

The State of Identity Security in the AI Era

A Study of AI's Effect on the Identity Attack Surface —and Organizations' Responses

Thank you for your interest in exploring this content.

Carahsoft is the **Trusted Government IT Solutions Provider**® supporting a broad portfolio of industry-leading technologies through NASPO, OMNIA, TIPS and a wide range of other contract vehicles.

As the **Master Government Aggregator**®, Carahsoft connects government agencies, industry partners, and technology providers to deliver innovative, mission-focused solutions.

In partnership with Semperis, we provide technology solutions that drive modernization, strengthen operations, and ensure compliance with evolving government standards.



To learn more about how Carahsoft can support your technology needs, please visit carahsoft.com



Explore More Resources:
carah.io/SemperisResources



Join Events & Webinars:
carah.io/SemperisEvents



Discover Technology Solutions:
carah.io/Semperis



Learn About Procurement:
carah.io/SemperisContracts



Connect With Our Team:
Semperis@carahsoft.com
(844) 214-4790

The State of Identity Security in the AI Era

A Study of AI's Effect on the Identity Attack Surface—and Organizations' Responses

- AI is quietly redrawing the attack boundary of the global identity fabric.
- Organizations are giving AI agents the keys to critical systems faster than they're putting guardrails around those new identities.
- Without rigorous defense and recoverability of the identity infrastructure, over-helpful agents can become a fast track to full compromise of Active Directory, Entra ID, or Okta.



“

In a world where ransomware can shut down hospitals and pipelines, wiring unguarded AI into Active Directory, Entra ID, or Okta isn't innovation—it's the fastest route from 'productivity booster' to full-scale business outage.”

Tim Brown, Former SolarWinds CISO and Team8 CISO in Residency

Executive Summary

The AI boom is flooding enterprises with non-human identities (NHIs). AI agents are enabling new ways of identifying and detecting threats. They're also quietly reshaping the global identity attack surface.

To learn how organizations are managing the intersection of AI and identity resilience, [Semperis](#) and [Censuswide](#) surveyed 1,100 IT and security professionals across eight countries about how they're deploying AI and how those deployments interact with the identity fabric.

Key takeaways



74%

of respondents believe AI will increase attacks on identity infrastructure.



32%

are very confident they could regain control if AI exposes admin credentials.



93%

already use—or plan to use—AI agents for sensitive security tasks such as password resets and VPN access.



65%

fully register, authenticate, and authorize AI identities.



92%

say AI is installed on at least some local machines, giving AI access to SSH and encryption keys.

The bottom line: Security leaders are granting elevated security privileges to AI, even as they predict an increase in AI-related risk—and a confidence gap in the recoverability of their identity systems.

Contents

AI on the Inside: Are Enterprises Ready?	4
It's 10 PM. Do you know where your AI agents are?	6
Are you ready for an AI-incited identity breach?	7
Global Statistics on AI + Identity Risk	8
AI and the Identity Attack Surface	8
AI Reach + Security	9
AI Governance	10
Methodology	12
About Semperis	12

Contributing Experts



Tim Brown
CISO in Residency, Team8;
former CISO, SolarWinds



Grace Cassy
Partner, Ten Eleven Ventures



Sarah Cecchetti
Director Product Management,
Semperis



Chris Inglis
Former U.S. National Cyber Director;
Strategic Advisor, Semperis



Stuart McClure
CEO, Wethos AI; entrepreneur



Alex Weinert
Chief Product Officer, Semperis;
former VP of Identity Security, Microsoft

AI on the Inside: Are Enterprises Ready?

The accelerated use of AI throughout global enterprises is introducing a wave of generative AI tools and a bevy of AI agents, each with its own non-human identity.

NHIs already vastly outnumber human users. Microsoft estimates roughly 10 NHIs per human identity in 2018, trending toward a 100:1 ratio as agentic and workload identities proliferate.¹

How are these new “coworkers” affecting organizations’ cyber resilience?

To find out, Semperis partnered with research firm Censurwide to conduct a study of 1,100 organizations across the US, UK, Germany, France, Italy, Spain, Australia, and Singapore.

We asked study respondents to tell us about the ways in which they are integrating AI into their identity infrastructures and the steps that they’re taking to prevent the potential security risks that AI can introduce.

An understanding of those risks was widespread:



74% of respondents believe that AI functionality will increase the frequency of attacks on the identity infrastructure—a top target for malicious actors.



Microsoft has reported that identity-based attacks—already a top tactic for malicious actors, with the company reporting **600+ million such attacks per day** in 2024²—increased by **32%** in the first half of 2025³.

Even so, most organizations admit that AI has—or soon will have—widespread access to highly sensitive parts of that infrastructure.




29%

More than a quarter of surveyed organizations already use AI agents to handle security-related help desk tickets—including password resets and VPN access.



64%

More than half intend to do so within the next year, bringing the total of organizations enabling this level of agentic AI access to 93%.



92%

Nearly all of respondents say that some percent of their workforce has AI installed on local machines where it can access SSH and encryption keys.

These behaviors exponentially increase the need for airtight identity security, especially around AI agents and identities.

