Survey Report

Agents of Change

The Confidence Gap in Managing Non-human Identities

opentext™

DARKREADING

Seven in ten executives say their identity and access management (IAM) services are mature. Yet roughly half still can't apply different policies to non-human identities (NHIs). That's problematic today, but as entities such as microservices and AI agents gain traction, it'll become a crisis.

Conducted by Cybersecurity Dive on behalf of OpenText, this survey of 152 corporate decision-makers reveals a widening gap between what organizations believe they know about securing autonomous AI agents and what they actually know. Respondents are caught in an awkward in-between stage. Agentic AI adoption remains modest, microservice usage beyond basic data services is still limited, and over-privileged access tops the list of concerns. Yet confidence remains high.

The tension between confidence and capability runs through every section of this report. It also highlights a divide between leaders and laggards. Some organizations genuinely have their house in order, while others lag behind the curve and may not realize it.

What do we mean by agentic AI?

The most basic problem is one of definition. The consensus is split on what agentic AI means. Half define it as software that perceives its environment, makes decisions, and takes actions autonomously. A third (34%) think of it more narrowly — as microservices that encapsulate decision-making logic and coordinate with other services to achieve defined goals. About one in seven (14%) equate it with adaptive bots or scripts. A tiny minority (2%) think it means any AI system that uses large language models.

That fragmentation matters more than it might seem. Organizations can't secure what they can't articulate. The gap between industry buzz and practitioner understanding has implications for every IAM decision downstream.

We treat agentic AI as the next evolutionary step in microservices, powered by LLMs. Traditional microservices do what they're programmed to do. Their behavior changes in predefined ways in response to inputs. Agentic microservices, by contrast, are adaptive and self-directive, often working together without human instruction. They operate non-deterministically, producing outcomes we can't always predict or explain. That raises the stakes considerably for identity management.

50%
define agentic AI as software that perceives its environment, makes decisions, and takes actions autonomously.

34%
think of it as microservices that encapsulate decision-making logic and coordinate with other services to achieve defined goals.

opentext™

The current footprint of non-human identities

The current footprint of agentic AI in corporate environments remains small. In 92% of companies, fewer than one in four NHIs is agentic. More than half (54%) of respondents say that fewer than one in ten NHIs is agentic.

Even broader microservice adoption is more modest than industry narratives suggest. Only data services enjoy “high use” (defined as 30% or more penetration) among a majority of companies. Gateway services reach that threshold for 39% of organizations, followed by 34% for domain services. Everything else (event processing, utility, caching, proxy, and aggregator services) falls well below those numbers. Yet the shift from monolithic infrastructure to microservices has been discussed for well over a decade. Again, reality doesn’t reflect rhetoric.

In the microservice environments that exist, AI agents account for a limited share. For most (67%), the share of microservices incorporating AI agents is under 20%.

In 92% of companies, fewer than one in four NHIs is agentic.

