REQUEST FOR COMMENT RESPONSE

ACCELERATING THE ADOPTION OF SOFTWARE AND AI AGENT IDENTITY AND AUTHORIZATION

April 2, 2026

I. INTRODUCTION

In response to the National Institute of Standards and Technology's ("NIST") request for comment on Accelerating the Adoption of Software and AI Agent Identity and Authorization ("draft publication"), CrowdStrike offers the following views.

We approach these comments as a leading international, U.S.-headquartered, cybersecurity provider that defends enterprises from globally distributed threats. Our AI-powered, cloud-native platform delivers real-time detection, automated protection, and threat intelligence across endpoints, cloud, identity, data, and AI workloads. Our insights are informed by applied, real-world experience across global industries through our software-as-a-service cybersecurity platform, incident response, managed threat hunting, and managed security services. CrowdStrike's recommendations are grounded in this operational experience.

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 identity and authorization is timely. AI agents represent both a critical advancement in cybersecurity capabilities and a potential new attack surface. The security - to include identity practices - of these systems is paramount as organizations across sectors rapidly deploy AI agents.

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.

While we do not have feedback on all questions in the draft publication, we did want to offer several responses that may be of use to NIST as it further examines AI agent identity.

A. *What enterprise use-cases are organizations currently using agents for? / What opportunities do agents present?*

The draft publication notes that AI agents for security are a real-world use case being leveraged today by organizations. 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.[1]

Such agents, along with agentic managed detection and response (MDR) capabilities, 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 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.

B. *What standards exist, or are emerging, to support identity and access management of agents? How might these need to be adapted to support new security risks or paradigms introduced by AI agents?*

[1] *CrowdStrike's Fall 2025 Release Defines the Agentic SOC and Secures the AI Era*, October 1, 2025, https://www.crowdstrike.com/en-us/blog/crowdstrike-fall-2025-release-defines-agentic-soc-secures-ai-era/.

We appreciate NISTs efforts to support the development of standards in this rapidly evolving space.[2] 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 likely 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 (both atypical behavior as well as behavior that strays from the intent of the human involved); and 3) enforcing policies and guardrails across an organization.[3] Security teams maintain control over AI agents by having robust visibility, and by continuously discovering misconfigurations, shadow AI, and access.

C. *How do we bind agent identity with human identity to support "human-in-the-loop" authorizations?*

CrowdStrike leverages agentic AI to automate the most time-intensive aspects of incident response, while ensuring analysts retain control over final decisions and customers can define when and how automated actions occur.

As noted in Section B, granular identity controls are essential for agent deployments throughout the enterprise. An emerging requirement for AI security is real-time identity. As agent identities increase, organizations require a modern control plane for

---

[2] CrowdStrike, on March 9, 2026, submitted comments to NIST's *Request for Information Regarding Security Considerations for Artificial Intelligence Agents*, providing additional information about the security of AI agents to include an overview of the threat landscape and security best practices.
[3] *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/.

all identities that is powered by real-time risk assessment and continuous dynamic authorization. Security teams can define fewer, more adaptable policies for human and non-human identities that adjust privileges based on real-time risk and contextual data when leveraging risk-aware permissions. This real-time approach can also grant access to SaaS and cloud resources the moment it's needed, and revoke it the moment it's not - creating a modern approach to continuous identity that can respond rapidly to threats.

Traditionally, "human-in-the-loop" architectures imply the ability for a person to intervene in any process. For certain AI use cases, this may be an appropriate control to maintain. However, the emergence of agentic workflows presents another application of the concept regarding data. That is, if AI agents are trained on data with a high degree of human touch – for example, every triage, escalation, and remediation from our managed security services team over 10 years has trained CrowdStrike's security agents. Building from data that is rooted in human judgement deeply integrates and improves the resulting AI use case.[4]

D. *What controls help prevent both direct and indirect prompt injections?*

Prompt injection, to include both direct and indirect prompt injection, is a new security challenge unique to large language models (LLMs) and AI agents. CrowdStrike has analyzed over 300,000 adversarial prompts and tracks more than 150 prompt injection techniques.[5] The prompt layer must be monitored and defended like any other critical layer of the AI tech stack.

