The Agentic Ecosystem Security Gap: 2026 CISO Report

Why 99% of Organizations Were Breached Through Their SaaS and AI Ecosystem Despite Record Security Investment

A survey conducted by Consensuswide
500 U.S. CISOs | January–February 2026

vorlon

Five Findings Every Security Leader Needs to See

The 2026 CISO Report surveyed 500 U.S. security leaders about their SaaS and AI ecosystem security posture, tooling, incidents, and preparedness. The data reveals a structural gap — not a vendor quality problem — between the security architecture most organizations have and the one the agentic era requires.

2026 CISO Report Results

99.4%

of organizations experienced at least one SaaS or AI ecosystem security incident in 2025 — despite running an average of 13 dedicated security tools.

1 in 3

enterprises experienced suspicious AI agent activity in 2025 — Year One of serious enterprise AI deployment.

83–87%

of security teams report limitations in every capability required to address the threat. This is structural, not selective.

Fewer than half

of CISOs claim comprehensive SecOps coverage for SaaS and AI across exposure management, incident response, and threat hunting.

98.2% concerned

about a SaaS supply chain breach in 2026. Only **.8%** feel adequately protected.

vorlon

The Agentic Ecosystem is Now Largest Attack Surface in the Enterprise

The agentic ecosystem is the converged layer of SaaS applications, AI agents, API integrations, non-human identities, and the sensitive data flows between them. It is not static. It expands every time an employee connects a new AI tool, every time a SaaS vendor adds an integration, and every time an AI agent is granted permissions to act.

It moves at machine speed. ServiceNow describes a routine IT ticket resolution — one employee, one VPN issue — where an AI agent autonomously touches identity systems, permissions, and configurations across multiple systems in minutes, with no human in the loop.

That same agent is also authorized to act in Okta, Slack, GitHub, DocuSign, SAP Concur, and Oracle HCM: identity, collaboration, code, contracts, expenses, and payroll. One agent. Dozens of systems. Each logs its own slice. Nobody sees the full picture.

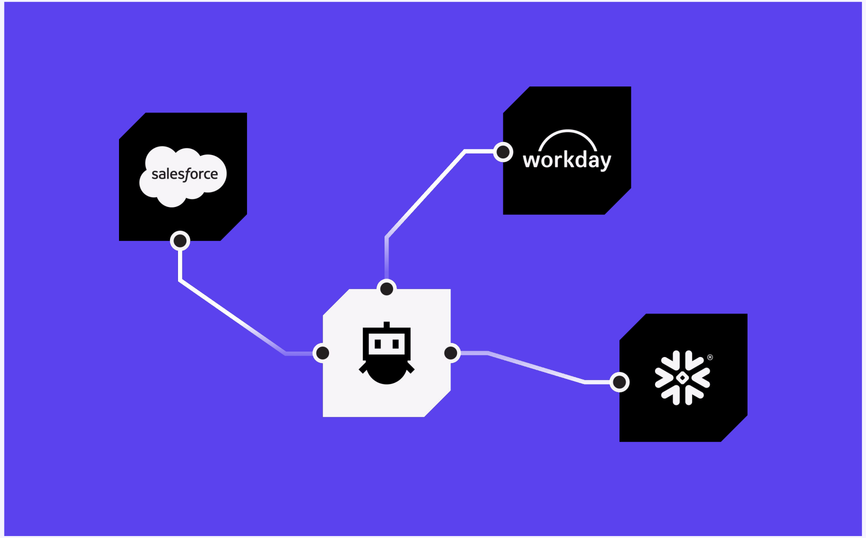

Survey Methodology

The Agentic Ecosystem Security Gap: 2026 CISO Report was conducted by Consensuswide, an independent research firm from January 27 to February 9, 2026. Consensuswide is a member of the Market Research Society (MRS) and the British Polling Council (BPC). They adhere to the MRS Code of Conduct and ESOMAR principles, a set of ethical guidelines for market, opinion, and social research. All statistics are verified against raw survey data.

