




by **Evan Schuman**
Contributor

Think agentic AI is hard to secure today? Just wait a few months

Feature

Feb 3, 2026 • 8 mins

Exponential expansion of autonomous agents in the enterprise may expand enterprise threat surfaces to an almost unmanageable degree — especially given poor foundations for non-human identity oversight.



Credit: Golovina Marina / Shutterstock

Early experimentation with agentic AI has given CISOs a preview of the **possible cybersecurity nightmares ahead** [<https://www.csoononline.com/article/4109999/agentic-ai-already-hinting-at-cybersecuritys-pending-identity-crisis.html>]. But with autonomous agent adoption expected to soar throughout 2026, CISOs' lack of visibility into agentic identities, activities, and decision-making is set to get far worse in quick measure.

Agentic use will vary by enterprise, but analysts, consultants, and security vendors agree that their numbers will expand far beyond CISOs' ability to maintain control as they simultaneously navigate the price of decades of **identity governance neglect for non-human identities (NHIs)** [<https://www.csoononline.com/article/3520884/personhood-cybersecuritys-next-great-authentication-battle-as-ai-improves.html>], including service accounts, OAuth tokens, embedded API keys, and automation credentials.

Ishraq Khan [<https://www.linkedin.com/in/ishraqkhann/>], CEO of coding productivity tool vendor Kodezi, sees most enterprises today housing 8 to 10 million such identities, a figure he projects will hit 20 to 50 million by year's end.

Jason Sabin [<https://www.digicert.com/blog/author/jason-sabin/>], CTO at DigiCert, predicts an even steeper rise, with enterprises' identity role calls increasing 10 times by January 2027.

"We need to rethink how identity and data provisioning is done and put in place the right processes that can scale with the growth of agentic identities," says **Justin Greis** [<https://acceligence.com/talent/profiles/justin-greis/>], CEO of consulting firm Acceligen and former head of the North American cybersecurity practice at McKinsey. "You simply cannot apply human processes to something that will scale at this rate."

Visibility is the bigger problem

As bad as that massively expanding identity universe is, the bigger problem may be how little visibility CISOs have into NHIs, with AI agents offering not just the fastest growth but the least visibility.

Jason Andersen [<https://moorinsightsstrategy.com/team/jason-andersen/>], principal analyst for Moor Insights & Strategy, estimates 25% NHI visibility for enterprise CISOs today. "The remaining 75% is in the shadows," he adds.

Those shadows include "semi-shadow" activities, such as third parties or lines of business that have been given permission to experiment with agentic AI but have not necessarily alerted IT or security teams about what they are doing.

Still, Andersen sees that number getting a lot worse, projecting visibility to drop to about 12% by year-end and then into the single digits by January 2028. “And then they’ll likely fix it,” he says, adding, “It’s a big frickin’ problem.”

Gartner analysts **Jeremy D’Hoinne** [<https://www.gartner.com/en/experts/jeremy-dhoinne>] and **Akif Khan** [<https://www.gartner.com/en/experts/akif-khan>] agree CISOs face urgent problems in this area today.

NHIs are going to be “several orders of magnitude larger than human identities and most organizations do not have a strong enough foundation to manage both machine and agentic identities,” Gartner’s Khan says.

Enterprise CISOs are “blind to what is happening. The numbers are going to be overwhelming,” D’Hoinne adds.

Forrester expects similar outcomes for CISOs. “There is going to be an explosion of non-human identities,” says Forrester analyst **Geoff Cairns** [<https://www.forrester.com/analyst-bio/geoff-cairns/BIO20052>]. “The exponential growth is indisputable.”

Kodezi’s Khan notes that the lack of a robust base for NHI governance — now including agentic AI — is a critical problem.

“Enterprises never solved non-human authentication so we don’t have the systems in place for a good secure environment. At its core, we never had the right foundation. That means that we will never have that perfect inventory,” he explains.

Cost effective fix: Do nothing

Kodezi’s Khan offers an interesting fix for that foundational problem: Don’t even try.

He argues it’s a money pit that will never be fully resolved. Instead, he suggests pouring resources into creating a strict identity strategy for every NHI going forward.

“Aim for containment rather than for perfection. You can’t really govern every identity, but if you start now, you *can* govern future actions,” he says, adding that, over the years, the percentage of uncontrolled identities will slowly drop as millions more identities are added. **Nik Kale** [<https://www.linkedin.com/in/nikkale>], principal engineer at Cisco and member of the Coalition for Secure AI (CoSAI) and ACM’s AI Security (AISec) program committee, agrees with that assessment. “If you are drowning, you don’t start by draining the ocean.”