opentext™

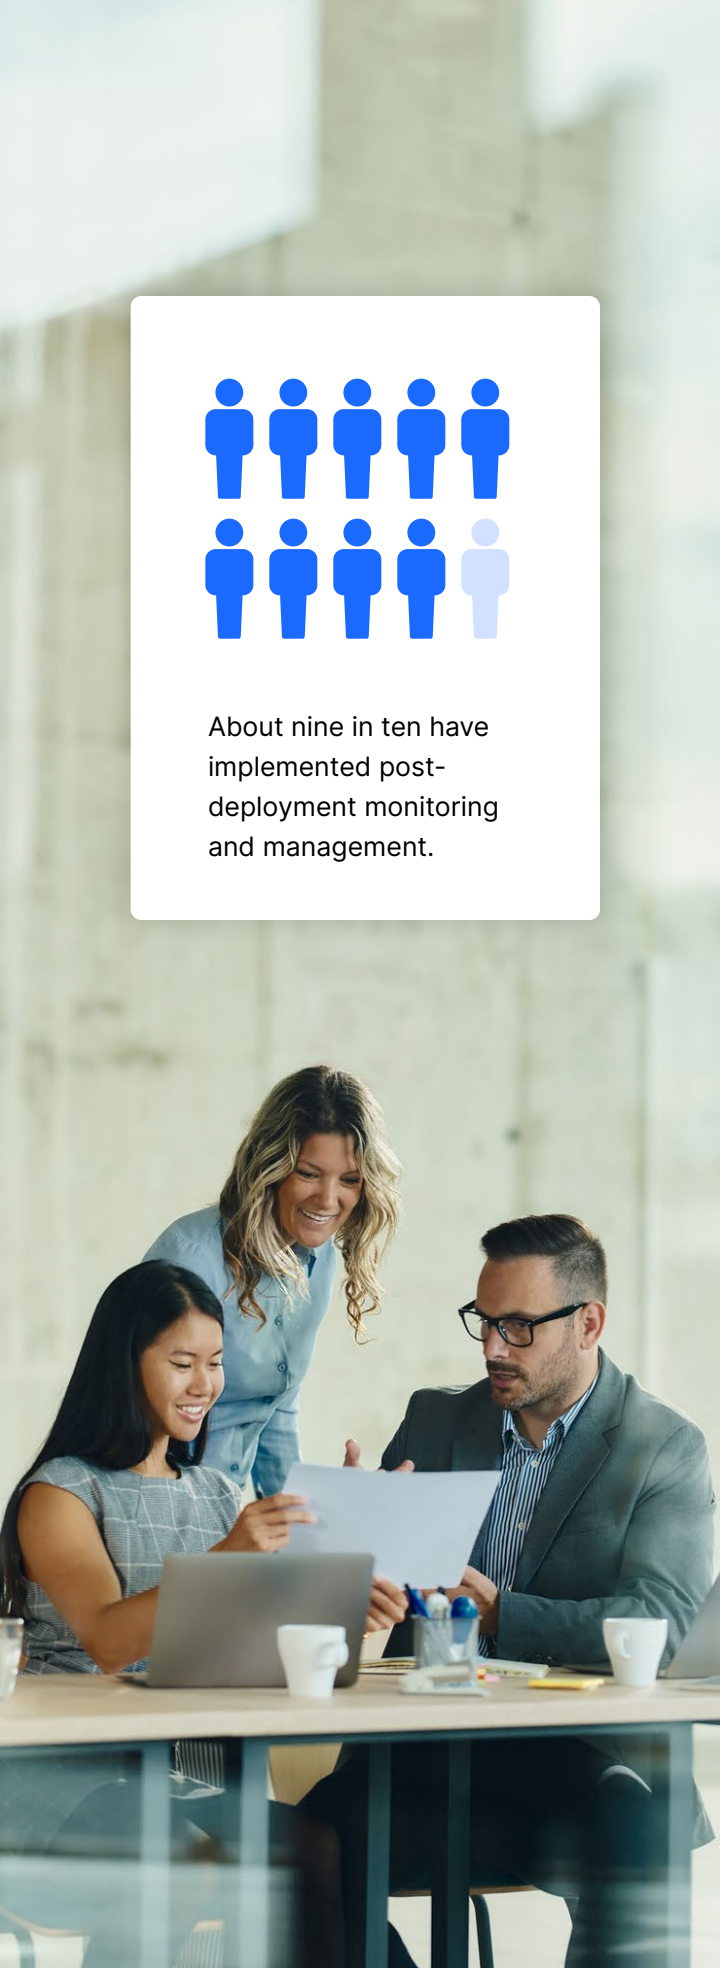

About nine in ten have implemented post-deployment monitoring and management.

How organizations onboard AI agents

Half of organizations that introduce new AI agents rely on a formal security review and approval process. However, a third (35%) say the process varies by team or use case, indicating inconsistent governance across business units. Only 9% run new agents through full identity management onboarding, while 6% rely on informal IT or security approval.

On paper, the broader security review infrastructure looks solid. About nine in ten have implemented post-deployment monitoring and management (89%), design-phase review (88%), and pre-deployment gates (87%) for microservices across their environments. Coding-phase review trails slightly at 84%.

But depth matters as much as breadth. Few enterprises apply these practices across all services. Only 36% of organizations apply the most pervasive practice, security monitoring. The gap between documented process and consistent execution is where risk accumulates, especially when introducing more autonomous and less predictable workloads.

IAM maturity: self-assessment vs. reality

Almost one in five respondents (18%) say their IAM is very mature (fully implemented and automated), while 52% are at the mature stage (implemented but partially automated). A quarter (25%) acknowledge they're still in the development stage, while a candid 5% admit they're even earlier in the process.

Yet executives are less confident in their ability to identify and apply differentiated policies for NHIs. Only 14% rate themselves very mature on this front, and just 39% consider themselves mature. Meanwhile, four in ten (40%) are still developing, and 6% are just starting out. That gap between confidence and capability becomes dangerous as organizations scale autonomous workloads.

We can see the gap between perception and reality in deployment statistics. **Sixty-three percent of organizations have deployed identity governance, and 61% have deployed data governance and secure API access.** The latter is especially relevant because APIs are the primary interface between services and agents. With more than a third of companies still lacking these controls, there is much work to do.

18%
say their IAM is very mature (fully implemented and automated).

