

Netskope AI Risk AI SecOps Agent

Data Sheet




Detect AI risk patterns that signal future disruption

The challenge

Individual AI activities can bypass alert thresholds and accumulate into significant business risk. Identifying which users, AI applications, agents, and MCP servers are accumulating risk becomes increasingly difficult through manual investigations as AI usage grows.

The solution

The Netskope AI Risk AI SecOps Agent analyzes telemetry from Netskope Skylight AI Security sensors across users, AI applications, agents, and MCP servers to identify risky patterns hidden within individual events. It prioritizes high-risk entities, investigates them automatically, and prepares cases with evidence, verdicts, confidence levels, and recommended actions.



Strategic advantages

- **Approve AI adoption with evidence:** Demonstrate to the board that business risk can emerge from patterns of AI activity that individual events may not reveal. When scaling AI adoption, leadership can rely on evidence-backed investigations.
- **Extend governance beyond initial approval:** Reassess approved AI applications and MCP servers daily based on behavioral changes and their security and compliance posture. Approved AI entities are continually evaluated as usage patterns, relationships, and risk exposure change.
- **Create a shared view of AI risk:** Send investigation verdicts to the Netskope AI Command Center as risk badges. Governance and SOC teams get a shared, evidence-backed view of AI risk, supporting coordinated and timely response decisions.

How it works

- **Connect the signals**


The agent correlates signals from Netskope Skylight AI Security, making connections across users, AI applications, and MCP servers. These inputs include data exposure from Netskope DLP, prompt violations from AI Guardrails, user trust from Advanced UEBA, app reputation from NG-SWG and CASB, MCP activity from Agentic Broker, and AI inventory data.

- **Detect risky patterns**

The agent evaluates activity daily across users, AI applications, and MCP servers to identify patterns that warrant investigation. Repeated violations, behavioral changes, persistence, and potential business impact are prioritized into a single queue. Only the highest-risk entities advance to investigation, while others are retained for monitoring and review.

- **Act with evidence**

The agent investigates prioritized entities automatically, assembles an evidence timeline, and returns a verdict of risky, legitimate, or inconclusive with a confidence level. Investigation outcomes appear as risk badges in Netskope AI Command Center. Analysts review the evidence and decide the appropriate response using existing controls such as trust score policies, watchlists, and inline protection.



Capabilities

Detect hidden risks in weak signals: Identify patterns, relationships, and repeated activity that indicate growing risk. Using this insight, security teams can intervene earlier, before risky activity escalates into a business disruptive incident.

Focus on high-risk behavioral change: Assess AI entities against their own behavior patterns rather than uniform thresholds, so analysts can identify deviation in risk patterns with accurate context and reduce time spent on low-value investigations.

Identify risk that emerges after AI approval: Reassess approved AI applications, agents, and MCP servers as behavior and exposure change, so organizations can detect risk that emerges after initial approval.

Distinguish productive AI use from accumulating risk: Evaluate activity against established behavior patterns, so analysts can prioritize critical changes and allow legitimate AI use to continue without unnecessary scrutiny.

Align governance and security teams on AI risk: Consolidate evidence, findings, and verdicts into a single risk view in the Netskope AI Command Center, so governance and SOC teams can prioritize and act with consistent context.



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.](#)