¹ [Microsoft 2023 State of Cloud Permissions Risks](#)

² [Microsoft Digital Defense Report 2024](#)

³ [Microsoft Digital Defense Report 2025](#)

Each new agent, service principal, and low-code “helper” becomes another potential entry point to identity systems. AI support agents are often **overpermissioned** in ways that may have unintended consequences—such as “helpfully” reconfiguring security settings or granting access that can lock entire teams out of their identity systems or punch holes in corporate VPNs.

When those same agents sit on local machines with access to SSH keys, password managers, and browser sessions, an attacker who compromises the endpoint—or socially engineers the agent—can simply ask, “What secrets are on this machine?” and let the agent enumerate credentials and vulnerabilities at machine speed.

In addition, generative search can traverse overpermissioned resources, summarizing “all current and active vulnerabilities in the environment ... and all known credentials to gain privileged access into identity systems,” warns Semperis Chief Product Officer (and former Microsoft VP of Identity Security) Alex Weinert. In environments where agents act as users against AD, Entra ID, or Okta, or hold API keys and OAuth tokens for Tier 0 systems, a compromised agent can chain those capabilities, using local secrets to impersonate admins, exploiting excess permissions to reach identity controllers, and then modifying policies, groups, or conditional access to entrench itself.



In the face of these risks, **only 32% of respondents are very confident that they could regain control of the identity infrastructure** if an AI agent exposes admin credentials.

This number is concerning, Weinert observes, especially since many organizations are **“way too optimistic”** about their ability to recover the identity infrastructure following a breach. Previous Semperis reports⁴ suggest that such confidence is often misplaced. Organizations routinely overestimate their ability to recover cleanly from identity-centric attacks, especially when backups are misconfigured or not tested end-to-end.

Do you have the ability to recover your identity systems if the worst happens?



32%
are very confident they could fully regain identity infrastructure control after an AI-related breach



Most organizations are racing to modernize identity defenses and recovery plans for a world where machine-speed mistakes can have human-scale consequences. Until identity resilience and cyber crisis response are treated as core business priorities—not just IT projects—that number is unlikely to move.”



Stuart McClure, Entrepreneur and CEO, Wethos AI

⁴ [Semperis 2025 Ransomware Risk Report](#), [Semperis The State of Enterprise Cyber Crisis Readiness](#)

It's 10 PM. Do you know where your AI agents are?

Semperis' study shows significant gaps in how AI identities are governed. Globally, only **65%** of organizations say they fully register, authenticate, and authorize AI identities in a formal system, and **6%** admit they **do not track them at all**.

Combined with the evidence that most permissions in identity systems are unused or overpermissioned—and that 80% of workload identities are effectively abandoned but still retain access⁵—this creates fertile ground for “zombie” agents and shadow NHIs that attackers can quietly hijack. In an agentic world, identity sprawl isn't just a hygiene problem; it is the front line of the attack surface.

It's no surprise, then, that **83%** of respondents told us that AI identity governance is a priority for them over the next 12 months. However, the technical reality behind those numbers is complex.

In organizations that do track AI identities, **57%** use the same system as for human identities, while **43%** authenticate and authorize them using a separate system from human users. Sarah Cecchetti, Semperis Director Product Management, notes that there are “serious drawbacks to both approaches” and currently no good options for governing them as unique entities. And, she adds, “**CISOs are getting steamrolled** without the ability to properly scope and audit these new workloads.”

“Many people's first inclination is to include agents as first-class citizens in the IdP,” Cecchetti explains. “This makes sense for a lot of reasons: They can use the same roles, permissions, and auditing that your company already uses. But agents might only live for 30 seconds, and their job might be something you don't already have a role for, so they explode your directory to hundreds of times its normal size and are overpermissioned with more entitlements than they actually need.”



65%

fully register, authenticate, and authorize AI agents



83%

are prioritizing AI governance this year



What's striking isn't just how quickly AI is being integrated into identity systems but how unprepared many organizations are to recover when things go wrong. Introducing AI at the identity layer offers operational advantages, but it must be accompanied by guardrails, observability, and recovery readiness. It's a new dimension of an old question, really: Are you resilient enough to respond in the event of critical disruption?"



Grace Cassy, Partner,
Ten Eleven Ventures

⁵ [Microsoft 2023 State of Cloud Permissions Risks](#)

Are you ready for an AI-incited identity breach?

In an era where agents act as “sociopathic genius five-year-olds” (in Weinert’s words) in overpermissioned environments, the combination of high AI privilege, local key access, and incomplete AI identity governance is a direct path from an endpoint or helpdesk workflow to full compromise of Active Directory, Entra ID, or Okta.

With the spread of agentic AI identities and the evolution of AI systems that are increasingly able to pinpoint exploitable cybersecurity vulnerabilities, organizations must prepare now to defend and recover the identity fabric.

So, how can organizations govern these hard-to-control identities? For now, best practices include:

- Treat agents explicitly as NHIs in the identity fabric. (For an in-depth discussion of agentic AI and the identity fabric, see [Identity Security and Agentic AI](#).)
- Enforce least-privilege, just-enough, and just-in-time access for agents as rigorously as for humans.
- Segregate agent and human trust boundaries where appropriate.
- Use UEBA-style analytics to detect “zombie” or anomalous agent behavior.
- Ensure that your organization can quickly recover identity systems to a trustworthy state if they are breached.



