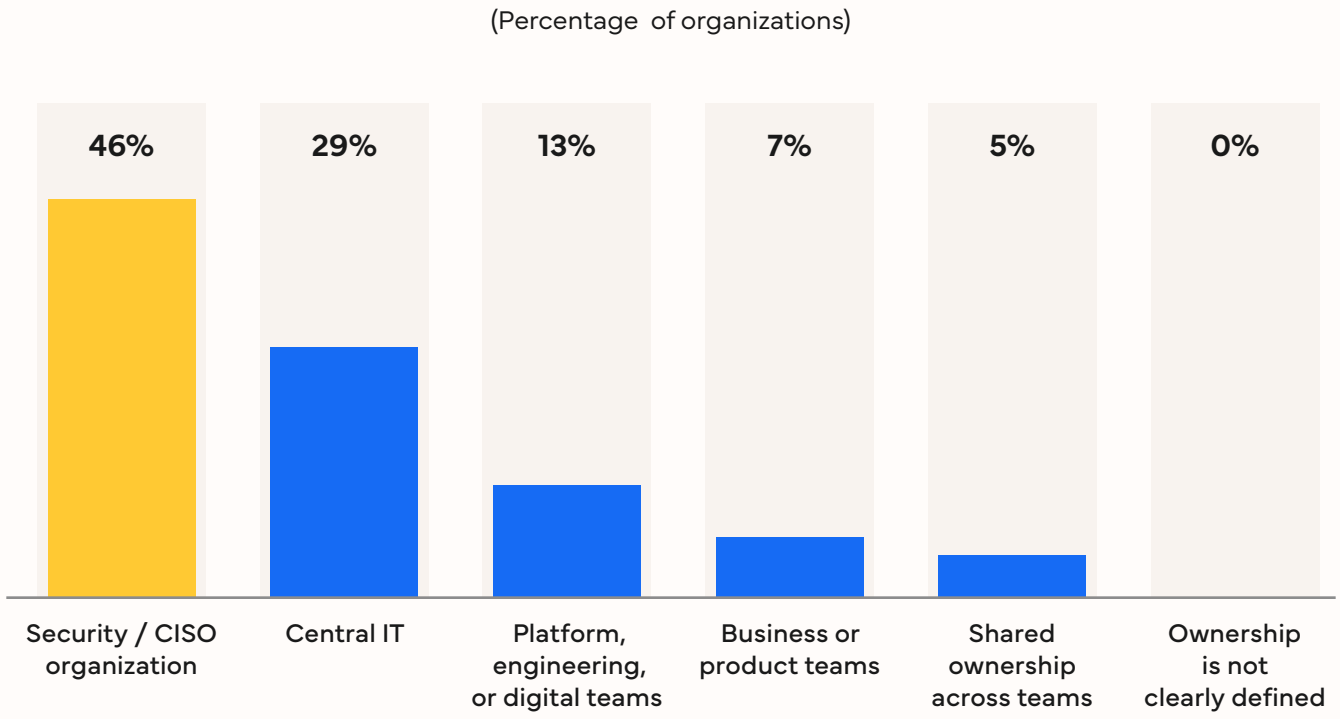

n = 860; Source: IDC *Worldwide IAM Security Survey*, May 2026 (EU25120009-S), unweighted

Panel B: Primarily responsible for AI governance

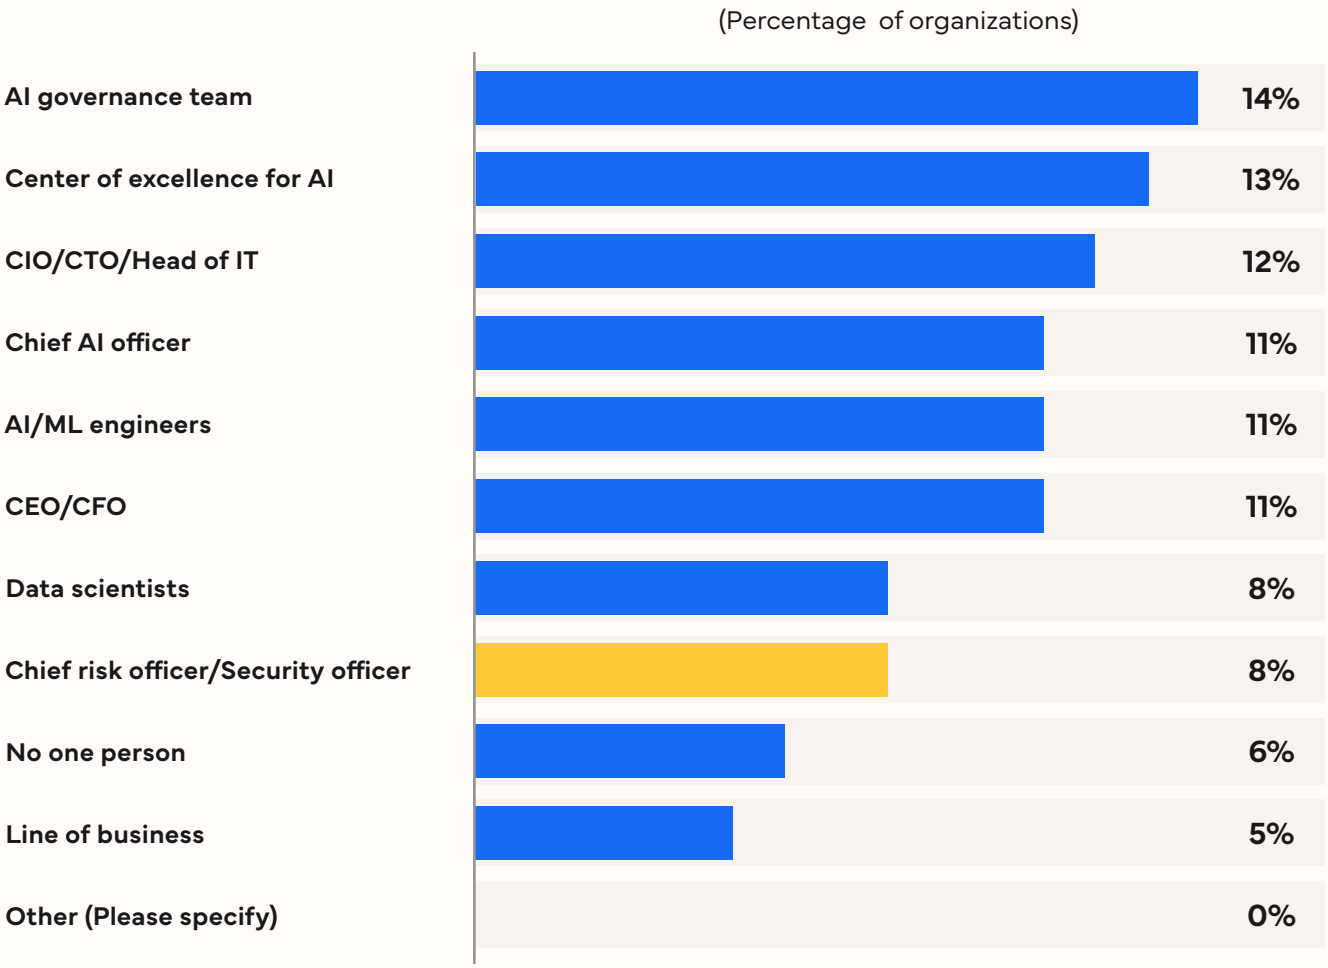

n = 744; Source: IDC *Unified AI Platforms and Governance Survey*, March 2026 (2600310), unweighted

The problem is that agentic identity sits exactly at the seam between those two ownership maps, one consolidated and one fragmented, which is why the panel experiences it as contested. The same surveys confirm the pressure behind the governance demand: security and compliance is the most cited challenge in implementing agentic AI initiatives (40.7%, IDC *Agentic AI Functional Use Case Study,* February 2026, n = 1,967), and security and data leakage (32.1%), and regulatory compliance, including the EU AI Act, (30.5%), are the top two AI risk concerns overall.

