Wednesday, June 3

8:30 am – 8:50 am

Leveraging Wiz Security to Enable Adoption of Agentic AI for Defensive Cyber Operations

Chris Saunders

Director of Public Sector Sales Engineering

Wiz, Inc

Abstract:

As the DoW transitions toward Agentic AI to automate defensive cyber operations (DCO), the primary challenge shifts from model performance to architectural trust and governance. At Wiz, we're building a platform that directly supports DISA and DoW's vision for Agentic AI in DCO. Our approach delivers sophisticated automation, secure integration, and strong governance across all your security tools and data. Wiz offers a mature, Agentic AI-ready platform to orchestrate and automate DCO, ensuring interoperability, secure integrations with engrained Zero Trust (ZT) principles, and robust governance for autonomous actions in high-stakes environments. Wiz provides a comprehensive security and governance framework that enables DISA to deploy autonomous agents with confidence. By leveraging our industry-leading Cloud Native Application Protection Platform (CNAPP) and AI-Security Posture Management (AI-SPM), Wiz provides the "connective tissue" required to orchestrate agents across diverse security tools while enforcing ZT principles and autonomous guardrails. Wiz addresses the lack of mature standards around AI by providing complete visibility. Our Security Graph maps every AI agent, its underlying Large Language Model (LLM), privileges granted, and the specialized tools/APIs it invokes into a single pane of glass. This allows DISA to visualize the "blast radius" of an agentic workflow across disparate DCO tools. Our Wiz Integrations Network (WIN) creates an open ecosystem, bringing together over 140 security tools and data sources. This enables the ability to orchestrate workflows and automate remediation across your existing security stack (e.g., SIEM, EDR, and SOAR), ensuring that the built-in multi-agent orchestration remains transparent and auditable. In a Zero Trust Architecture (ZTA), an AI agent is a non-human identity (NHI). Wiz differentiates by integrating directly

with existing authentication and authorization (AuthN/AuthZ) solutions to analyze agent permissions. We identify "overprivileged agents" that have broader system access than their defensive mission requires, enabling DoW to enforce strict identity bindings and prevent lateral movement if an agent is compromised via prompt injection or logic manipulation. We further enforce secure integration by monitoring service accounts for least privilege access suggesting changes to overly privileged accounts, strengthening AuthN/AuthZ. Wiz integrates with identity security solutions like CyberArk and Linx Security for Just-In-Time (JIT) access, directly supporting a robust ZTQ. Our platform builds in multi-layered defense, continuously verifying and authorizing every interaction. Wiz offers an AI bill of materials (AI-BOM), which tracks the lineage of agents, including the Model Context Protocol (MCP) servers. This framework provides DISA and DoW with a standardized method to govern the supply chain, ensuring only approved, vetted agents are permitted to execute actions in high-stakes environments. Beyond static configuration, Wiz provides runtime AI security. We monitor agent behavior for deviations from intended operational workflows, such as an agent attempting to exfiltrate data or modify firewall rules outside of its defined mission parameters. This provides the "human-in-the-loop" governance necessary for high-stakes DCO, where autonomous actions must be reconciled against real-time security policies. This means critical automated remediation actions can be reviewed and approved by people through interactive messaging, ensuring transparency and accountability. Features like approval requests and workflow variable governance help restrict sensitive data flow, providing policy-driven oversight and avoiding "black box" automation worries. Our Blue AI Agent and Green Agent concepts use agentic principles to automate threat investigation, assign verdicts, and create dynamic remediation plans, with human validation.