Netskope Insider Threat AISecOps Agent

Data Sheet

Daily insider risk monitoring with defensible evidence

The challenge

Evidence of insider risk exists but is scattered; analysts correlate timelines manually and every review starts from scratch. Manual investigations take days and remain incomplete, even after the risky user has left the organization with data they should not possess.

The solution

The Netskope Insider Threat AISecOps Agent runs a daily assessment of every watchlist user and calculates a combined risk score from behavior, data loss prevention (DLP), malware, access, and application signals. When a user's score demands attention, the agent surfaces a prioritized, evidence-backed insight. The analyst opens a case; the agent investigates it automatically and returns a verdict with recommended actions in minutes.

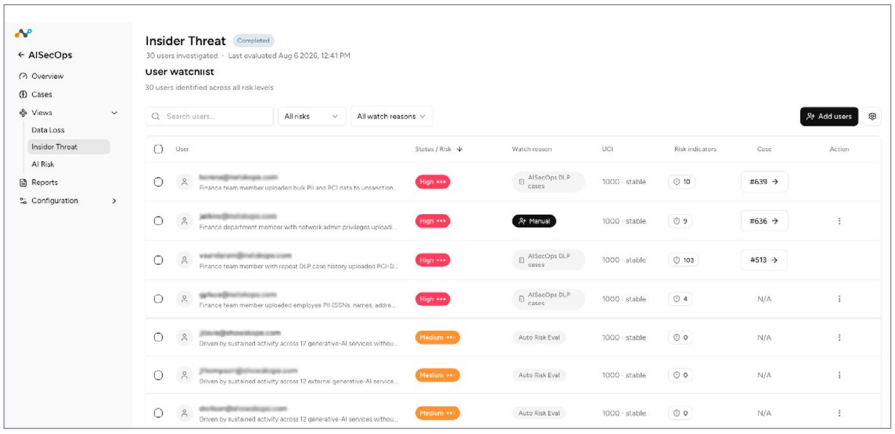

Netskope daily assessment dashboard

Strategic advantages

- **Human in the loop:** The agent surfaces daily insights; the analyst reviews the risk-prioritized list and decides which insights become cases. The agent never opens a case on its own.
- **A defined watchlist, not a company-wide watch:** The agent watches only the insiders the organization defines which keeps the program acceptable to employees, HR, and legal.
- **Coverage beyond HR lists:** Contractors and privileged third parties rarely show as flight risks on HR lists. Designated IDP groups put them under the same daily watch as employees, and the list maintains itself.

netskope

How it works

- **Monitor: inline visibility**
  Netskope operates inline, observing file moves, application use, policy triggers and behavioral signals generated across networks, applications, endpoints and AI. It does this, across web, application, network, endpoint, and AI. The scope includes only defined watchlists.
- **Assess: 24-hour assessment engine**
  Once a day, the agent scores each watched user across user confidence index (UCI), behavioral anomalies, DLP incidents, malware, access changes, and application activity. Notable scores surface an insight with evidence for and against risk; others log silently.
- **Act: automated investigative sweep**
  The analyst opens a case: the agent sweeps user and entity behavior analytics (UEBA), UCI, DLP, malware, and application signals, plus endpoint detection and response (EDR) and IDP where configured, and delivers verdicts, such as, risk, inconclusive, or legitimate, with evidence and recommended actions in minutes.

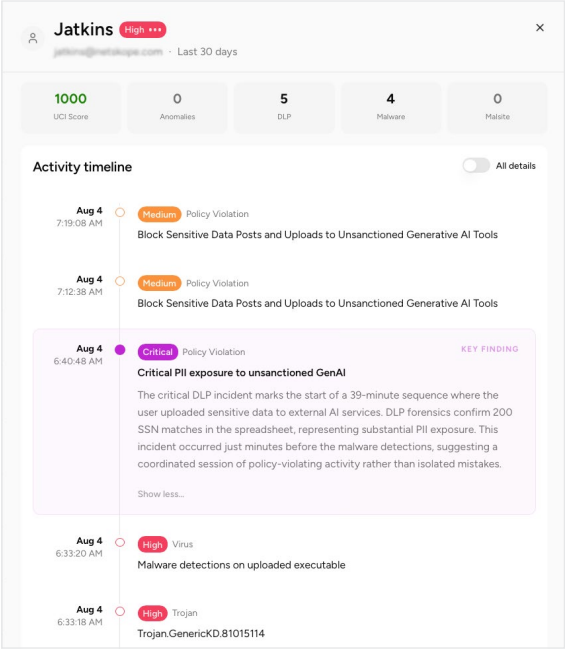

Capabilities

Combined daily risk score: Every 24 hours, the agent combines multiple signals and activity logs into a single risk score per watched user. One number per person replaces dozens of unconnected alerts.

Prioritized daily insights: The analyst starts each day with a list of watched users ranked by risk. Each entry shows evidence and context, so analysts can build a full picture.

Automatic case investigation: When the analyst opens a case, the agent re-examines the user's full history, pulls in extra context, and returns a verdict with recommended actions in minutes.

Evidence traceability: Every claim links back to exact data points. The analyst can check how the agent reached its conclusion, and the organization can defend that decision.

Shared case queue and workflow reach: Cases land in the DLP queue, with the same status, assignment, and RBAC. No new console to learn and results flow to SIEM and ITSM.