Figure 5

Key challenges in implementing agentic AI initiatives

A14. What key challenges do you face in implementing the current agentic AI initiatives? [Select all that apply]

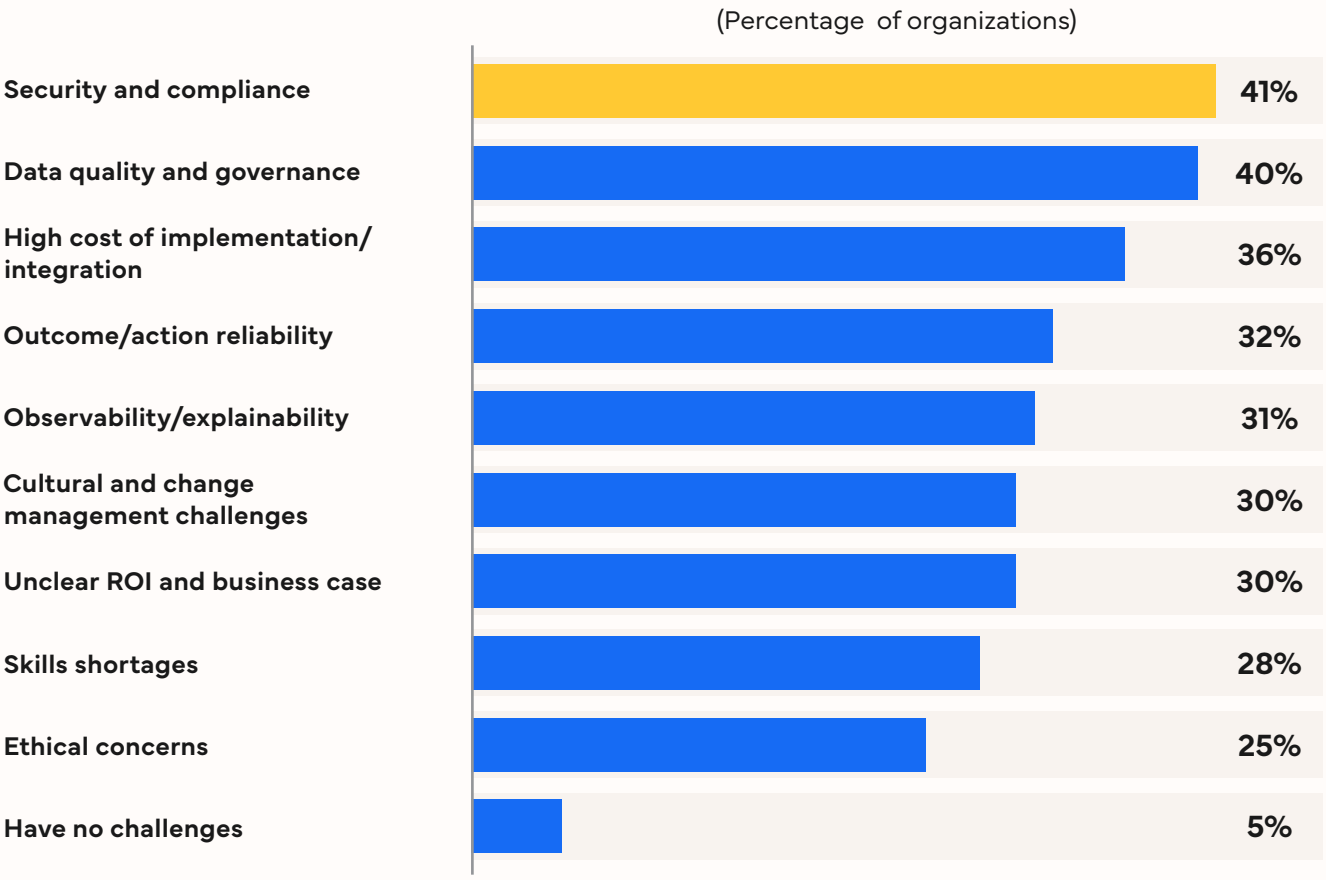

n = 1,967 (all organizations); multiple responses allowed, totals exceed 100%. Source: IDC *Agentic AI Functional Use Case Study*, February 2026 (AM25110001), weighted

The regulatory clock reinforces the urgency. The EU AI Act's obligations for high-risk AI systems, including the Article 12 requirement for automatic event logging over the system's lifetime, apply from December 2, 2027. NIST's control overlays for securing AI systems, with dedicated overlays for single-agent and multi-agent deployments, are in active development, and NIST recommends organizations map their deployments to the AI Risk Management Framework now rather than wait for finalized standards. Regulation is already the dominant force on identity budgets: 39.9% of organizations name regulatory or external requirements as the primary driver of IAM and identity security investment, and another 45.8% name them as one of several key drivers (IDC *Worldwide IAM Security Survey*, May 2026, n = 860). Governance frameworks are no longer something the market must invent; they are something the organization must operationalize.

Figure 6

Regulatory and external requirements as a driver of identity investment

Q1. To what extent are regulatory or external requirements influencing your organization's IAM and identity security investments?

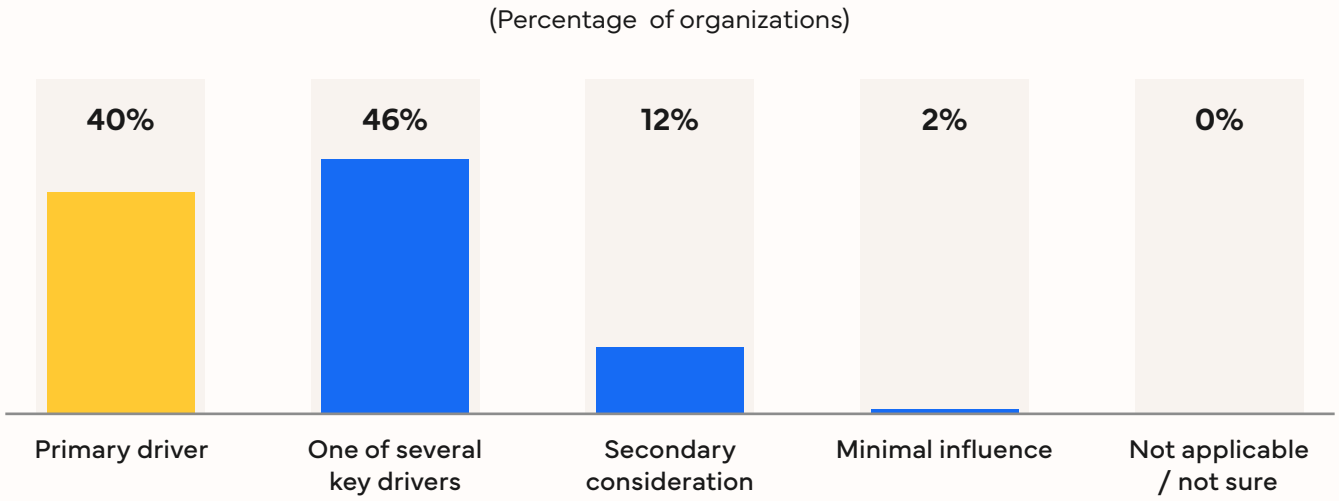

n = 860 (all organizations); Source: IDC *Worldwide IAM Security Survey*, May 2026 (EU25120009-S), unweighted

