REQUEST FOR INFORMATION RESPONSE

SECURITY CONSIDERATIONS FOR ARTIFICIAL INTELLIGENCE AGENTS

March 9, 2026

I. INTRODUCTION

In response to the National Institute of Standards and Technology's ("NIST") request for information on Security Considerations for Artificial Intelligence Agents ("consultation"), CrowdStrike offers the following views.

We approach these questions from the standpoint of a leading international, US-headquartered, cloud-native cybersecurity provider that defends globally distributed enterprises from globally distributed threats. CrowdStrike offers insights informed by multiple practice areas: cyber threat intelligence; proactive hunting, incident response and managed security services; and an AI-powered software-as-a-service cybersecurity platform and marketplace. Accordingly, this perspective is informed by CrowdStrike's role in protecting organizations from data breaches and a variety of other cyber threats.

II. COMMENTS

We appreciate NIST's ongoing efforts to provide resources to organizations which in turn will better protect the nation from cybersecurity threats. NIST's consultations have a history of developing into useful tools, and exploring Artificial Intelligence (AI) agent security is timely. AI agents represent both a critical advancement in cybersecurity capabilities and a potential new attack surface. The security of these systems is paramount as organizations across sectors rapidly deploy AI agents.

While the consultation is focused on the security of AI agents, it is important to understand the role, and encourage the use, of agentic AI in security. In today's threat landscape, defending against AI-accelerated adversaries, and securing AI systems themselves, requires cybersecurity operating at machine speed.

One of the most immediate areas Agentic AI can improve cybersecurity practices is leveraging agents to eliminate bottlenecks in the Security Operations Center (SOC). By deploying specialized agents to tackle time-intensive tasks, security teams can reclaim

a speed advantage, close persistent labor and response gaps, and shift from reactive to proactive defense. Over the past several months, CrowdStrike has launched a series of specialized agents for this purpose including those that can prioritize vulnerabilities and analyze malware.

Such agents are central to a profound change that's underway now to modernize traditional SOCs for the emerging era of the Agentic SOC. A Next-Generation Security Information and Event Management (Next-Gen SIEM) capability will enable organizations to leverage these agents by exposing them to all relevant security data and positioning them to perform workflows like threat hunting and remediation.

As NIST continues work on agentic AI, CrowdStrike recommends agentic AI use in cybersecurity for the protection of AI systems is encouraged.

We do not have feedback on all questions in the consultation, but we do want to offer several responses that may be of use to NIST as it further examines the security of AI agents.

A. *What are the unique security threats, risks, or vulnerabilities currently affecting AI agent systems, distinct from those affecting traditional software systems?*

AI agents are just one part of a growing, and complex, AI tech stack that is changing the threat landscape. As organizations embed AI into core business processes, the attack surface will expand to include AI models, training data, agents, and supply chains. Today, organizations must have increased visibility, and robust security, of AI operations so that adversaries are not able to amplify risk and create exploitable gaps. CrowdStrike's 2026 *Global Threat Report* Attacks found that AI-enabled actors increased by 89% year-on-year, while average breakout time fell to 29 minutes, with the fastest observed at 27 seconds.[1] This shows that the current threat landscape consists of adversaries operating at machine-speed.

In parallel, cloud-conscious intrusions rose 37%, with a 266% increase among state-nexus actors. These trends indicate that AI models, training data, agents, and supply chains should be treated as high-value assets that require robust protection. Security architectures must be prepared to defend against rapid lateral movement,

[1] *CrowdStrike 2026 Global Threat Report*, https://www.crowdstrike.com/en-us/global-threat-report/.

legitimate credential abuse, and cross-domain evasion as baseline conditions rather than edge cases.

Identity and trust relationships are increasingly exploited by adversaries. Valid account abuse accounted for 35% of cloud incidents, and adversaries increasingly subvert federated trust to obtain persistent access. AI systems compound these risks: model development depends on complex software supply chains, while agentic and prompt-driven systems introduce novel manipulation vectors.

Security must be built into AI systems. AI systems depend on a complex tech stack that includes hardware, GPUs, cloud workloads, training data, token factories, and AI applications and agents. Each layer of this stack must ensure security from development through deployment and use. The layers of the new AI stack are the attack surface of the future, and it must be protected with a proactive and holistic security approach.

Secure AI infrastructure, which includes AI agents, must therefore be designed as an integrated, cross-domain security architecture that preserves confidentiality and integrity of AI applications and sensitive data without materially degrading performance. This requires strong, granular identity governance (including non-human identities or AI agents), hardened software supply chains, and unified telemetry across endpoints, cloud, SaaS, and orchestration environments.