500

U.S. CISOs surveyed

500+

employee organizations

17

questions

All major

industry verticals represented

Company size range: 500 – 10,000 employees (1 firm with over 10K)

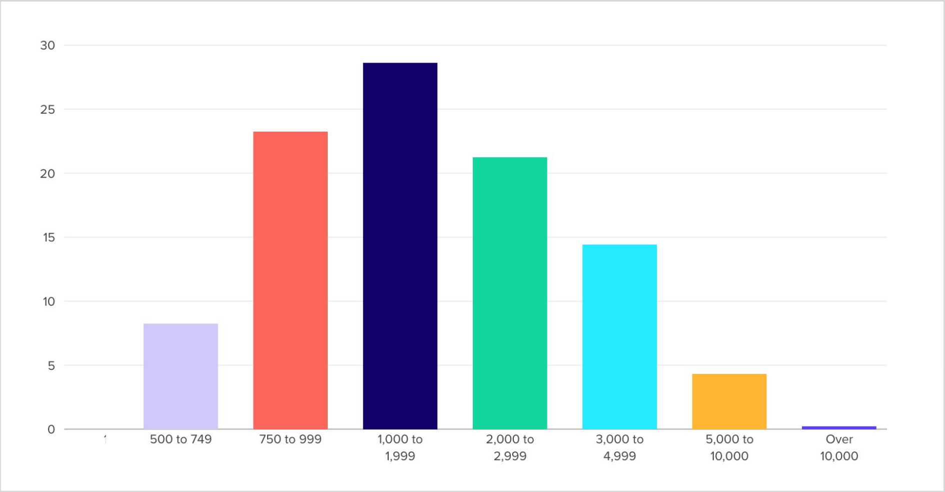

GARTNER'S #1 CYBERSECURITY TREND FOR 2026

Gartner Named Agentic AI Oversight the #1 Cybersecurity Trend for 2026

> Agentic AI is rapidly being used by employees and developers, creating new attack surfaces. No-code/low-code platforms and vibe coding expand this further, driving unmanaged AI agent proliferation, unsecured code and potential regulatory compliance violations.
>
> Cybersecurity leaders must identify both sanctioned and unsanctioned AI agents, enforce robust controls for each and develop incident response playbooks to address potential risks.

Alex Michaels - Director Analyst, Gartner

Alex Michaels, "Gartner Identifies the Top Cybersecurity Trends for 2026," Gartner Newsroom, February 5, 2026.
www.gartner.com/en/newsroom/press-releases/2026-02-05-gartner-identifies-the-top-cybersecurity-trends-for-2026

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact.

YEAR ONE: AI AGENTS
ALREADY A BREACH VECTOR

One in Three Enterprises Hit by an AI Agent Security Incident in 2025

In 2025 — the first year of serious enterprise AI deployment — security incidents involving AI agents were already occurring at about the same rate as social engineering via SaaS attacks (33.6%) and SaaS supply chain attacks (30%). These are not hypothetical future risks. They are current ones.

75.4%

of CISOs say AI agents are a critical or significant security risk

30.4%

experienced suspicious AI agent activity in 2025

30.8%

experienced unauthorized data exfiltration through SaaS-to-AI integrations

CISOs Confident in AI Tools They Know. Long Tail is Where Visibility Breaks Down.

80 to 85% of CISOs report confidence in understanding what data their deployed big name AI tools — ChatGPT, Claude, Copilot, Gemini — can access. But for the long tail of shadow AI, confidence drops sharply. When asked about "other AI tools" beyond the big names, confidence drops to 65.4%, with 25% reporting not very confident or not confident at all.

