

ServiceNow Research Study Warns Security Leaders: Ungoverned Agentic AI Exposure Is Compounding Inside Enterprises

2026-10-20

90% of organizations worldwide already had an agentic AI security incident in the last year, finds new research surveying 3,700+ global security leaders

Only 13% of businesses have a defensible security posture; ServiceNow calls on enterprises to close the gap between AI chaos and AI control

SANTA CLARA, Calif.--(BUSINESS WIRE)-- **ServiceNow (NYSE: NOW)**, the Super Intelligence Control Tower for business reinvention, today warned that ungoverned agentic AI exposure is compounding inside the enterprise, based on **new global research** of senior security leaders. The warning comes as AI agents move out of pilots and into the systems that run hospitals, airports, factory floors, financial institutions, and beyond. Environments where a single ungoverned action can mean devastating impacts on society. Nine in ten global organizations have already experienced an agentic AI security incident in the past year, and only 36% of those businesses contained it in under 24 hours.

"The agentic incidents making headlines this year are a warning about what can happen when AI is deployed without the right controls," said Ben de Bont, CISO, ServiceNow. "Organizations shouldn't have to choose between moving quickly and managing risk. The ones that get this right will build governance into how they deploy agents from the start, so they can realize the benefits of AI without losing visibility or control."

Inside the AI chaos: what the research found

According to **Deployed but ungoverned: the state of agentic security**, while many security leaders self-assess as having established agentic security postures, only 13% have defensible programs when measured against five key domains: attack surface, accountability, security by design, visibility, and audit readiness.

Most agentic AI is currently governed by legacy frameworks and human-driven systems, with less than half of organizations using governance frameworks that cover key elements of AI agent deployments. Just 44% of security teams globally design security into an agent's development before it's built, which means for the majority, permissions and guardrails are being implemented on an ad-hoc, after-the-fact basis. That's reactive by design.

- Scale and blind spots compound organizations' exposure
 - Forty-nine percent of organizations now run between 100 and 999 AI agents, and 35% run more than 1,000.
 - Fifty-seven percent of security teams have a full, accurate record of every AI agent in production and have processes to track changes.
 - Twenty-nine percent of security teams found AI agents running in production that had never been registered or known. What gets missed can surface later as shadow AI, discovered only after it's already been running unsupervised.
- Machine-speed risk is outpacing human-speed response
 - Seventy-six percent of organizations say they maintain a complete, current map of every agent's access, but only 19% can see more than 90% of their agents in real time.
 - Even organizations that can see a problem often can't stop it fast enough. Two-thirds (66%) cannot act automatically when an agent behaves anomalously, which means a human still has to intervene before anything changes.

Closing the gap to restore control

ServiceNow's **Autonomous Security** offering helps move companies from reactive, fragmented security to enable prevention-first, autonomously governed security that can operate faster than AI creates risk. Operating on the **Shift Zero** imperative, ServiceNow is supporting organizations to achieve as close to zero exposure at all times as possible, enabling enterprises to answer, with proof, what systems are doing, why, and who/what is accountable, while AI moves as fast as the business needs. This includes the rules and guardrails to safely adopt agentic AI, including a **kill switch** and **AI Passport**.

"Innovation and security have to move together, not in sequence," said Nadir Izrael, GVP, Cybersecurity and Risk, ServiceNow. "That's what Autonomous Security and Shift Zero are built for: governance embedded the moment an agent is built or deployed, not bolted on after. Security that runs at machine speed stops being the brake on AI

adoption and becomes what makes scaling it safe."

Methodology

Deployed but ungoverned: the state of agentic security was conducted by ServiceNow in conjunction with Phronesis Partners. The study surveyed more than 3,700 senior security and technology decision-makers, VP level and above, across 20 countries. Respondents represent midsize and large organizations, with 97% reporting annual revenue of \$250M or more.

About ServiceNow

ServiceNow (NYSE: NOW) is the Super Intelligence Control Tower for business reinvention. The ServiceNow platform integrates with any model, any cloud, and any data source to orchestrate, govern, and secure work across the agentic enterprise. ServiceNow gives customers the enterprise context and control to build and scale confidently. With more than 100 billion workflows running on the platform each year, ServiceNow helps organizations turn fragmented operations into coordinated, autonomous work for IT, CRM, Security, Employee Experience, and more. Learn how ServiceNow puts AI to work for people at www.servicenow.com.

Media Contact

Nicole Rosenberg

401.594.3673

press@servicenow.com

Source: ServiceNow