What organizational leaders should be asking for: A governance, policy, and liability framework mapped to standards such as NIST guidance, ISO/IEC 42001, the CSA AI Controls Matrix, and the EU AI Act, and consolidation of point tools into an orchestrated platform. Settle internal ownership of the program early; a contested mandate is itself a security gap.

"There is risk I do not control: Shadow AI and vendor-embedded AI"

The need: Shadow AI detection and controls over identities your vendors introduce.

Three forces put agents in your environment that you never provisioned. First, end users and developers can spin up agents on their own. Second, AI is increasingly embedded in the SaaS tools the organization already buys, running under identities the customer never sees. For vendor-embedded AI, control leans on contracts and monitoring more than on direct telemetry. A third force involves no agent at all: employees moving corporate data into consumer AI platforms under identities the organization does not control.

The exposure compounds in two distinct ways. First, on consumer tiers of major AI assistants, conversation data may be used for model training by default, subject to terms the employee has typically never read, and the organization never negotiated. Second, and less intuitively, the "share" features built into these platforms can quietly convert a private conversation into a public document. In July 2026, users discovered that Claude conversations and artifacts they had shared via link, intending them for a specific colleague, had been indexed by Google Search once those links circulated on the open web. The indexed material included business strategies, cloud infrastructure diagrams, and clinical trial planning documents. An employee took a routine collaboration step without intending to make the information public, but the system's settings assumed discoverability was desired.

This exposure is an identity problem, not merely a data-loss problem. A personal AI account is an unmanaged identity holding corporate data outside every boundary the organization can see, with no sponsor, no lifecycle, no revocation path, and no audit trail.

The same registry-and-ownership discipline this paper describes for agents is the eventual answer here as well, but the immediate controls are nearer to hand: visibility into which AI platforms employees actually use, enterprise tenants with training-data protections and sharing restrictions in place of ungoverned personal accounts, and user education that a share link is a publication decision, not a private handoff. Shadow AI detection and governance emerged in this research as a High tier area of need; third-party and vendor AI identity risk a s a Moderate area of need.

Voice of the customer

"It's one thing to identify ChatGPT/Gemini/Copilot use, but it's more difficult to identify third-party vendors that are embedding AI agents into the SaaS apps that we could be using."

Security Architect, Highly Regulated Enterprise

"What's driving the visibility gap here is the ability to keep people contained and not use commercially available products without approval or prior to analysis."

Security Leader

"We will need network or agent-based monitoring tools that can detect AI patterns on user workstations, browsers, etc."

Identity Leader, Hospitality

"This is not just shadow AI, this is totally dark AI without the right network-level observability solutions."

C-Suite, Education and Care Services

"The harder areas are browser extensions, unsanctioned SaaS tools, personal AI accounts, unmanaged plugins, service accounts, and business-led projects outside of IT Security."

C-Suite, Financial Services

"Shadow AI differs from traditional shadow IT in that it can access data, generate content, call APIs, create workflows, and act as a non-human identity."

C-Suite, Financial Services

The survey base shows most organizations are still early in this work: 26.7% describe their AI governance program as focused on discovering and inventorying AI usage, explicitly including cataloging shadow AI, while only 11.7% have progressed to measuring and optimizing value (IDC *Unified AI Platforms and Governance Survey*, March 2026, n = 744). A quarter of the market is still finding its AI, let alone governing it.

Figure 7

Current state of AI governance programs

D05. Which of the following best describes the current state of your organization's AI governance program?

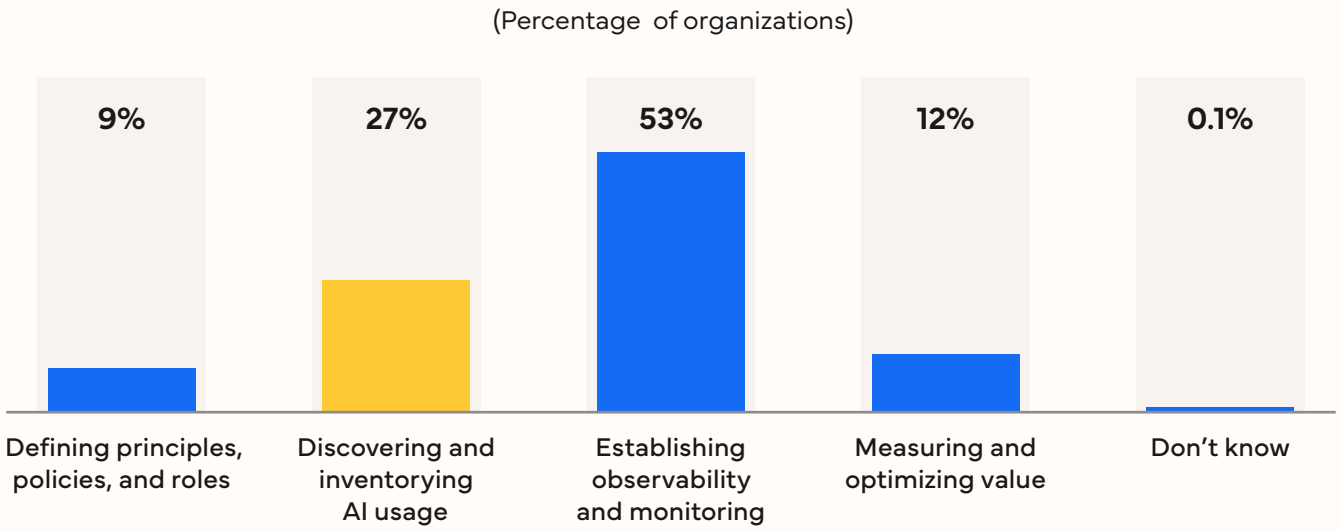

n = 744 (all organizations); Source: IDC *Unified AI Platforms and Governance Survey*, March 2026 (2600310), unweighted

What organizational leaders should be asking for: Shadow AI discovery and monitoring; visibility and contractual controls over vendor-introduced AI identities, including no data retention, ephemeral access, and audit rights written into the agreement; and managed enterprise tenants for the consumer AI platforms employees already use, so that corporate data flows through identities the organization governs rather than personal accounts it cannot see.