B. *What technical controls, processes, and other practices could ensure or improve the security of AI agent systems in development and deployment?*

Security teams increasingly require specialized capabilities to protect the enterprise throughout the AI lifecycle. An emerging security category called AI Detection and Response (AIDR) enables security teams to: monitor AI behavior and interactions; secure AI agent identities and access; discover shadow AI; prevent sensitive data leakage to AI systems; protect cloud-based AI applications; and implement guardrails on AI system actions.

AI agents are often quickly deployed across SaaS environments by employees, without centralized tools to govern them. While the intent is productivity, the result can be organizational risk. Organizations that do not leverage an AIDR tool could lack crucial visibility into which agents exist, what they can access, and how they behave over time. AI agents deployed without guardrails introduce new attack surfaces that adversaries are eager to exploit.

Emerging best practices to securing AI agents include: 1) mapping each AI agent to its human creator (and leveraging more granular identity controls and permissions where possible); 2) detecting anomalous behavior; and 3) enforcing policies and guardrails across an organization.[2] Security teams maintain control over AI agents by having robust visibility, and by continuously discovering misconfigurations, shadow AI, and access.

C. *Which cybersecurity guidelines, frameworks, and best practices are most relevant to the security of AI agent systems?*

Standard cybersecurity practices also apply broadly to AI agent security. We recommend several approaches to strengthen both cyber and AI agent security:

- **Zero Trust Architecture:** Zero Trust is an impactful concept for increasing cybersecurity in AI environments. By removing implicit trust and continuously validating every stage of digital interaction, organizations can significantly reduce the risk of unauthorized access. Zero Trust principles can be applied across the AI infrastructure to include agents.

- **Identity Threat Detection and Response (ITDR):** Identity has become a primary attack vector for adversaries. In AI environments, this extends beyond human identities to include AI agents. Human, and non-human, identities can have access to sensitive resources and must be governed effectively. We recommend implementing ITDR solutions that continuously monitor user and system activity, detect unusual behavior, and alert security teams to potential compromise.

- **EDR and Next-Gen SIEM:** Endpoint Detection and Response (EDR) capabilities are a core pillar of most contemporary sophisticated security programs. EDR provides granular visibility of potential threats. This enables holistic, real-time threat detection and proactive threat prevention. EDR coupled with a Next-Gen SIEM solution allows organizations to leverage advanced telemetry and observability specifically designed for AI environments to provide granular visibility into activities across the AI infrastructure ecosystem. These capabilities allow for monitoring of data access patterns, compute resource usage, model behavior, agents, and network communications to detect anomalies that may

[2] *How CrowdStrike Secures AI Agents Across SaaS Environments*, August 5, 2025, https://www.crowdstrike.com/en-us/blog/how-crowdstrike-secures-ai-agents-pervading-saas-environments/.

indicate compromise. High-fidelity telemetry is essential for threat hunting across AI environments and enabling rapid response to emerging threats to AI agents.

As NIST explores potential guidelines on securing AI agents, it is worth noting that fixed technical requirements may quickly become obsolete given the pace of AI advancement and the adaptive nature of threat actors exploiting AI systems. Prescriptive standards can lead to a compliance-first mindset, where meeting narrow criteria takes precedence over meeting the desired security outcomes. We recommend an adaptive, risk-based approach that prioritizes flexibility, encourages innovation, and allows organizations to tailor implementation to their risk profile and domain-specific needs. Any AI agent specific guidance should aim to proactively harmonize with applicable AI, cybersecurity, privacy, and data protection existing NIST guidelines.

III. CONCLUSION

NIST's effort on AI agent security represents a valuable step forward in helping organizations manage the opportunities and risks presented by AI. As AI continues to evolve rapidly, it is crucial that security guidance remains flexible while addressing core security principles. As NIST moves forward, we recommend continued engagement with stakeholders.

IV. ABOUT CROWDSTRIKE

CrowdStrike (NASDAQ: CRWD), a global cybersecurity leader, has redefined modern security with the world's most advanced cloud-native platform for protecting critical areas of enterprise risk – endpoints and cloud workloads, identity and data.

Powered by the CrowdStrike Security Cloud and world-class AI, the CrowdStrike Falcon® platform leverages real-time indicators of attack, threat intelligence, evolving adversary tradecraft, and enriched telemetry from across the enterprise to deliver hyper-accurate detections, automated protection and remediation, elite threat hunting, and prioritized observability of vulnerabilities.

Purpose-built in the cloud with a single lightweight-agent architecture, the Falcon platform delivers rapid and scalable deployment, superior protection and performance, reduced complexity, and immediate time-to-value.

CrowdStrike: We stop breaches.

Learn more: https://www.crowdstrike.com/.

CONTACT

We would welcome the opportunity to discuss these matters in more detail. Public policy inquiries should be made to:

Drew Bagley
VP & Counsel, Privacy and Cyber Policy

Elizabeth Guillot
Senior Manager, Public Policy

Email: policy@crowdstrike.com

***