52%
are at the mature stage (implemented but partially automated).

opentext™

Organizations struggle with consistent credential and access management

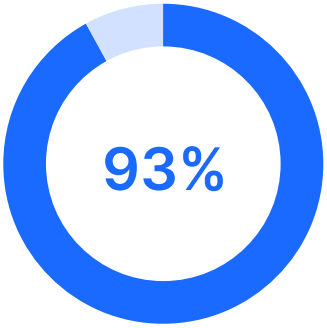

93% of organizations have assigned distinct identities rather than shared credentials to at least some AI agents.

Organizations provision credentials for non-human identities in various ways. **Short-lived or auto-rotating credentials lead the field at 82%. Scoped or time-limited tokens follow at 74%.** Both figures are encouraging, yet long-lived credentials, such as static keys, passwords, and tokens, persist in 65% of organizations, sitting alongside the newer approaches. The average organization juggles 3.2 credential methods simultaneously, underscoring the complexity of real-world environments and the difficulty of fully retiring older practices.

Unique identity usage is also patchy. While 93% of organizations have assigned distinct identities rather than shared credentials to at least some AI agents, only 27% have done so for all or most agents. Two-thirds have limited distinct identity assignment to a few pilots or experiments.

That uneven discipline extends to access reviews. **Only 13% review continuously, 23% monthly, and 43% quarterly.** At the other end, 11% review only when triggered by an incident. In an environment where autonomous agents can silently accumulate privileges, that lag creates real exposure.

A parallel gap in human user authentication

Consistency challenges aren't limited to non-human identities. **Two-thirds of organizations have deployed passwordless authentication fully or partially for human users.** That leaves a third still piloting, planning, or not considering it at all. Passwordless authentication is an excellent defense against token harvesting, which is becoming a common attack vector against multi-factor authentication (MFA). You cannot proxy a credential that is cryptographically bound to a physical device.

The ShinyHunters vishing campaigns that escalated in January and February 2026 illustrate why passwordless matters. Attackers posing as corporate IT staff used phone calls to direct employees to fake login portals, intercepting SSO credentials and MFA codes in real time via adversary-in-the-middle proxy kits. FIDO2 and passkey-based authentication would have rendered the attack chain ineffective.[1]

Perhaps the starkest gap in the survey data is the disparity between protections for internal and external users. Eighty-nine percent of organizations have deployed MFA for employees, but only 20% have extended it to external users. Single sign-on is at 80% for employees and 20% for external users. **Automated provisioning and deprovisioning also shows a large delta between internal and external, at 52% and 15% respectively.** The protective perimeter has a clear boundary, and everything outside it is considerably less secure.

89%
of organizations have deployed MFA for employees.

20%
have extended it to external users.

1. "Tracking the Expansion of ShinyHunters-Branded SaaS Data Theft | Google Cloud Blog." Google Cloud Threat Intelligence, cloud.google.com/blog/topics/threat-intelligence/expansion-shinyhunters-saas-data-theft

The most straightforward safeguard, human approval for sensitive actions, leads the field, with 86% of organizations using or piloting it.

Security controls in practice

The simpler the control for NHIs, the broader the adoption. The most straightforward safeguard, human approval for sensitive actions, leads the field, with 86% of organizations using or piloting it. Automatic credential rotation and regular identity health checks or attestation follow at 77% and 73%, respectively. Isolated execution environments or sandboxes sit at 70%, and real-time access revocation at 60%.

But as controls become more sophisticated, adoption drops sharply. **Only 48% are using or piloting behavioral anomaly detection for agents (the kind of monitoring that could catch one acting outside its intended scope).** Decentralized identity approaches and just-in-time access provisioning each account for roughly 30%. Non-human identity governance will eventually require these more advanced tools.

Meanwhile, confidence in containing digital minions on the loose is higher than the controls would justify. Three in ten executives say they're very confident their systems can contain an agent acting outside its intended boundaries, while four in ten are moderately confident. Another 30% are only slightly confident.

One factor almost certainly feeding that confidence is that 95% of organizations have never experienced a security incident or near-miss caused by a non-human identity or AI agent. The two percent that have seen an incident cited misconfigured identity or access controls. The absence of incidents doesn't equal the absence of risk. It may simply reflect the still-small footprint of autonomous agents.

The concern executives voice most often reinforces the gap: over-privileged access or privilege escalation tops the worry list at 51%, followed by difficulty tracking or auditing agent actions (36%) and uncontrolled identity sprawl (25%).

Where organizations stand on zero trust

This is the fourth time OpenText has asked a zero-trust tracking question in its research. Because populations and methodologies have varied over the years, we should treat differences in numbers as directional rather than precise. With that in mind, several trends stand out.

- **Employee MFA saw the biggest absolute jump: 61% (2022) to 63% (2024) and 89% (2026).**

- While it could be doing better, passwordless access has finally gained meaningful traction, rising from 12% in 2022 to 17% in 2024 and 41% in 2026. Last year's report noted that low adoption "reflects the lack of organizational confidence we find in passwordless tech." That gap, at least, appears to be closing.

