




Radware Agentic AI Security & Governance

You can't secure what you can't see.



Background

Artificial Intelligence is transforming how organizations operate, innovate and compete. Across every industry, AI agents are increasingly being used to automate workflows, assist employees, interact with customers and perform complex tasks autonomously.

Unlike traditional AI applications that generate responses and require human oversight, agentic AI systems can make decisions, use tools, access enterprise resources and execute multi-step actions with minimal human intervention.

As organizations accelerate AI adoption, a new reality is emerging: AI is no longer confined to centrally managed platforms. Employees are creating their own agents, connecting to external AI services, deploying developer-hosted assistants and integrating AI capabilities into day-to-day workflows—often without formal review, governance or security approval.

This rapid expansion creates a pressing challenge for security, governance and compliance teams: How do you secure an AI ecosystem that you cannot fully see, understand or govern?

Protecting an organization's agentic ecosystem is imperative. Before organizations can protect AI agents, however, they must first discover them, understand how they operate and establish governance across the entire AI ecosystem.

Radware Agentic AI Protection helps organizations discover, govern and protect their AI environment by providing comprehensive visibility, shadow AI detection, AI governance and compliance reporting, risk posture management and real-time protection against agentic AI threats.

The Challenge: The Rise of Shadow AI

AI adoption is happening faster than most organizations can track.

While some AI deployments are formally approved and managed through centralized IT and security processes, many others emerge organically. Employees experiment with commercial AI services, create custom agents, connect external tools or deploy developer-hosted assistants to solve business challenges. These unsanctioned deployments, commonly referred to as **shadow AI**, introduce visibility, governance and security blind spots.

Organizations frequently struggle to answer critical questions:

- Which AI agents currently exist in our environment?
- Which agents have been approved and which are shadow AI?
- Which employees and business units are using them?
- Which tools and MCP servers do these agents access?
- Which enterprise resources and sensitive data can they reach?
- Are we compliant with AI governance regulations and standards?
- Where are our highest AI-related risks?
- Are agents behaving according to their intended goals?

Without clear visibility and governance, organizations cannot effectively assess risk, demonstrate compliance or secure their AI ecosystem.

Radware Agentic AI Protection addresses these challenges by providing a complete lifecycle approach to AI security. It starts with discovery and governance and extends into risk management and real-time protection.

Radware Agentic AI Protection Solution

Radware Agentic AI Protection delivers end-to-end visibility, governance, compliance and security for enterprise AI ecosystems.

The platform continuously discovers approved and shadow AI agents; maps relationships among users, agents, tools and enterprise resources; evaluates AI security posture; generates audit-ready compliance reports; and protects against emerging agentic AI threats through intent-aware behavioral protection.

Organizations gain complete control over their AI ecosystem from a single platform.

Aligning With Industry Best Practices

Radware's solution is aligned with the **OWASP GenAI Security Project**, particularly the [OWASP Top 10 for LLM Applications \(2025\)](#), [AIVSS Scoring System for Agentic AI Core Security Risks](#) and [OWASP Top 10 for Agentic Applications for 2026](#)

These frameworks provide a structured methodology for identifying and mitigating the unique risks posed by autonomous AI systems.

By adopting these standards, Radware ensures our security posture evolves in step with the rapidly advancing capabilities of agentic AI, allowing us to bridge the gap between traditional cybersecurity practices and the demands of intelligent, self-directed systems.

Solution Core Capabilities



AI Discovery and Shadow AI Visibility

Discover your entire AI ecosystem.

Organizations cannot govern or secure what they cannot see.


Radware continuously discovers AI agents operating across the enterprise, including both sanctioned deployments and shadow AI. The solution builds a complete inventory of agents, tools and AI services while revealing how they interact with users and enterprise resources.

Key capabilities include:

- Continuous discovery of approved and shadow AI agents
- Identification of AI usage across employees, teams and business units
- User-to-agent visibility and ownership mapping
- Agent-to-agent interaction visibility
- Agent-to-tool relationship mapping
- Tool-to-resource dependency analysis
- Discovery of MCP and non-MCP tools
- Visibility into enterprise resource access including DBs, CRMs, email systems and file repositories
- Long-term activity tracking and behavioral analytics
- Interactive Connection Map for complete AI ecosystem visibility
- Transformation of raw AI telemetry into structured and human-readable context

The result is a continuously updated inventory of the organization's AI landscape and the relationships that define it.

Figure 1:
Inclusive Agents &
Tools graph map



AI Governance & Compliance

Establish trust, accountability and control.