When asked…How confident are you, if at all, that you know which AI tools are accessing data stored in which SaaS applications? CISOs…

Named AI Tools (ChatGPT, Claude, Copilot, Gemini)

82.8%	very or somewhat confident in data access visibility
15.7%	not very or not at all confident
42%	very confident

Other AI Tools

65.4%	very or somewhat confident in data access visibility
25%	not very or not at all confident
21%	very confident

99.4% of Organizations Experienced a SaaS or AI Ecosystem Security Incident in 2025

Only 3 of 500 CISOs reported no security incidents from the list below. Every other organization — regardless of company size, industry, or security investment — experienced at least one incident last year. No category was rare. CISOs experienced each incident type at roughly the same rate, between 27% and 34%. This is not a tail risk. It is the baseline.

99.4%

of organizations experienced at least one SaaS or AI ecosystem security incident in 2025

3 of 500

Only **3 of 500** CISOs reported zero incidents

Incident breakdown by category:

Social engineering via SaaS	33.6%
SaaS-to-AI data exfiltration	30.8%
Suspicious AI agent activity	30.4%
Supply chain attack via SaaS vendor	30%
Compromised OAuth tokens	27.4%
Third-party SaaS vendor breach	27.4%

High Confidence. Higher Breach Rates. Something Doesn't Add Up.

The survey reveals a set of contradictions that expose something more troubling than a simple confidence gap. CISOs are not just overestimating their protection — they are claiming capabilities and experiencing outcomes that cannot simultaneously be true. This is not a reflection on individual security leaders. It is evidence that an entire category of tools creates the appearance of coverage without delivering the substance of it.

What CISOs Claim	
89.2%	claim strong or comprehensive OAuth token governance
78.6%	claim a comprehensive, real-time data flow map across SaaS and AI
77%	claim comprehensive behavioral monitoring with data-layer context

What Actually Happened	
27.4%	were breached through compromised OAuth tokens or API keys
86.8%	say they cannot see what data AI tools are exchanging with SaaS applications
30.8%	experienced unauthorized SaaS-to-AI data exfiltration

Configuration audits look like monitoring. Permission reviews look like governance. Single-application detection looks like ecosystem visibility. The gap between what these tools report and what they actually see is where breaches live.

Security Was Built for the Front Door. The Threat Moved to the Engine Room.

The tools most enterprises rely on were designed to monitor the front door: application configurations, user login events, permission settings — built for human users, browser-based access, and application-by-application risk. The attack surface has moved to the engine room: the runtime layer where AI agents move sensitive data between systems, where OAuth tokens grant persistent cross-platform access, where a single compromised integration can cascade silently across an entire SaaS and AI supply chain.

THE FRONT DOOR	THE ENGINE ROOM
Application configurations	AI agent actions at runtime
User login events	OAuth token data flows
Permission settings	API-to-API integrations
Browser-based access	MCP server communications
App-by-app risk management	Cross-app data movement
What legacy tools see	What legacy tools miss

83–87%: THE STRUCTURAL LIMITATION

Limiting Factors of Current SaaS and AI Security Tools Affect 83–87% of Orgs

When CISOs were asked to rate their current tooling across the 11 limitations shown on the right, every factor was rated as having some level of limitation (mindor, moderate, or major) by 83–87% of organizations. The range spans only four percentage points. This is not evidence that some tools are better than others. It is evidence that the entire existing architecture shares the same structural deficiencies.

Limiting Factors of Current SaaS and AI Security Tools	% Reporting
Cannot see sensitive data flows across applications	87%
Cannot see what data AI tools are exchanging with SaaS apps	86.8%
Focus on configuration/compliance, not runtime threats	86.2%
Too many siloed tools, no unified view	85.8%
Lack behavioral analytics and anomaly detection	85.8%
Alerts lack context and clear remediation guidance	85.6%
Limited or no coverage of AI tools and integrations	85.4%
Cannot coordinate response across SaaS applications	85.4%
Cannot detect new or risky integrations	84.8%
Cannot detect OAuth token or API key abuse	84.8%
Cannot distinguish human from non-human behaviors	83.4%

The Tools Most Often Cited as the Answer Were Built for the Front Door

SSPM (SaaS Security Posture Management) and SSE/SASE (Security Service Edge / Secure Access Service Edge) are the two categories most frequently positioned as the solution to SaaS and AI ecosystem security. Both are legitimate investments. Neither was built to see what happens in the engine room. SSPM audits what permissions exist — not what agents do with those permissions at runtime. SSE/SASE monitors what traverses the network perimeter — not API-to-API data flows that bypass the network entirely.