The one thing buyers already agree on: Every agent needs a named human sponsor and a managed lifecycle

This, the strongest consensus in the study, is procedural and actionable today, and it requires no decision about what an agent identity ultimately is. The pattern is consistent across organizations: gated intake, an approved business purpose, scoped access, expiration or recertification, and a clean revocation path, with ownership binding accountability to a named person.

A subset of respondents, often unprompted, pushed one step past the problems this paper has cataloged, to the one that arrives next: multi-agent chains. That is when one agent invokes another, and in which those agents invoke a tool, and the questions of whose authority is being exercised, and who is accountable for the outcome, blur with every hop. The blur is not abstract. Every control described in this paper assumes an answer to a simple question: Who is acting?

Delegation chains complicate the answer at each link.

- The sponsor who approved agent A never approved the actions of agent B
- Privileges granted for one task flow silently into another
- Revoking the first agent may or may not stop the second
- When something goes wrong, the investigation must reconstruct a chain of authority that no system recorded at the time

No established governance solution yet exists, and the evidence shows how early this territory is. Coverage of the agent-to-agent frontier was among the thinnest in the coded corpus. That, in itself, is a signal of where current practice is least developed.

It is the most mature organizations in the study, those already running a distinct governed identity class, whose next area of need is agent-to-agent readiness. The standards bodies have reached the same frontier from the other side: NIST's control overlays for securing AI systems include a dedicated multi-agent overlay. Its concept paper on agent identity names task-scoped privilege delegation in multi-agent systems as a requirement. And the OWASP Top 10 for Agentic Applications reserves entries for insecure inter-agent

communication and cascading failures. Industry is moving too: The Agent2Agent (A2A) protocol, initiated by Google and now governed by the Linux Foundation with backing from leading cloud service providers (CSPs) and enterprise software vendors, reached version 1.0 in March 2026, giving agents a standard transport for discovering and invoking one another, with cryptographically signed agent cards to verify an agent's identity across organizational boundaries. But a transport standard answers how agents talk and how an agent proves what it is. It does not answer whose authority a delegated action carries, which sponsor is accountable at the third hop, or how a revocation propagates down the chain. This is a problem that will become mainstream in short order.

This is also why the consensus matters beyond its own terms. A chain of agents that each carry a named sponsor, a scoped grant, and a recorded delegation lineage is a chain that can be reconstructed, constrained, and revoked. A chain of borrowed credentials is unaccountable at the first hop and unrecoverable by the third. The consensus is not just the answer to today's problem. It is the foundation on which the next one will be solved.

What organizational leaders should be asking for: A turnkey sponsor-and-lifecycle wrapper and an agent registry that binds every agent to an owner and a lifecycle, deliverable before the architecture question is settled, together with a path to the highest level of AI maturity appropriate for your organization.

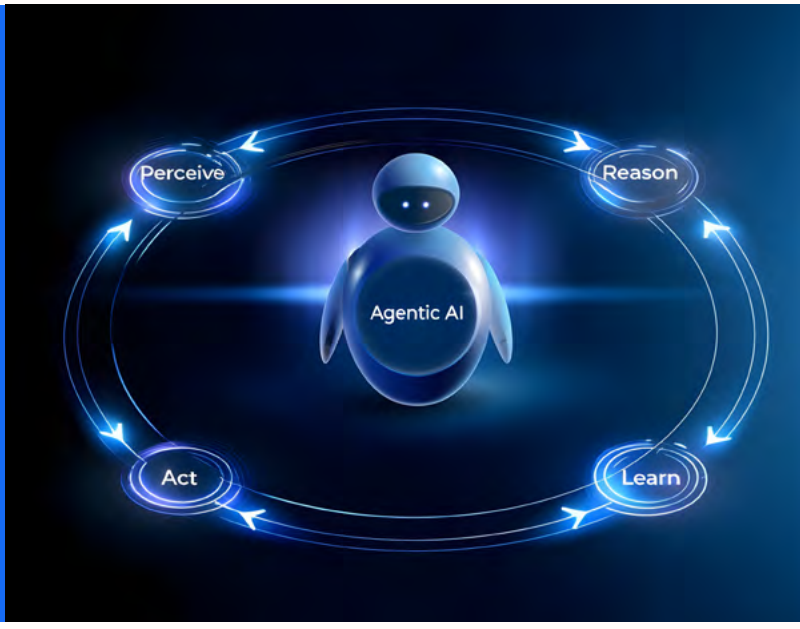

Voice of the customer

"Every agent should have a clear business purpose, a human owner, and approved data access with clear guidelines."

C-Suite, Financial Services

"We manage agent identity as a lifecycle issue, not just a technical credential."

Security Architect, Financial Services

"Each agent type has a sponsor who is accountable for its actions."

Security Architect, Technology

"The next wave is multi-agent chains where agent A invokes agent B invokes a tool, and suddenly the question of whose authority is being exercised and who's accountable for the outcome gets really murky."

Security Leader, Technology and Media

"Because it is an assistant to a human being, the business owner and the system operator are liable."

Security Architect, Enterprise IT

"The internal message is simple: we cannot secure or govern AI agents unless we can identify them, assign ownership, limit access, and monitor behavior."

Security Architect, Financial Services

"The risk of deferring is that we end up with unmanaged agents, excessive permissions, unclear ownership, and poor auditability."

Security Architect, Financial Services

Start by understanding your organization's AI maturity level

The maturity curve

Organizations in this study modeled agent identity in visibly different ways. Some treat agents as a distinct, first-class identity type. Others reuse the existing service-account or application-identity construct, often while acknowledging that the fit is poor. And some run agents on ad hoc or borrowed identities, including the invoking user's credentials.

The essential finding is that this variation is a maturity gap, not a standing disagreement. Working definitions co-vary with capability and move in one direction, from borrowed and service-account constructs toward a purpose-built, distinct identity class. In the coded corpus data, service-account language peaks at Stage 2 **(Table 1, next page)** and falls to zero by Stage 4, while distinct-class language rises steadily to its peak at Stage 4. The destination is shared; the starting point is not. Your current approach to agent identity is therefore a useful indicator of your organization's overall AI maturity, and the practical move is to locate your stage and take the next step it indicates. Understand the four stages **(Table 1, next page)**.

Table 1

Four-stage maturity model of agent identity management

Stage	Where you are	Your next step
1. Ad hoc/emerging (n = 3)	Agents run on borrowed or embedded identities; often early or small; you cannot see what is running.	Start with discovery and inventory; stand up the registry; consider managed services for capacity.
2. Extending constructs (n = 2)	Agents run as service accounts, a known stopgap; strategy is forming; you are wary of lock-in.	Separate agents from service accounts; build an agent lifecycle; protect portability.
3. Formalizing a model (n = 5)	You are building the governed class; the most demanding stage; weighing extend versus build.	Automate least-privilege and just-in-time access; add risk scoring; make the architecture decision deliberately.
4. Distinct governed class (n = 3)	The distinct class is in place; focus shifts to runtime; you are beginning to eye multi-agent chains.	Adaptive access and drift monitoring; third-party controls; agent-to-agent readiness.

n = 13 (senior security and identity leaders); Source: IDC *qualitative study Converging Identity and AI*, 2026

What organizational leaders should expect from vendors: Stage-aware engagement, not a one-size pitch, and a path that lets you act now and mature on your own terms. A vendor who cannot tell you which stage their offering serves is asking you to do that mapping yourself.