Defending against prompt injection attacks demands a multi-layered approach that addresses both technical controls and organizational processes to limit the size of the attack surface and detect and stop injection attacks. Organizations should implement an AIDR program that includes:[6]

---

[4] *Inside the Human-AI Feedback Loop Powering CrowdStrike's Agentic Security*, February 10, 2026, https://www.crowdstrike.com/en-us/blog/inside-the-human-ai-feedback-loop-powering-crowdstrike-agentic-security/.

[5] *Taxonomy of Prompt Injection Methods*, https://www.crowdstrike.com/en-us/resources/infographics/taxonomy-of-prompt-injection-methods/.

[6] *Indirect Prompt Injection Attacks: A Lurking Risk to AI Systems*, December 4, 2025, https://www.crowdstrike.com/en-us/blog/indirect-prompt-injection-attacks-hidden-ai-risks/.

- **Prompt Injection Detection**: Deploy specialized prompt injection detection systems capable of identifying and blocking malicious prompts, both direct and indirect.
- **Input Validation and Sanitization**: Implement robust filtering of AI system inputs and external data sources to limit the total addressable attack surface for indirect attack.
- **Content Security Policies**: Establish clear policies about what types of content AI systems can process and from which sources. Implement allowlisting for trusted data sources and treat external content with appropriate suspicion.
- **Privilege Separation**: AI tools that are enterprise-managed should have minimal access to sensitive data and limited capabilities to take actions. Separate the read and write permissions, and require explicit user confirmation for high-risk actions.
- **AI Use Monitoring and Access Control**: Shadow AI exacerbates the attack surface for indirect prompt injection. Deploy solutions to illuminate employee AI tool use, and enforce governance policy and access controls to prevent unauthorized AI tool use.
- **User Education**: Train employees to recognize risks associated with AI tool adoption, and establish clear policies about sanctioned versus unsanctioned AI applications.

E. *How can zero-trust principles be applied to agent authorization?*

The draft publication notes that one of the goals of this effort is to apply existing identity standards and best practices to AI agents. CrowdStrike agrees with this approach and there are several standard cybersecurity practices that also apply broadly to AI agents. 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 permissions across every stage of digital interaction, organizations can significantly reduce the risk of unauthorized access. Further, controls to revoke permissions when underlying conditions or risks change are increasingly important. Zero Trust principles should be applied across 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.

F. *Relevant standards and guidelines*

CrowdStrike supports NIST's approach of leveraging existing best practices, standards, and NIST's previous work to inform agent identity. We welcome the inclusion of System for Cross-domain Identity Management (SCIM), which provides an external layer of security and governance outside of the agent, OpenID Connect, and current OAuth standards (currently 2.0/2.1). OAuth 2.0/2.1 and OAuth extensions provide flexibility that can be tailored to different security identity architecture choices and provide guidance on fundamental security principles, as well as a consistent identity paradigm and audit logs.

CrowdStrike also encourages alignment with NIST's recent Request for Information about Securing AI Agent Systems.

G. *Agentic Architectures*

Broadly, the use case selected for this draft publication is a single human to a single agent. We believe that multi-agent authentication and authorization pose additional considerations and challenges. We understand that NIST may be treating this draft publication as an initial approach to the topic of AI and identity, but we believe NIST could improve future drafts by noting plans to address multi-agent use cases.

We welcome NIST's inclusion of a graphic depicting a standard *interactive* agent architecture (Figure 1). NIST should contemplate graphics for other architectures like those where agents are *delegated* (human user initiates a task and the AI agent works over a long period of time to complete the task); *autonomous* (human user gives an AI agent a charter, but the agent assigns itself work and runs indefinitely; and *chained* (AI agent invokes another AI agent).[7]

[7] See, for example, *The four MCP use cases: Summary of a discussion at IIW*, October 29, 2025, https://sgnl.ai/2025/10/the-four-mcp-use-cases-summary-of-a-discussion-at-iiw/.

III. CONCLUSION

NIST's effort on AI agent identity 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 guidance remains flexible while addressing core security principles. As NIST moves forward, we recommend continued engagement with stakeholders and harmonization with applicable AI, cybersecurity, privacy, and data protection existing NIST guidelines.

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
Chief Privacy and Policy Officer

Elizabeth Guillot
Senior Manager, Public Policy

Email: policy@crowdstrike.com

***