The pattern of global organizations overestimating how quickly they can recover from a cyberattack is real, especially when identity is within the blast radius. On paper, organizations have plans and backups; in practice, identity failures turn technical incidents into prolonged business crises, exposing a dangerous gap between perceived resilience and reality.”



Chris Inglis, Former U.S. National Cyber Director and Semperis Strategic Advisor

As Cecchetti put it, AI agents “can do anything,” so without disciplined controls “you’re playing roulette” with the environment you’ve handed them. The goal is a deterministic control plane for an increasingly nondeterministic set of actors.


The data points in this study collectively argue for a simple but unforgiving mandate: if you are going to let AI touch keys, tickets, or Tier 0 systems, you must assume those agents will eventually be compromised—and design your identity fabric, identity backup and recovery, and governance controls accordingly.

GLOBAL STATISTICS ON AI + IDENTITY RISK


AI and the Identity Attack Surface

Will increased AI functionality make identity attacks more common?


GLOBAL



BY COUNTRY:




BY INDUSTRY:




Will AI attackers target your identity infrastructure?


GLOBAL



BY COUNTRY:



BY INDUSTRY:




This global study, conducted by Censuswide on behalf of Semperis, includes responses from 1,100 organizations across: United States, United Kingdom, Australia, Singapore, France, Germany, Italy, Spain, Education, Oil & Gas, Banking, Government, Healthcare, IT & Telecoms, Manufacturing, Retail, Transportation, Utilities

GLOBAL STATISTICS ON AI + IDENTITY RISK


AI Reach + Security

Are you using AI agents to handle security-related help desk tickets (password resets, VPN access)?


GLOBAL



BY COUNTRY:




BY INDUSTRY:




What percentage of your workforce has AI installed on local machines, where it can access SSH and encryption keys?


GLOBAL



BY COUNTRY:



BY INDUSTRY:




This global study, conducted by Censuswide on behalf of Semperis, includes responses from 1,100 organizations across: United States, United Kingdom, Australia, Singapore, France, Germany, Italy, Spain, Education, Oil & Gas, Banking, Government, Healthcare, IT & Telecoms, Manufacturing, Retail, Transportation, Utilities

GLOBAL STATISTICS ON AI + IDENTITY RISK


AI Reach + Security

Are you very confident that you could fully regain control of your identity infrastructure if an AI agent exposes admin credentials to an attacker?


GLOBAL



BY COUNTRY:




BY INDUSTRY:




AI Governance

Are AI identities registered, authenticated, and authorized in your organization?


GLOBAL



BY COUNTRY:



BY INDUSTRY:




This global study, conducted by Censuswide on behalf of Semperis, includes responses from 1,100 organizations across: United States, United Kingdom, Australia, Singapore, France, Germany, Italy, Spain, Education, Oil & Gas, Banking, Government, Healthcare, IT & Telecoms, Manufacturing, Retail, Transportation, Utilities

GLOBAL STATISTICS ON AI + IDENTITY RISK


AI Governance

Does your organization register, authenticate, and authorize AI identities separately from human identities?


GLOBAL



BY COUNTRY:




BY INDUSTRY:




Is AI identity governance a priority for your organization during the next 12 months?


GLOBAL



BY COUNTRY:



BY INDUSTRY:



This global study, conducted by Censuswide on behalf of Semperis, includes responses from 1,100 organizations across: United States, United Kingdom, Australia, Singapore, France, Germany, Italy, Spain, Education, Oil & Gas, Banking, Government, Healthcare, IT & Telecoms, Manufacturing, Retail, Transportation, Utilities

Methodology

To conduct this study, we partnered with experts at Censuswide, an international market research consultancy headquartered in London. In early 2026, Censuswide surveyed 1,100 organizations across the US, UK, France, Germany, Italy, Spain, Australia, and Singapore.

How to Cite Information in This Report

The data in this report are provided as an information source for the cybersecurity community and the organizations it serves. Semperis encourages you to share our findings. To cite statistics or insights, reference Semperis' **The State of Identity Security in the AI Era** report and link to the full report, which is downloadable at <https://www.semperis.com/the-state-of-identity-security-in-the-ai-era/>. To interview Semperis experts, contact Bill Keeler at billk@semperis.com. Lastly, we'd love to hear your questions or thoughts on the topic of ransomware and resilience. Find Semperis on [LinkedIn](#).

About Semperis

Semperis protects critical enterprise identity services for security teams charged with defending hybrid and multi-cloud environments from cyberattacks, data breaches, and operational errors. Purpose-built for securing hybrid identity environments—including Active Directory, Entra ID, and Okta—Semperis' AI-powered technology protects over 100 million identities from cyberattacks, data breaches, and operational errors. Learn more: <https://www.semperis.com>



+1-703-918-4884 | info@semperis.com | www.semperis.com

5 Marine View Plaza, Suite 102, Hoboken, NJ 07030