Take immediate action with a definition-agnostic operating model

There's a repeatable, ready-to-deploy governance model that matters more than anything else. Organizations at every maturity level are converging on the same operational practices: similar procedures, similar controls, similar end-states. The path is clear enough that organizations can act on it today.

The operating model you adopt should treat each AI agent as a governed object. It should have a designated owner and a defined lifecycle (creation, change, retirement). All controls, including access, audit, compliance, should attach to that agent object, instead of attaching to whatever user or service identity happened to invoke it. In other words, govern the agent, not just the identity behind it.

The central design principle: Govern the agent itself. Assign it an owner, manage it from creation to decommission, and hang all your security/compliance controls on the agent, not on the user account, service account, or identity that happens to be running it.

The organizing element is a thin abstraction layer, an agent registry, that sits above identity providers and platforms, and holds the policy (Table 2, page 42). In outline:

- Each agent has a single record, sub-agents included, with parent and delegation lineage captured where one agent invokes another
- Creation passes a mandatory intake gate
- Ownership binds to a named human sponsor

- Privilege is bound to the grant rather than the primitive, with a preference for just-in-time and zero standing access
- The lifecycle runs to a clean revocation
- An observability and audit layer records what the agent did
- Containment and an agent-aware incident-response path allow rapid isolation
- Oversight is risk-tiered to the action rather than the identity
- Agents are disclosed as non-human where they interact

The model is definition-agnostic because the controls do not change across the contested positions; only the mapping beneath them does.

- Where agents are treated as a distinct identity class, the registry becomes that class's system of record
- Where they are managed as service accounts, it tags and governs those accounts and flags where the model strains
- Where they still run on borrowed user credentials, it tracks and constrains the agent-to-human binding and marks it as the stopgap to migrate

The model is also stage-portable.

A stage-one organization **(Table 1, page 36)** can stand up a thin version, namely a registry, a sponsor, an expiration, and a revocation path, and thicken it with just-in-time privilege, drift monitoring, and eventually agent-to-agent delegation as it matures, without re-platforming. The policy content, what an approved purpose or a scoped grant means, is where the standards attach, so the same model carries the compliance story.

The mapping is direct:

- The registry satisfies the AI system inventory that the NIST AI Risk Management Framework expects
- The lifecycle-to-revocation component answers the improper offboarding risk at the top of the OWASP Non-Human Identities Top 10
- The observability and audit layer produces the lifetime event records that EU AI Act Article 12 requires of high-risk systems from December 2, 2027

An organization that stands up the registry is not just adopting a best practice; it is building the evidence base its regulators and auditors will ask for.

One clarification matters here, because "wait for the standards to settle" is the most tempting wrong conclusion to draw from this paper. Two different things get bundled together under "agent identity," and only one of them can afford to wait:

1. **Knowing who each agent is.** This means a record for every agent: it exists, a named person owns it, it has an approved purpose, it is allowed to touch these systems and this data, and it expires on a set date unless renewed. This is what the registry provides. It requires no new platform and no architecture decision, and it can start now.

2. **How each agent proves who it is.** This is the plumbing question: whether agents get their own account type inside the identity provider, or credentials built on an emerging standard such as SPIFFE and SPIRE. Here the tools and standards genuinely are still settling, and waiting for them to converge is reasonable.

Treating these as one decision is how organizations talk themselves into doing nothing, and doing nothing has a cost that compounds. An agent running on a borrowed login or a shared service account cannot be told apart from everything else using that credential. When it misbehaves, no one can say with certainty which actor acted. When it needs to be shut off. Revoking the shared credential breaks everything else that depends on it. Its permissions cannot be trimmed without trimming its neighbors'. And it leaves no per-agent record of what it did, which is precisely the record that auditors and new regulation ask for. Meanwhile the agent population keeps growing, so every month of delay adds more unlabeled identities to a pile someone will eventually sort out by hand.

The registry, then, is not a way to postpone the real decision. It delivers the benefits of distinct agent identity now, ownership, scoping, clean revocation, and an audit trail, while leaving the plumbing choice open. Regardless of which standard ultimately applies, what changes is the connection underneath the registry, not the governance built on top of it.

Blast radius, and why machine speed makes it the variable to govern.

An agent's blast radius is the full extent of what it can reach and affect if it is compromised, misdirected, or simply wrong, as well as the systems it can touch, the data it can read or alter, the actions it can complete, and the other agents and tools it can invoke, all before anyone intervenes. Two properties make blast radius the governing variable for agentic risk.

First, blast radius is set by identity and privilege, not by the agent's intent or the quality of its code. What an agent is permitted to do defines what it can do at its worst. Its radius is the union of its standing privileges, its inherited or cached credentials, and its delegation reach, and every borrowed credential or unmanaged sub-agent extends that radius silently.

Second, machine speed converts blast radius from potential into realized damage faster than human oversight can operate. A human who holds excessive privilege misuses it one action at a time, at a pace that alerts, reviews, and colleagues can catch. An agent can complete hundreds of actions in the time it takes to read this sentence. Machine speed does not change what an agent may do; it changes how much of what it may do is already done before anyone can intervene. And because the speed cannot be governed down without forfeiting the value of the agent, the radius is the variable that remains governable.

This is why the operating model works on radius at every layer. Privilege bound to the grant trims what any single task can reach. Zero standing privilege collapses the radius to near zero between tasks. Risk-tiered autonomy places human approval exactly where the radius is large or the action irreversible, while letting small-radius work run at machine speed. Delegation lineage keeps the radius of a multi-agent chain from multiplying invisibly hop by hop. And the containment path caps how long any radius stays exposed. An organization that cannot state an agent's blast radius cannot risk-tier its autonomy; the registry's record of permitted systems and data is where that statement lives.

Voice of the customer

"The blast radius concept is not well understood."

Security Architect, Enterprise IT

"The speed of AI makes this harder because new capabilities appear faster than policy and governance processes can adapt."

Security Architect, Financial Services

"Use risk-ranked autonomy so low-risk actions run at machine speed while only high-risk, irreversible or uncertain decisions undergo human approval."

C-Suite, Education and Care Services

"Responding to agents should be an order of magnitude faster than responding to humans. We also have agents monitoring agents along with humans in the loop."

Security Architect, Technology

"Absolutely not[.] AI agents should never be allowed to present as human."

Security Leader, PE-Backed Enterprise

"It is a constant uphill battle to ensure the risk posture of AI is aligned with our risk tolerance."

Security Architect, Highly Regulated Enterprise

"AI can identify and recommend much faster than human review, but it should never have material access without a human confirming and approving before the changes are committed."

C-Suite, Financial Services

"Revocation means removing API keys, tokens, roles, secrets, and delegated permissions, and confirming the agent cannot act through another identity."