39%

of organizations use an SSPM tool

42.8%

Of those, 42.8% say it only detects within individual applications — or functions primarily as a configuration and compliance audit tool, not real-time cross-platform threat detection

What They're Designed For	What They Can't See
SSPM: Configuration and compliance auditing	What an AI agent does with the access it has
SSE/SASE: Network perimeter monitoring	API-to-API data flows that bypass the network
Both: Application-layer, human-speed threats	Agent runtime, OAuth persistence, cross-SaaS data movement

No Industry Consensus on Who Owns the Impact Assessment

When a SaaS vendor announces a breach, there is no industry consensus on who owns the impact assessment. Responses span nine organizational functions with no single team cited by more than 21.8%, suggesting this new attack surface has yet to find a settled home in the enterprise.

Function	Ownership
SaaS security team	**21.8%**
IT security leadership	**13.8%**
Data security	**11.6%**
IT operations	**10.6%**
Security operations	**10.6%**
Risk and compliance	**9.8%**
Cloud security	**9.2%**
Security engineering	**7.2%**
Application owner (e.g., Salesforce Admin)	**5.2%**
No defined owner	**0.2%**

Fewer Than Half of Security Teams Have Comprehensive SaaS and AI Coverage in Any Core SecOps Workflow

Security operations teams have spent years building mature workflows for endpoints and physical infrastructure. That investment has not extended to the SaaS and AI ecosystem. Fewer than half of CISOs reach comprehensive coverage in any of the three core SecOps workflow areas. The good news: 93%+ plan to add or expand coverage across all three, with nearly half intending to do so within 12 months.

	Exposure Management	Threat Hunting & Investigation	Incident Response
Comprehensive coverage	**41.8%**	**44%**	**38.2%**
Partial or minimal	**55.2%**	**53.6%**	**59.6%**
Plan to expand	**93.8%**	**93.4%**	**93.4%**
Within 12 months	**47.6%**	**47.2%**	**45%**

99.2% are Concerned About a SaaS Supply Chain Breach in 2026. Only 0.8% Feel Protected.

Following high-profile SaaS and AI supply chain breaches in 2025, including the Salesforce ShinyHunters vishing attack, the Salesloft/Drift OAuth hijack, and the Gainsight supply chain compromise, 99.2% of CISOs report concern about a similar incident in 2026.

99.2%

concerned about a SaaS supply chain breach in 2026

0.8%

feel adequately protected against one

Supporting stats:

46.6%	call it a top priority risk
30%	already experienced a supply chain attack in 2025
51.2%	Only 51.2% have an automated IR playbook for active SaaS exfiltration
48.8%	would rely on manual response — human-speed processes against API-speed breach cascades

Budgets Are Increasing. The Question is Whether the Architecture Changes With Them.

More than 86% of organizations plan to increase their SaaS security budget in 2026, and 84% plan to increase their AI security budget. But budget directed at the same tool categories will compound operational complexity without closing the coverage gap. 99.4% were breached despite running an average of 13 dedicated security tools. More tools in the same categories will produce the same results.

SaaS Security Budget 2026 vs. 2025:

- Increase significantly (>25%): 7.4%
- Increase moderately (10–25%): 42.4%
- Increase slightly (<10%): 37%
- Stay the same: 7% | Decrease: 6.2%
- **Total increasing: 86.8%**

AI Security Budget 2026 vs. 2025:

- Increase significantly (>25%): 12.8%
- Increase moderately (10–25%): 42.6%
- Increase slightly (<10%): 28.8%
- Stay the same: 10.4% | Decrease: 5.4%
- **Total increasing: 84.2%**

Securing the Agentic Enterprise Requires Visibility at the Ecosystem Layer — Not the Application Layer

The organizations that close the security gap in the agentic era will be those that extend security operations to the ecosystem execution layer — where AI agents act, OAuth tokens persist, and sensitive data moves across interconnected systems at machine speed. Comprehensive coverage means the ability to see, detect, and respond across the full agentic ecosystem — not one application at a time.