“The ratios tell you why this is so ungovernable. These identities are growing much faster than the discovery capabilities,” Kale notes. “It becomes a math problem at that point.”

As for the path forward, Kale advises not to try to fix the legacy situation.

“You just have to contain it, segment it, assume it’s compromised and that it’s hostile territory,” he says. “The plan needs to be containment plus a clean slate going forward. Inventory all non-human identities. Identify which have standing versus just-in-time access. Assign ownership to every one of them. No product required — just a terrifying spreadsheet.”

Kale adds that cleaning IDs from now on will deliver a better benefit to CISOs. “In my opinion, the ratio matters less than the governance gap. Whether it’s 200:1 or 500:1, if IAM [identity access management] only manages 44% of them, the attack surface is already unmanageable,” he says.

But he stresses that NHIs — especially when agentic — can be particularly difficult to find, let alone control.

“Most organizations are undercounting by two to three times because machine identities are scattered across cloud consoles, repos, config files, and secrets managers that nobody’s aggregating,” Kale says. “Agentic AI is a multiplier, not an addition. Agents spawn subagents, create credentials dynamically, and establish agent-to-agent auth chains. One agent deployment can generate dozens of new machine identities.”

Sanchit Vir Gogia [<https://greyhoundresearch.com/svg/>], chief analyst at Greyhound Research, sees a reckoning ahead.

“The enterprise control plane has quietly shifted from humans to machines, while governance stayed behind,” he says. “Once nonhuman identities outnumber humans by hundreds to one, identity stops being an administrative discipline and becomes the operating system of trust. The failure mode is not that there are too many identities; it is that enterprises cannot assert intent, ownership, and accountability for what those identities are doing at runtime.”

Moreover, the situation is intensifying thanks to today’s business environment.

“This is compounded by incentive structures that reward speed and uptime while penalizing breakage, which leads teams to overpermission machines by default,” Gogia says.

“Overpermission is invisible until it is catastrophic. At that point, audits, roles, and reviews offer comfort but not control.”

Agentic didn’t start the fire

None of this situation was caused by agentic AI, Gogia underscores.

“Enterprises did not enter a machine identity crisis because of agentic AI. They entered it years ago through service accounts, embedded API keys, long lived tokens, and automation credentials that were created to keep systems moving and then quietly forgotten,” he says. “What agents change is velocity and reach. They inherit trust and then operationalize it at machine speed. A legacy identity that once represented a contained risk now becomes an execution layer across systems, vendors, and workflows.”

Gogia adds: “The most dangerous assumption in enterprise security today is that valid identity implies safe behavior. In machine-driven environments, credentials are often correct and activity is authorized, yet outcomes are harmful. Machines do not follow joiner-mover-lever models. They do not pause for approvals. They operate continuously and propagate actions automatically.”

As a result, decision-making agents, layered into operations, achieve a rate of action that “collapses the window for detection,” he says. “The failure shifts from prevention to detection lag. By the time humans understand what happened, the agent already did it.” This should — and likely will — cause a **rethinking from both enterprise CISOs and CIOs [https://www.csoonline.com/article/4089732/rethinking-identity-for-the-ai-era-cisos-must-build-trust-at-machine-speed.html]**, he says.

“This moment tests leadership alignment. CIOs are under pressure to deploy agents for productivity and scale. CISOs are staring at accountability gaps, forensic complexity, and cascading blast radius. If these agendas diverge, the enterprise ends up with autonomy without responsibility. Boards will ask who owns an agent, who sets its boundaries, and who answers when it causes harm,” Gogia explains.

“The next phase of governance will require responsibility mapping for agents, separation of duties for high impact actions, and clear human checkpoints where judgment truly matters,” he adds. “Incident response must also evolve toward reconstructing chains of machine decisions, not just tracing logins.”

Artificial Intelligence[https://www.csoonline.com/artificial-intelligence/]

Security[https://www.csoonline.com/security/]

Identity and Access Management[https://www.csoonline.com/identity-and-access-management/]

Access Control[https://www.csoonline.com/access-control/]

Sponsored Links

Join IGEL Now & Next in Miami, March 30–April 2. This is ultimate event to discover what's next in endpoint cybersecurity & innovation. Register now!