Security Architect, Financial Services

Table 2

A definition-agnostic operating model for agent identity

Component	Function	Why it is definition-agnostic
Agent registry (system of record)	One record per agent: sponsor, technical owner, business purpose, permitted data and systems, environment, model and version, creation and expiration, status, and parent or delegation lineage for prime and sub-agent relationships	Describes the agent as an entity irrespective of credential type
Mandatory intake gate	No production without a named sponsor, approved purpose, scoped access, and an expiration	The gate is the same whether a distinct identity or a service account is provisioned
Ownership binding	A named human sponsor accountable for behavior and continued purpose	An organizational mapping, not a credential feature
Privilege bound to the grant	Least privilege, scoped, just-in-time, zero standing between tasks	Constrains what the agent may do; the underlying mechanism can vary by platform
Lifecycle to revocation	Time-bound, auto-expire unless recertified, clean decommission of tokens, keys, and roles	The same process applies across constructs
Observability and audit	Log actions, data touched, system, and time; behavioral and drift monitoring	The accountability layer holds whatever the identity turns out to be
Containment and incident response	Fast revoke, isolate, and kill path; agent-aware playbook	Works regardless of how the agent authenticates
Risk-tiered autonomy	Gate only high-blast-radius or irreversible actions; log the rest	Oversight attaches to the action's risk, not the identity type
Disclosure	Agents identified as non-human where they interact	Independent of construct

Source: IDC *qualitative study Converging Identity and AI*, 2026

What to look for when buying solutions

"I will not buy on a promise, and my team cannot keep up"

The need: Provable, no-lock-in solutions backed by services.

Proof-of-concept and proof-of-value are prerequisites for the buyers in this study. They stress-test candidate tools against real identity workflows and edge cases before committing. The stress test deserves a resilience dimension too. Once identity is the control plane, the identity stack is critical infrastructure, and a candidate tool should be evaluated on how it fails, how fast a compromised agent can be contained, and how the program recovers, not only on how the tool performs when everything works.

Lock-in is a real fear because identity tools embed in workflows fast, so buyers want portability, open integration, and a clear exit.

- Portability is usually framed as a procurement concern; it is equally a resilience concern
- A registry and governance layer that can be exported, re-pointed, and rebuilt on new plumbing survives a vendor failure, an acquisition, or a platform change
- A single point of failure with a contract attached is less than ideal

- And capacity matters: capacity-strained and smaller teams need managed services and upskilling to operate what they buy, because a control no one can run is a control that fails quietly

The broader buyer base behaves the same way, and it is already voting for resilience. Identity resilience or recovery is the second most selected identity security priority for the next 12 to 24 months, named by 52.7% of organizations, yet fewer than half (47.4%) have a recovery plan for an IAM system failure or identity compromise actually in place; 37.6% have one in progress and 14.0% are still developing one (IDC *Worldwide IAM Security Survey*, May 2026, n = 860). Among organizations running multiple IAM and identity security solutions, the leading reasons are that tools were selected by different teams with different priorities (26.7%), deployment model differences such as legacy on-premises environments (24.1%), and that no single vendor provides sufficient depth across all required identity security domains (24.0%) (n = 821), a finding that independently corroborates the panel's "no single vendor covers the framework" complaint.

Figure 8

Recovery plans for IAM system failure or identity compromise

N3. Does your organization have a plan to recover from an IAM system failure or identity compromise?

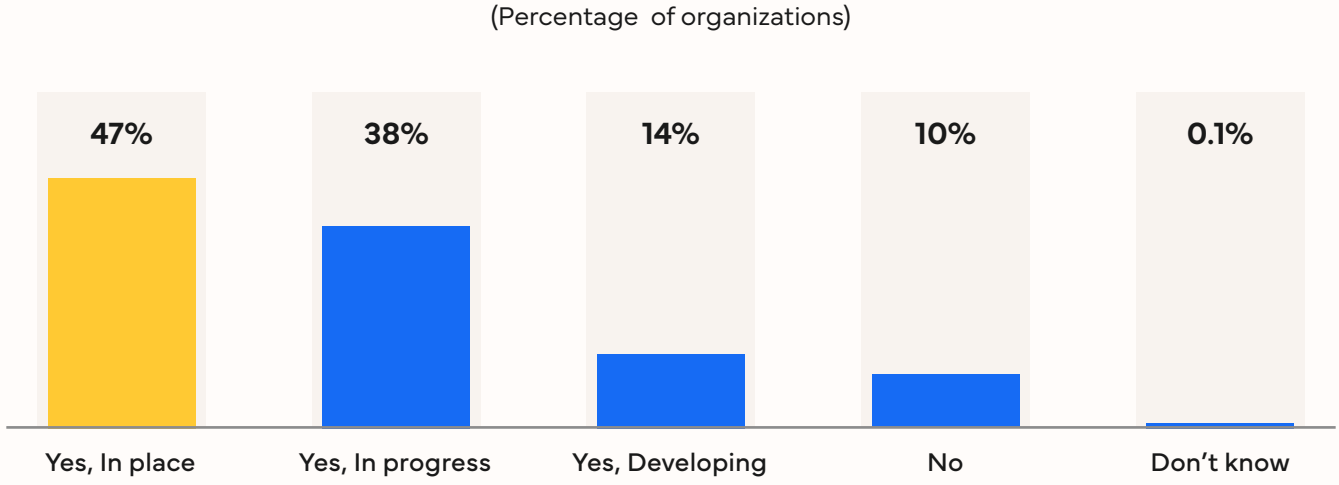

n = 860 (all organizations); Source: IDC *Worldwide IAM Security Survey*, May 2026 (EU25120009-S), unweighted

Voice of the customer

"Proof-of-concept and proof-of-value are critical before purchase decisions are finalized."

C-Suite, Education and Care Services

"As a small team, the pace of change exceeds our ability to absorb and prepare for it."

Security Leader, PE-Backed Enterprise

"For vendors, we look for evidence that their models are being tested against current adversarial techniques, not just known deepfake samples."

Security Architect, Financial Services

"Greatest identity decision is the procurement and commissioning of a new kind of IAM tool that will allow authentication, authorization, just-in-time and just-enough-authorization."

C-Suite, Education and Care Services

"Vendor lock-in is a big concern because identity tools become embedded in workflows quickly."

Security Architect, Financial Services

"To make the case for AI-augmented identity, we must demonstrate the technology at a small scale with a clear value proposition/ROI while demonstrating clear security controls at a reasonable and forecastable price."

IAM Leader, Hospitality

"We can never secure our organization 100%, but we can implement the necessary risk management controls to prevent catastrophe if or when bad things happen."

Security Leader, Entertainment and Media

What organizational leaders should be asking for: Pilotable, measurable, portable solutions, and managed services plus enablement if your team is capacity-strained. Structure the evaluation as a weighted scoring matrix built from your own frameworks and requirements rather than the vendor's demo script. Then ask every candidate vendor three questions: What can you prove in a pilot against my real workflows? What does exit look like? And who runs this if my team cannot?

