Netskope Enables Security Teams to Stop Risky AI Agent Actions Before They Execute

September 15, 2026

Part of the newly renamed Netskope Skylight AI Security suite, Agent Action Control classifies agentic actions by risk and blocks the ones that matter before they execute

SANTA CLARA, Calif., Sept. 15, 2026 (GLOBE NEWSWIRE) -- Netskope (NASDAQ: NTSK), a leader in modern security and networking for the cloud and AI era, today announced Netskope Skylight Agent Action Control, a new capability that classifies every action an AI agent attempts, and applies granular controls to stop high-risk actions before they execute. The capability meets security teams' need for a policy-based way to govern, rather than react to, what autonomous agents are permitted to do.

Agent Action Control arrives as Netskope renames its AI security portfolio to better communicate the company's broad capabilities across AI security. Formerly Netskope One AI Security, Netskope Skylight is the new family name for the suite of Netskope AI security products including Netskope Skylight AI Command Center, Netskope Skylight AI Guardrails, Netskope Skylight AI Gateway, Netskope Skylight Agentic Broker, and Netskope Skylight AI Red Teaming.

Netskope Skylight Agent Action Control

Today, 91% of organizations cannot stop a risky agent action before it executes[1], and 54% reported a confirmed or suspected AI agent security incident in the past year[2]. With Netskope Skylight Agent Action Control, teams can more effectively govern what an agent is attempting to do thanks to:

- **Control over nine different types of action:** To enable granular policy definition, every agent action is classified into one of nine intent-based categories prior to execution. These include, access control changes, configuration changes, credential and secret manipulation, data destruction, infrastructure provisioning, potential data exfiltration, potential external communication, remote code execution, or source code change.
- **Risk-based policy profiles by agent type:** Block, allow or alert based on low, medium, high or critical risk categories, with the option to send a default or custom notification to the end-user for user coaching when the action is blocked. Profiles attach to specific agents, so a coding assistant and a chat application do not have to operate under the same rules.
- **Granular policy alerts:** Every action is logged, with security teams able to filter policy alerts by cost exposure, source code changes, infrastructure updates, or external communication to align investigation effort with organization risk prioritization.
- **No new console and no new agent:** Netskope Skylight Agent Action Control runs on network traffic the Netskope platform already inspects, so security teams gain enforcement over agent behavior without deploying a separate tool or standing up a second console.

The result is a policy that governs what an agent is allowed to do before it acts. When a coding agent working from a vague prompt attempts to delete a production repository, for example, the action is classified as data destruction at a critical risk level, and the call is stopped before it reaches the repository. The security team gets a record of the attempt instead of an incident report. This approach provides governance and controls over agents at risk of "authority drift", or whose harness or instructions are too vague or permissive.

> *"AI agents tend to act first and explain later, and most security teams only learn what happened after it is done," said John Martin, Chief Product Officer, Netskope. "Agent Action Control puts a decision in front of every action an agent takes, so a team can say yes to agentic AI without saying yes to the one action that could cost them a production system."*
>
> *"As enterprises accelerate adoption of AI agents, we're finding that probabilistic controls are sometimes insufficient to protect the enterprise, but even occasional failure is unacceptable," said Dr. Grace Trinidad, Research Director for AI Security and Trust at IDC. "These hardened, policy-based, deterministic controls are the backstop that prevents AI agents from causing an enterprise incident."*

Netskope Skylight Agent Action Control will be available at the end of the current quarter. To learn more about Netskope Skylight, visit *netskope.com/skylight*.

Head over to The Lens (Netskope's blog) to read more about Skylight, or to take a deeper look at Netskope Skylight Agent Action Control.

About Netskope
Netskope (NASDAQ: NTSK), a leader in modern security and networking for the cloud and AI era, addresses the needs of both security and networking teams by providing optimized access and real-time, context-based security for the AI ecosystem inclusive

of agents, applications, tools, LLMs, people, devices, and data. Thousands of customers, including more than 30% of the Fortune 100, trust the Netskope One platform, its Zero Trust Engine, and its powerful NewEdge network to reduce risk and gain full visibility and control over cloud, AI, SaaS, web, and private applications — providing security and accelerating performance without trade-offs.

Learn more at netskope.com, on LinkedIn, and on Instagram.

Media Relations Contacts:
press@netskope.com

Investor Relations Contacts:
ir@netskope.com

---

[1] Netskope, 2026 AI Risk and Readiness Report.
[2] Gravitee, The State of AI Agent Security 2026