- **Identity governance climbed steadily: 36% to 47% to 63%. Privileged access to cloud infrastructure: 33% to 50% to 55%.**

- Context-based access policies remain the persistent laggard, inching the needle from 21% to 26% and then 32%. Yet context-aware, intent-based access control is exactly the capability that executives say they most want from future IAM solutions (see "What an ideal IAM solution would look like"). And yet it remains the least-deployed zero-trust control in practice.

The road ahead: investment, barriers, and risk perception

Executives expect mostly moderate growth in AI agents over the next 12 to 24 months. **Half envision moderate expansion — between 2x and 5x current levels. Four in ten expect only modest growth of 1x to 2x. Just nine percent anticipate significant growth of 5x to 10x, and notably, no one predicts either rapid growth above 10x or flat growth.**

Budget and staffing plans align with measured expectations. Seventy-two percent anticipate a modest increase of 10% to 24% in resources dedicated to managing AI agents and non-human identities over the next 12 months. One in five expects a moderate increase of 25% to 49%. Only 2% envision a significant increase of 50% or more, and 5% plan no increase at all.

Nearly all executives expect they'll need to change their IAM approaches to support AI agents at scale, but 59% think they can get by with only moderate adjustments to existing frameworks. Just over a third (36%) eye major changes, including three percent who foresee a complete paradigm shift.

The top barriers to making their IAM AI-ready are the complexity of managing short-lived or dynamic identities, which concerns 49% of respondents. Close behind, 41% say current tools and platforms don't support the level of governance required. Regulatory and compliance uncertainty concerns 35%. Organizational challenges also loom, with 27% fretting over unclear ownership of agent identities.

When asked which NHIs pose the greatest risk, the answer is overwhelmingly agentic, at 77%, far above any other category. Other automation or script accounts follow at 41%, and supply-chain or B2B services at 32%.

Perhaps the most revealing data point comes from the self-assessed readiness questions. Eight in ten (82%) agree they "understand the potential security risks that AI agents pose." Yet only 53% agree that their IAM tools are "equipped to handle the expected growth in AI agents." That 29-point gap between awareness and preparedness captures the entire confidence problem.

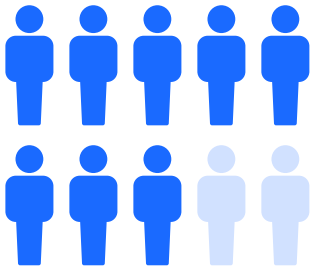

Eight in ten (82%) agree they "understand the potential security risks that AI agents pose."

opentext™

What an ideal IAM solution would look like

OpenText asked executives what capabilities an ideal IAM solution for AI agents would offer that no vendor currently delivers well. They gave a wide range of answers that clustered into three broad categories.

The largest group focused on monitoring and detection capabilities, including real-time anomaly identification, transparent causal logs, and continuous behavior tracking. A second cluster described controls and prevention capabilities, including instant access revocation, intent-based guardrails, and automated policy enforcement. The third centered on efficiency: seamless cross-environment interoperability, scalability as agent counts grow, and workflow automation.

Some responses captured the wish list vividly:

> "Security teams can instantly ditch agent access with one dashboard."

> "It should detect strange or risky behavior in real time."

> "Having dynamic scoping and intent-based guardrails with a human in the loop triggers."

> "It should be contextual and task-based, not tied to rigid role definitions."

These responses collectively portray how far current tooling falls short of what practitioners need. They also highlight the kind of context-aware, dynamic governance that a world of autonomous agents will eventually demand.

Time to start moving

The gap between confidence and capability is the central finding of this survey. Organizations believe they are ready for the age of non-human identities, but the data, considered in full, suggest otherwise.

But the slow walk toward agentic AI may be a gift. Because adoption remains modest and growth expectations are measured rather than explosive, organizations still have time to close the maturity gaps in their IAM foundations before the pressure builds in earnest.

Governance frameworks designed for a human-centric era won't stretch far enough to cover the extensive use of NHIs. The fundamentals must be in place before organizations attempt to scale. With the agentic train still approaching the station, the question is whether organizations will be ready to board in time.

opentext™

OpenText™ Identity and Access Management (IAM) offers enterprise-scale solutions that secure access, enforce compliance, and enable digital trust across hybrid and cloud environments. From identity lifecycle automation and privileged access control to adaptive authentication and governance, OpenText IAM (NetIQ) helps organizations reduce cyber risk, support Zero Trust initiatives, and meet evolving regulatory demands.

Learn more

03.26 | 248-000172-001