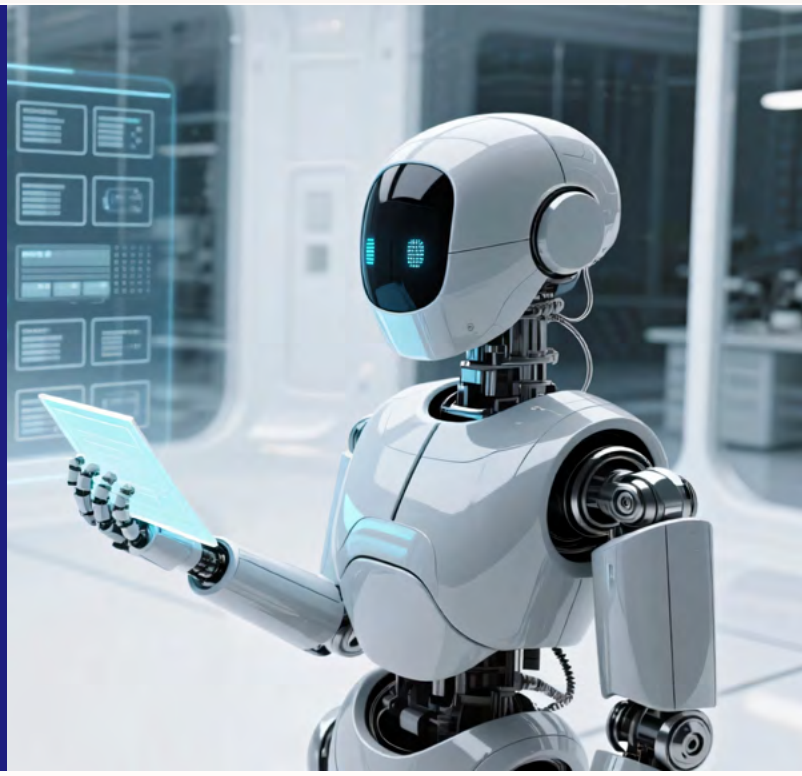

The four critical demand areas

Four capability areas reached the Critical tier in this study, combining near-universal reach with the highest evidence volume **(Table 3, below)**. Together they define the market's table stakes: not a single feature but a cluster.

Table 3

The four critical capability areas

Four Critical Capability Areas	Tier	Reach	Coded Quotes
Governance, risk, and policy framework for agentic AI	Critical	13/13	174
AI-driven threat detection and identity threat response	Critical	12/13	164
Agent and non-human identity lifecycle and ownership	Critical	13/13	157
Privilege right-sizing and just-in-time access	Critical	12/13	156

n = 13; Source: IDC *qualitative study Converging Identity and AI*, 2026

The full picture is eighteen capability areas ranked by demand strength and tiered Critical, High, Moderate, and Selective **(Table 4, below)**. Demand is a cluster, not a single feature: organizations want agents and non-human identities treated as first-class identities that can be discovered, governed across a lifecycle, privilege-bounded, and monitored, under a governance and liability framework most concede they lack.

Table 4

The demand map: All 18 capability areas, ranked by demand strength

Capability area	Tier	Reach	Quotes	What buyers should ask vendors for
Governance, risk, and policy framework for agentic AI	**Critical**	13/13	174	Provide governance frameworks, risk and blast-radius management, regulatory mapping, and clarity on agent liability.
AI-driven threat detection and identity threat response (ITDR)	**Critical**	12/13	164	Detect anomalous identity and agent behavior, drive risk-based access, and shorten detection-to-response.
Agent and NHI identity lifecycle and ownership	**Critical**	13/13	157	Model agents as a first-class identity type and govern them from creation through expiry and revocation, each with a named owner.
Privilege right-sizing and just-in-time access (PAM for AI)	**Critical**	12/13	156	Right-size agent privilege automatically, eliminate standing access with ephemeral and JIT credentials, and manage secrets at machine scale.
Agent and NHI discovery and visibility	**High**	13/13	117	Automatically find every agent and non-human identity, maintain an authoritative real-time inventory, and show what each can access.
Workforce enablement and managed services	**High**	12/13	113	Help build new identity-AI skills and offset capacity strain through automation and managed services.
Procurement support: vendor trust, transparency, lock-in	**High**	12/13	108	Model and training-data transparency, pilotable proof of value, references, and freedom from lock-in.
Shadow AI detection and governance	**High**	12/13	107	Discover unsanctioned AI tools and agents, then block or bring them into governance.

Capability area	Tier	Reach	Quotes	What buyers should ask vendors for
Identity platform consolidation and orchestration	**High**	11/13	106	Consolidate identity onto an integrated fabric or orchestration layer and fill the gaps point tools leave.
Human-in-the-loop and tiered autonomy controls	**High**	12/13	101	Configurable human approval and risk-ranked autonomy so low-risk work runs at machine speed while high-risk actions are gated.
Identity governance automation and entitlement cleanup	**High**	12/13	100	Use AI to map and clean up entitlements and automate provisioning and access certification, with accuracy buyers can trust.
Continuous and adaptive authentication	**Moderate**	12/13	82	Move beyond point-in-time MFA to continuous, behavioral, low-friction authentication.
Deepfake and synthetic-identity defense/identity proofing	**Moderate**	12/13	71	Detect deepfakes and synthetic identities and harden identity proofing against AI-augmented attacks.
Accountability, audit logging, and agent disclosure	**Moderate**	12/13	69	Immutable, explainable audit trails of agent actions, and disclosure that an actor is non-human.
Third-party and vendor AI identity risk	**Moderate**	12/13	62	Visibility and control over identities introduced by vendor AI, backed by due diligence and contractual controls.
Zero trust enablement	**Selective**	12/13	43	Shift from static rules to continuous, context-aware access decisions and help clear legacy blockers.
Incident response and containment for agents	**Selective**	11/13	37	Revoke, isolate, and kill compromised agents fast, with agent-aware IR playbooks.
Post-deployment identity monitoring	**Selective**	11/13	34	Continuous identity monitoring and drift detection after an agent goes live.

Note: Demand strength combines breadth with intensity (coded-quote volume). Source: IDC *qualitative study Converging Identity and AI*, 2026; n = 13

Recommendations

What good looks like

For your organization (act now, mature deliberately):

- Stand up an AI governance board that owns the standards for access, usage, and approved AI, and route agent architecture through its approval before production. Governance bodies of this kind were named by 7 of 13 respondents, and a single accountable body resolves the contested-ownership problem this study documents.
- Establish a distinct logical identity for every agent now via a registry: sponsor, business purpose, scoped access, expiry, revocation. No architecture decision is required to start.
- Stand up automated discovery and an authoritative inventory before anything else. You cannot govern what you cannot see.
- Push the logical identity now; stage the technical primitive as standards converge.
- Locate your stage on the maturity curve and take the next step it indicates, rather than attempting Stage 4 controls from a Stage 1 posture.

What to demand from your vendors (key capabilities in NHI management):

- Immediate visibility through discovery assessments, followed by the sponsor-and-lifecycle operating model delivered as a turnkey pattern mapped to the standards.

- Immediate support and solutions for the four Critical areas as the cross-stage floor: governance, risk, and policy framework for agentic AI; AI-driven threat detection and identity threat response; agent and non-human identity lifecycle and ownership; and privilege right-sizing with just-in-time access.
- Proven value through staged pilots and peer benchmarks; protected portability; and engagement matched to your organization's AI maturity level.

The bottom line

Security leaders are not asking vendors to define agent identity for them. They are asking for help seeing what they already have, governing it with tools that fit, and proving value without locking them in.

