

# Netskope Skylight Agent Action Control

Data Sheet



*Control agent actions. Scale AI securely*

## The challenge

Effective AI strategies involve agents, but most act without oversight, calling tools and moving data unsupervised. Nearly all organizations - 91% - can't stop a risky action before it happens.<sup>1</sup> Almost the same proportion (88%) have had a confirmed or suspected incident.<sup>2</sup>

## The solution

Netskope Skylight Agent Action Control lets you scale agents with confidence. It classifies agent actions into nine risk categories - from data destruction to remote code execution - then blocks risky actions or requires human approval before execution. It runs on traffic Netskope already inspects, so you scale AI without managing new tools.

### How it works

- **Identify:** Agent Action Control recognizes known AI agents, giving visibility into agents performing unintended actions.
- **Classify:** Attempted actions are classified into one of nine intent-based categories, including data destruction and infrastructure provisioning, and assigned a risk level based on what it touches.
- **Set policy:** Security teams build policy that pair each category and risk level with a decision: block or user alert, then apply that policy to specific agents.
- **Enforce controls:** Netskope One NG-SWG carries out enforcement on API traffic, stopping the action before it ever reaches its target.

## Strategic advantages

### Every action is classified.

Every attempted action sorts into one of nine intent-based categories covering access control, configuration, credentials, data destruction, infrastructure, exfiltration, external communication, code execution, and source code, enabling granular policy definition.

### Simple policy profiles by agent type.

Risk profiles are attached to the intent-based categories, and be applied to one, or multiple agents, so a coding assistant and a personal agent do not need to share the same rules.

### No new tool to deploy.

Agent Action Control runs on traffic already inspected by NG-SWG, with no additional console to manage.


### Full activity logging by default.

Every classified action is logged and available for review, giving security teams a clean record, without needing a separate console.

**“By 2028, 25% of enterprise breaches will be traced back to AI agent abuse, from both external and malicious internal actors.”**

– Gartner, Top Strategic Predictions for 2025 and Beyond, 2025

## Netskope Skylight Agent Action Control



1. 91% figure source: [Netskope, 2026 AI Risk and Readiness Report](#)

2. 88% figure source: Agat, AI Agent Security in 2026: What Enterprises Are Getting Wrong

## The Netskope difference

Providing one control plane for every agent action, Agent Action Control runs as a single sensor inside Netskope Skylight, the family of AI security products built on the Netskope One platform. Netskope Skylight follows one approach across the AI stack: discover what is running, govern what it is allowed to do, and secure it at the point of use. Agent Action Control covers the govern and secure steps for agent activity across agentic traffic, classifying every action an agent attempts and enforcing the policy a security team sets for it.

That classification runs on network traffic already inspected by NG-SWG, so the same traffic carrying data loss prevention and threat protection also carries agent action classification.

This means a security team evaluating agent risk turns on a capability inside infrastructure it already runs instead of standing up a new console or tool. The team gets a view of agent activity in the same place it already sees AI app use, model traffic, and data exposure. A control plane built this way grows with how an organization uses AI: agent by agent, action by action, and product by product, as Netskope Skylight adds new sensors over time.

| YOUR NEEDS                                | THE NETSKOPE SOLUTION                                                                                                                     |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| One inventory of agent activity.          | See every agent action across API traffic in one view, instead of piecing together partial signals from disconnected tools.               |
| Intent-based policy, not a binary choice. | Set block, allow, or alert per action category and risk level. Teams needn't choose between blocking every agent or trusting every agent. |
| Tailor policies for agents for use cases. | Apply a stricter profile to agents coding use cases, and a lighter one to personal agents, instead of one rule for all.                   |
| Intent-based policy, not a binary choice. | Set block, allow, or alert per action category and risk level. Teams needn't choose between blocking every agent or trusting every agent. |
| No new console to manage.                 | Classification runs on NG-SWG traffic already in place, so there is nothing new to deploy.                                                |
| Nine defined intent-based categories.     | Cover the range of what an agent can attempt, from data destruction and credential manipulation to remote code execution.                 |
| A clean record when action                | Get proof that an attempted action never executed, instead of a forensic investigation after the damage is done.                          |



Interested in learning more?

Request a demo

Netskope, a leader in modern security and networking, addresses the needs of both security and networking teams by providing optimized access and real-time, context-based security for people, devices, and data anywhere they go. 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.](#)

©2026 Netskope, Inc. All rights reserved. Netskope, NewEdge, SkopeAI, and the stylized “N” logo are registered trademarks of Netskope, Inc. Netskope Active, Netskope Cloud XD, Netskope Discovery, Cloud Confidence Index, and SkopeSights are trademarks of Netskope, Inc. All other trademarks included are trademarks of their respective owners.