As AI adoption grows, governance requirements are rapidly expanding. Organizations increasingly need to demonstrate accountability, traceability and control of AI systems to internal stakeholders, auditors, regulators and customers.

Radware provides complete transparency across the AI lifecycle.

Key capabilities include:

- Continuous agent inventory and lifecycle tracking
- Full audit trail of agent activity
- Prompt, interaction and execution traceability
- User accountability and ownership tracking
- Tool usage monitoring
- Resource access visibility
- Governance reporting aligned to organizational policies
- Executive and auditor-ready reporting

This enables organizations to demonstrate responsible AI usage and implement effective governance programs at scale.



Compliance Reporting

Align with leading AI frameworks.

The AI regulatory landscape continues to evolve rapidly. Radware helps organizations prepare for and demonstrate compliance with leading AI governance and risk management standards through automated reporting and evidence collection.


Supported frameworks include:

- ISO 42001
- NIST AI Risk Management Framework (AI RMF)
- EU AI Act
- Internal AI governance policies

Key capabilities include:

- On-demand compliance reporting generation
- Audit-ready documentation
- Executive and regulatory reporting support

Figure 2:
Reports tuned for
global standards



Organizations gain the visibility required to demonstrate governance maturity while reducing compliance preparation efforts.



AI Security Posture Management (AI-SPM)

Understand and prioritize risk.

Effective AI security begins with understanding where risk exists.

Radware continuously evaluates AI posture and provides contextual risk analysis across agents, tools and enterprise resources. Key capabilities include:

- Dynamic AI Risk Graph
- Agent risk scoring
- Tool risk scoring
- Multi-agent attack path analysis
- Sensitive data exposure identification
- Resource access risk assessment
- Compliance posture evaluation

By understanding where risk exists, organizations can focus resources on the areas that matter most.



Security

Ensure intent-aware behavior-based detection, prevention and response.

Once organizations have visibility into their AI ecosystem and understand their risk posture, they can effectively protect AI operations.

Radware provides real-time behavioral protection that goes beyond static guardrails by understanding agent intent, goals and actions. Key capabilities include protection against:

- **LLM Attacks** – Prompt injection, jailbreaking, use of toxicity, unsafe use and more
- **Agent Behavior Hijack** – Manipulation of an agent's goals to the benefit of the attacker (including indirect prompt injections)
- **Tool Misuse & Exploitation** – Tricking agents into using their tools in harmful or unintended ways
- **Memory & Context Poisoning** – Corrupting agent memory or context to distort reasoning and decision-making
- **Supply Chain Attacks** – Infected agents or tools attacking others further down the supply chain
- **Rogue Agents** – Malicious or compromised external agents acting autonomously to deceive or disrupt protected agent's behaviour

Unlike traditional security technologies that focus on rules and signatures, Radware leverages intent-aware behavioral analysis to identify emerging threats that target autonomous AI systems.



Broad Platform Integration

The Radware solution integrates with leading enterprise AI platforms and custom-built environments.

Supported environments include:

- Microsoft 365 Copilot
- Microsoft Copilot Studio
- Azure AI Foundry
- AWS Bedrock
- Anthropic Claude Code
- OpenClaw
- N8N
- Custom enterprise agents

Integration methods include:

- Out-of-path API integration
- Inline enforcement through OpenAI-compatible frameworks
- Lightweight client-side integrations for developer-hosted agents

Summary

The first challenge of enterprise AI security is visibility. Then comes active security. Organizations must first discover sanctioned and shadow AI; understand how agents interact with users, tools and enterprise resources; establish governance programs and demonstrate compliance with emerging standards. Only then can they effectively reduce risk and defend against agentic AI threats.

Radware Agentic AI Protection delivers this complete lifecycle approach by enabling organizations to: **Discover -> Govern -> Protect -> Comply.**

From shadow AI discovery and AI governance to risk management and intent-aware security, Radware provides the visibility and protection organizations need to confidently embrace the agentic AI era.

Interested in Radware Agentic AI Protection?

Secure your organization against the challenges of the agentic AI era with Radware's comprehensive real-time protection.

Contact Radware

This document is provided for information purposes only. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law. Radware specifically disclaims any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. The technologies, functionalities, services, or processes described herein are subject to change without notice.

© 2026 Radware Ltd. All rights reserved. The Radware products and solutions mentioned in this document are protected by trademarks, patents and pending patent applications of Radware in the U.S. and other countries. For more details, please see: <https://www.radware.com/LegalNotice/>. All other trademarks and names are property of their respective owners.

