CIS, Astrix, and Cequence Release New AI Security Companion Guides

Partnership delivers practical guidance for securing LLMs, agents, and MCP environments

CLIFTON PARK, NY, UNITED STATES, April 21, 2026 /EINPresswire.com/ -- The Center for Internet Security, Inc. (CIS®), Astrix Security, and Cequence Security today announced the release of three new CIS Critical Security Controls® (CIS Controls®) Companion Guides designed to help enterprises secure rapidly evolving AI environments.

Co□authored by experts across all three organizations, the guides extend the CIS Critical Security Controls into AI systems where large language models (LLMs), autonomous agents, and Model Context Protocol (MCP) integrations introduce new and unique risks. Each guide focuses on a distinct layer of the AI ecosystem, offering targeted guidance aligned with how modern AI systems operate:

- AI LLM Companion Guide: Provides guidance for securing large language models, including risks related to prompts, context handling, and exposure of sensitive information.

- AI Agent Companion Guide: Outlines controls for managing autonomous and semi-autonomous agents, focusing on safe tool execution, governed autonomy, and appropriate access to enterprise systems.

- MCP Companion Guide: Details protections for Model Context Protocol environments, emphasizing secure tool access, management of Non-Human Identities (NHIs), and auditable interactions across the protocol layer.

Download the guides

As AI becomes deeply embedded in production workflows – from copilots to autonomous task execution to tool-integrated systems – security teams are confronting risks that traditional controls were never built to address. These include data leakage, unbounded agent autonomy, credential misuse, and unsafe or inappropriate execution of tools. The new Companion Guides offer practical, prioritized guidance that reflects how AI is actually deployed in modern enterprises.

“These guides reflect a shared effort to bring clarity to an area where organizations are seeking direction,” said Curtis Dukes, Executive Vice President and General Manager of Security Best Practices at CIS. “By combining our collective expertise, we translated the CIS Controls into concrete steps that help teams secure AI systems across the model, agent, and protocol layers.”

Astrix contributed deep expertise in securing AI agents, MCP servers, and NHIs, including API keys, service accounts, and OAuth tokens that connect AI systems to enterprise resources.

“AI agents introduce a new operational surface that organizations must understand before they scale,” said Jonathan Sander, Field CTO of Astrix Security. “Collaborating with CIS and Cequence allowed us to build guidance that addresses identity, authorization, and execution risks in a way that’s both actionable and aligned with how enterprises work today.”

Cequence brought extensive experience in securing enterprise applications, data, and APIs, shaping guidance around visibility, governance, and control over what AI systems can access and execute.

"As AI systems interact more directly with applications and APIs, the security implications become increasingly critical," said Shreyans Mehta, CTO and Co-Founder of Cequence Security. "This partnership enabled us to create guidelines that codify what we've learned about deploying agentic AI at the world's largest enterprises without sacrificing security, governance, or scale, giving organizations a framework for enabling agentic AI safely."

How the Companion Guides Support Organizations

Together, the three Companion Guides give security and IT teams a unified way to apply the CIS Controls to AI systems that behave and evolve differently from traditional software. By extending the Controls into environments powered by LLMs, autonomous agents, and MCP-based integrations, the guidance helps organizations understand where risks emerge and how to address them with guidance that reflects real-world deployment patterns.

The guides:

- Adapt the CIS Controls to AI-driven architectures, helping teams secure LLMs, agentic systems, and MCP interfaces without adopting a new framework.

• Provide clear, prioritized recommendations that support responsible AI adoption across development, deployment, and operational phases.

• Blend the strengths of all three organizations by combining standards leadership with deep expertise in agentic AI and API-centric security.

• Cover the full AI security stack, from model inputs and context handling to agent reasoning, tool execution, and protocol-level access.

Join CIS, Astrix, and Cequence on May 13 at 1:00 p.m. ET for From Prompts to Protocols: The Security Blueprint for Enterprise AI. The teams will highlight key insights and offer guidance for security teams, developers, and AI practitioners.

Register Now

For more information about the partnership and the guides, visit cisecurity.org.

###

About CIS
The Center for Internet Security, Inc. (CIS) makes the connected world a safer place for people, businesses, and governments through our core competencies of collaboration and innovation. We are a community-driven nonprofit, responsible for the CIS Critical Security Controls® and CIS Benchmarks® guidelines, globally recognized best practices for securing IT systems and data. We lead a global community of IT professionals to continuously evolve these standards and provide products and services to proactively safeguard against emerging threats. Our CIS Hardened Images® provide secure, on-demand, scalable computing environments in the cloud. CIS is home to the Multi-State Information Sharing and Analysis Center® (MS-ISAC®) organization, the trusted resource for cyber threat prevention, protection, response, and recovery for U.S. State, Local, Tribal, and Territorial government entities, and the Elections Infrastructure Information Sharing and Analysis Center® (EI-ISAC®) organization, which supports the rapidly changing cybersecurity needs of U.S. election offices. To learn more, visit cisecurity.org or follow us on X: @CISecurity.

About Astrix Security
Astrix secures the full lifecycle of AI agents and the Non-Human Identities (NHIs) that power them, extending traditional IAM to govern the modern AI attack surface. While agents and other NHIs outnumber humans 100:1, they remain under the radar, creating the biggest blind spot in our identity perimeter. Astrix provides a unified solution for continuous discovery of all AI agents and NHIs, security and remediation of excessive privileges, protection against real-time threats, and responsible adoption of new agents with 'secure by design' guardrails like agentic just-in-time access. This enables enterprises to adopt AI responsibly while accelerating productivity.

Astrix is trusted by leading organizations including Workday, NetApp, Priceline, Figma, HubSpot, Workato, and many more. To learn more, visit https://astrix.security/.

About Cequence Security
Cequence protects the applications and data that power enterprises in the agentic era. More than a decade of bot defense and API security experience has established Cequence as the leader of safe and secure agentic AI adoption. The Cequence platform delivers deep insight into user, entity, and agent behavior, enabling organizations to secure and control agentic AI interactions while protecting against bad actors and rogue agents. Cequence delivers value in minutes rather than days or weeks with a highly scalable, no-code approach. Trusted by the largest and most demanding private and public sector organizations, Cequence protects more than 10 billion daily API interactions and 4 billion user accounts. To learn more, visit www.cequence.ai.

Kelly Wyland
Center for Internet Security
email us here
Visit us on social media:
LinkedIn
Instagram
Facebook
YouTube
X