

Idira Secure AI Agents

Industry-First Privilege Controls for AI Agents

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 nondeterministic by nature, require privileged access to enterprise resources, and operate at machine speed, often with little human supervision. This combination of factors introduces significant risks 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 significant uptick in AI agent identities requires them 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.

Agentic AI Is a Top Concern

Two-thirds of CISOs surveyed in financial 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 we already have AI agents running in our organization?
- How can we quickly respond if an AI agent is compromised?
- How can we ensure AI agents operate securely?
- How can we govern AI agents to ensure compliance?

Why Now Is the Time for Agentic AI

Speed of Growth

Overall agent adoption is projected to nearly double from ~43% of organizations in 2025 to ~76% in 3 years.²

Setting Secure Foundations

Today's low-code and 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.

1-2. *Securing Agentic AI: Identity as the Foundation of Defense*, CyberArk, November 3, 2025.

Idira Secure AI Agents

With an identity-first approach to agentic AI security, you can discover, manage, and secure AI agents with the right level of privilege controls to keep your sensitive resources safe. That's where Idira™ Secure AI Agents by Palo Alto Networks comes in. This comprehensive solution addresses the following key areas to secure AI agents.

Discovery and Context

The Idira 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.




Figure 1. AI Agents discovered over time

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 level of privilege and revoked automatically, helping ensure zero standing privileges.

Lifecycle Management and Compliance

Manage your agents at every step while getting full visibility into the actions they are taking. For example, see the actions that agents perform for a specific user. Get comprehensive access reviews that highlight the resources each user can access. Know which permissions are in use. And, improve lifecycle management with automated ownership workflows.

Threat Detection and Response

Leverage functions that can flag abnormal agent behaviors so you can immediately suspend that agent. Idira Secure AI Agents provide information on agent actions and how they are used. If an agent behaves abnormally or exceeds its role, you can immediately suspend it.

See how to secure the AI agents in your organization. [Learn more.](#)



3000 Tannery Way
Santa Clara, CA 95054
Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087
www.paloaltonetworks.com

© 2026 Palo Alto Networks, Inc. A list of our trademarks in the United States and other jurisdictions can be found at <https://www.paloaltonetworks.com/company/trademarks.html>. All other marks mentioned herein may be trademarks of their respective companies.
idira_sb_idira-secure-ai-agents_042126