None of this is exotic. A named owner, an accurate inventory, scoped access, and a clean revocation path are security fundamentals, back to basics, applied to a new class of identity that moves at machine speed.

Organizations converge on procedural governance, a sponsor and a lifecycle for every agent, long before they converge on the architecture of what an agent identity is. Act on the consensus now. The destination is shared; the starting point is not.

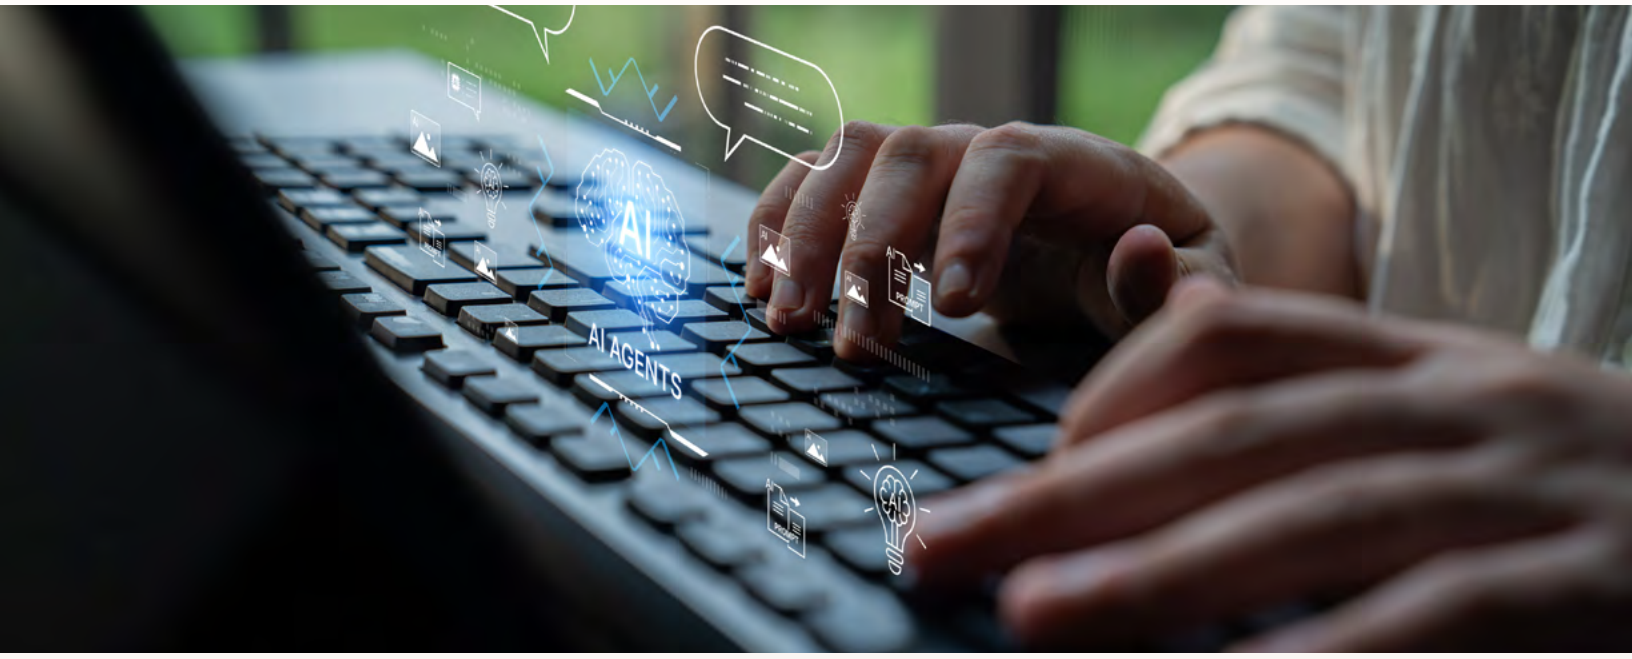

Methodology

This IDC White Paper draws on the IDC qualitative study Converging Identity and AI: a three-session, asynchronous expert panel of 13 senior security and identity leaders (12 full participants, 1 partial) spanning financial services, healthcare, education and care services, technology and media, entertainment and media, hospitality, private-equity-backed enterprise, highly regulated enterprise, and enterprise IT. Participants answered 31 open-ended items in their own words across three consecutive daily sessions.

Responses were segmented into 1,810 sentence-level units and inductively coded, producing 2,007 code assignments across 21 themes, 100 granular concepts, and 168 sub-codes. From the coded corpus, IDC derived a demand map of 18 vendor capability areas comprising 49 discrete needs, rated for demand strength by combining breadth (respondent reach) with intensity (coded-quote volume), and a four-stage maturity model in which each stage is populated by at least two respondents. Every claim traces to a source quote. Quotes are lightly copyedited for spelling, capitalization, and punctuation only; wording is unchanged, and any inserted word is marked in brackets.

Panel findings are corroborated where noted with quantitative results from three IDC surveys: the IDC *Worldwide IAM Security Survey (May 2026*, n = 860 identity and security decision makers), the IDC *Unified AI Platforms and Governance Survey (March 2026*, n = 744), and the IDC *Agentic AI Functional Use Case Study (February 2026*, n = 1,967, weighted).

About the IDC analyst

Grace Trinidad

Research Director,
AI Security and Trust, IDC

Grace Trinidad is research director and global lead for IDC's AI Security and Trust program, where she leads research on AI security and how trust shapes adoption, procurement, and brand value across the digital economy. Her coverage spans the trustworthiness of AI and generative AI, security, privacy, and compliance, with an emphasis on how organizations operationalize trust as AI enters core enterprise decision-making. Trinidad helped create IDC's Trust Perception Index, now used to evaluate trust across leading cloud, database, and enterprise PC vendors.

More about Grace Trinidad

Message from the sponsor

GUIDEPOINT® SECURITY

GuidePoint Security offers this research as a deep dive at the intersection of artificial intelligence and Identity and Access Management. Organizations can use the information as a foundation for more confident AI adoption, and to secure innovation, accelerate cybersecurity operations, and defend against sophisticated AI-powered threats.

GuidePoint Security is the trusted advisor and partner that organizations rely on to protect, enable, and lead their cybersecurity efforts. We help companies of every size solve their most complex cybersecurity challenges, mature security architectures, and proactively reduce enterprise risk. Our purpose-built solutions ensure compliance, safeguard critical assets, and align security with business objectives. Backed by deep expertise, strong partnerships, and advanced technologies, we deliver scalable, adaptive cybersecurity capabilities built for today's threats and ready to evolve with the threat landscape, so our customers operate securely, confidently, and without compromise. Stronger Together. Protecting What's Next.

Learn more

IDC Custom Solutions

IDC Custom Solutions produced this publication. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis that IDC independently conducted and published, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies.

This IDC material is licensed for external use and in no way does the use or publication of IDC research indicate IDC's endorsement of the sponsor's or licensee's products or strategies.

IDC.com | IDC on LinkedIn | IDC Blog

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries. IDC's analysis and insight help IT professionals, business executives, and the investment community make fact-based technology decisions and achieve their key business objectives.