Secure AI Agents

Industry-first privilege controls for AI agents

Solution Brief

Agentic AI is a top concern

Two-thirds of CISOs surveyed in financia services and software rank agentic AI among their top three cybersecurity risks—and more than one-third name it as their top concern, ahead of ransomware and supply-chain threats.[1]

Organizations need answers

- Do I already have AI agents running in my organization?
- How can I quickly respond if an AI agent is compromised?
- How can I ensure AI agents operate securely?
- How can I govern AI agents to ensure compliance?

Challenges of Agentic AI

Agentic AI is moving from experiment to enterprise staple, with organizations rapidly deploying autonomous agents across critical workloads. This shift brings major security concerns: unchecked access, visibility gaps, unpredictable behavior, and soaring complexity. CISOs worry adoption will outpace their ability to manage the risks.

New identity class with new risks

AI agents inherit the risks of human and machine identities while also introducing unique threats. AI agents are non-deterministic by nature, require privileged access to enterprise resources, and operate at machine speed, often with little human supervision. This combination of factors introduces significant ris s for organizations without adequate security controls.

Expanded attack surface

Agents often require privileged access to sensitive systems, data, and cloud services. Each new agent multiplies the attack surface, and many are “shadow agents,” spun up outside of IT or security oversight.

Scale and management

The huge uptick in AI agent identities need to be onboarded, managed, and deprovisioned without added burden. Security teams, already stretched thin, will be challenged to keep pace – forcing a tough choice between slowing innovation and accepting unmitigated risk.

Compliance and audit

Without a clear audit trail of agent activities — what they accessed, why, and under whose direction — you cannot ensure the same level of compliance for your AI agents as your human and machine identities. As regulations evolve for AI, this will only become more important.

CYBERARK®
A PALO ALTO NETWORKS COMPANY

Why now is the time for Agentic AI

Speed of growth

Overall adoption of agents is projected to nearly double from ~43% of organizations in 2025 to ~76% in 3 years.[1]

Setting secure foundations

Today's low-code/no-code platforms mean almost anyone can develop and deploy an agent. Organizations need to implement strong security controls early.

Regulations are coming

Regulatory bodies are beginning to demand stricter oversight. Don't be left scrambling when compliance becomes mandatory.

AI agents discovered overtime

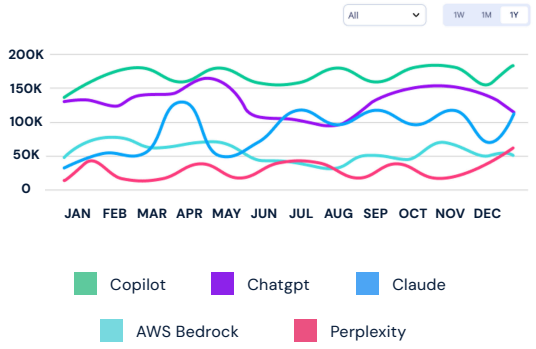

[1]Securing Agentic AI: Identity as the Foundation of Defense (CyberArk, 2025)

CyberArk Secure AI Agents

With an identity-fir t approach to agentic AI security, our Secure AI Agents solution lets you discover, manage, and secure AI agents with the right level of privilege controls to keep your sensitive resources safe. The comprehensive solution addresses four key areas to secure AI agents:

Discovery and context.

The Secure AI Agents solution provides a centralized repository for AI agents and detects agents running across SaaS, cloud, and developer environments including AWS Bedrock and Microsoft Copilot Studio. Each agent is enriched with context, such as ownership, purpose description, status, and permissions, helping you understand who owns each agent, what it does, and what it can access.

Secure access.

Our AI Agent Gateway provides AI agents secure access to resources by leveraging identity security controls. Permissions are granted only for a specific task, with the right l vel of privilege, and revoked automatically, helping to ensure zero standing privileges.

Lifecycle management and compliance.

Manage your agents at every step while getting full visibility into the actions they are taking. See what actions agents perform and on behalf of what user, get comprehensive access reviews that highlight what resources each user can access, know which permissions are in use, and improve lifecycle management with automated ownership workfl ws and more.

Threat detection and response.

It's important to have functionality that can flag abnormal a ent behaviors, so you can immediately suspend that agent. Secure AI Agents provides information on agent actions and how they are used. If an agent behaves abnormally or exceeds its role, you can immediately suspend it.

Learn more about securing AI agents in your organization.

CYBERARK®
A PALO ALTO NETWORKS COMPANY