WHAT SECURITY TEAMS NEED

Continuous discovery

All SaaS apps, AI agents, integrations, and non-human identities — including shadow AI and shadow integrations

Cross-app data flow mapping

How sensitive data moves between every SaaS app, AI tool, and integration in near real time

Behavioral monitoring for human and non-human identities

Data-layer context that distinguishes a compromised agent from normal operation

AI Agent Flight Recorder

Forensically complete, cross-SaaS audit trail of every agent action, mapped to sensitive data and blast radius

Blast radius calculation

Answering the board-level question in minutes: which data, which systems, which identities are at risk

Cross-app coordinated response

Native SecOps integration across exposure management, threat hunting, and incident response

The Agentic Workforce is Already Here. Security Needs to Catch Up.

The data in this report does not describe a future threat. It describes 2025 — Year One of serious enterprise AI deployment. One in three organizations experienced an AI agent security incident. Nearly all experienced some form of SaaS or AI ecosystem compromise. And the tools in place share a structural limitation: they were built to govern access, not to record what happened after access was used.

The agentic era does not require abandoning existing security architecture. It requires extending it to a layer it was never designed to see. Organizations that recognize this shift now — and build security operations that cover the engine room, not just the front door — will be the ones that avoid the next wave of AI-era breaches.

"The agentic workforce is already here. Vorlon exists to make sure it doesn't operate in the dark."

Amir Khayat
CEO and Co-Founder
Vorlon

ABOUT VORLON

Vorlon. Built for the engine room.

Most security was built for the front door. The threat has moved to the engine room. AI agents move freely across systems. OAuth tokens transfer sensitive data between applications at machine speed. One compromised integration cascades across your SaaS supply chain.

Vorlon is the Agentic Ecosystem Security Platform. Its patented DataMatrix™ technology builds a live model of how sensitive data, identities, and integrations interact across your agentic ecosystem — giving security teams the visibility, forensics, and remediation to manage sensitive data exposure and deploy AI at scale.

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact.

vorlon

SOC 2 Type II certified
Backed by Accel

Recognized by Gartner:
Emerging Tech: Intelligent Simulation Accelerates Proactive Exposure Management, 2025

www.vorlon.io

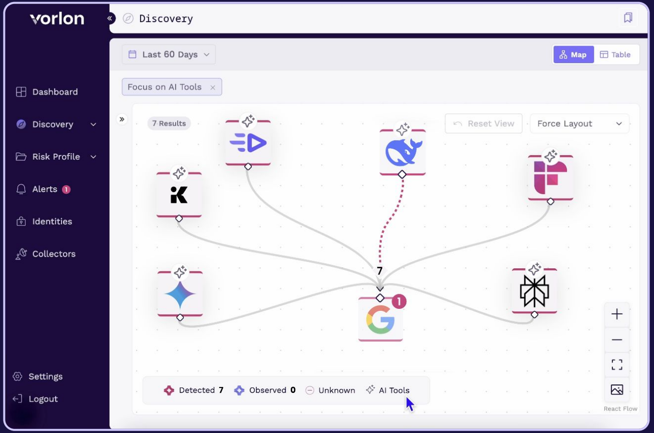
vorlon
Discovery
Last 60 Days
Map
Table
Focus on AI Tools
Dashboard
Discovery
Risk Profile
Alerts 1
Identities
Collectors
7 Results
Reset View
Force Layout
7
1
Settings
Logout
Detected 7
Observed 0
Unknown
AI Tools
React Flow