

2026 REPORT

State of AI Agents in Security and Network Operations

From Alert Overload to Trusted Agentic Operations



Executive Summary

Security and network operations teams now detect incidents faster than they can resolve them, and manual effort alone no longer keeps pace with growing workloads. Investigations compete for scarce analyst time, and important cases go untouched. Behind that unfinished work are possible intrusions, unresolved service faults, and exposure that lasts longer than it should. Increasing operational capacity has become a security priority.

AI agents expand the work teams can delegate. While traditional scripted automation follows predefined paths, an agent can decide what to check next as evidence changes. It can investigate an alert, assemble the findings, and carry out a response within its approved authority. Analysts can supervise the work and handle consequential decisions without directing every intermediate step. This is the core value of agents in security and network operations: to complete work that would otherwise consume valuable analyst time or remain in the queue.

However, the same ability to reason and act quickly creates risk. An agent that misreads evidence can carry out the wrong response if its permissions allow it. Teams need to define what it may change, when human approval is required, and who can intervene. Delegating work does not transfer accountability.

Based on a survey of 921 security, network, and IT professionals, this report shows where agents can relieve operational pressure, what tasks teams are prepared to delegate, and the controls needed to make that delegation work.

Key Findings:

- **42% of organizations let alerts go uninvestigated**
Workloads have increased for 81% of teams, yet only 11% consistently keep up. As a result, 42% see alerts, cases, or tickets go uninvestigated. For 24%, sensitive data, systems, or users remain exposed longer than necessary.
- **Detection is improving faster than resolution**
60% of teams have become more effective at detecting threats and network issues, compared with 49% who report more effective response and resolution. For incidents requiring more than one team to investigate, 31% report containment times of at least a day.
- **Context and action are fragmented across tools**
Only 9% describe their tools as highly integrated, and 51% cite scattered signals and context as a leading obstacle to faster resolution. Analysts must pull those records together manually before they can decide what happened and how to respond.
- **Teams want agents to increase their capacity**
56% cite capacity pressure as a reason to consider agents; executive or board expectations rank last at 18%. For the coming year, 47% plan to adopt agents for specific workflows, compared with 32% planning to add staff. Triage offers a practical starting point because agents can organize and prioritize cases for analyst review without changing operational systems.
- **Teams grant agent authority in steps**
68% would not let agents make changes on their own: 38% rule out agent-executed changes entirely, and 30% require approval first. 23% would accept low-risk automatic actions within guardrails.
- **Agents need shared context and consistent rules**
Asked what an agent platform must provide, 55% put shared context first and 51% said open integration with existing tools. 46% want one governance model that covers vendor-built, internally built, and managed-service agents alike. Each new agent should build on the same case record rather than open another history to reconcile.

Teams want agents to take on more work and expand capacity, but most would not let them make operational changes on their own. Agents already run in production for 29% of teams, yet only 9% are very confident they can govern them. The challenge is to delegate enough work to relieve capacity pressure while keeping consequential actions under control. The following pages examine where agents can deliver that value, how to define their authority, and how to measure reduction of the team's workload.

Workload has grown faster than the team

For many operations center teams, the morning shift begins with a queue that holds more work than they will clear that day. Analysts choose what can wait before they have even gathered enough evidence to know which alert, case, or ticket carries the greatest risk.


Workload has increased for 81% of security and network operations teams over the past year and only 11% consistently keep up. The resulting pressure is leading to 45% of respondents reporting analyst burnout, attrition, or unfilled roles. When experienced people leave, teams lose knowledge they rely on to triage, investigate, and resolve problems.

Teams point to three main sources of pressure: 41% cite higher alert, case, or ticket volume, 34% name staffing constraints, and 30% report false positives, duplicates, or low-fidelity alerts. Much work remains manual: 61% estimate that triage, correlation, documentation, and reporting consume at least 40% of team time. Analysts sort duplicates, pull evidence from identity, device, and application systems, update tickets, and find owners before an investigation can advance. Repeating those steps across every case leaves the team’s most experienced people doing its lowest-value work instead of confirming incidents, containing exposure, or restoring a failing service.

The result: 42% see alerts, cases, or tickets go uninvestigated, 40% handle incidents more slowly than their targets, and 24% report that sensitive data, systems, or users remain exposed longer than necessary. An overlooked alert is a possible intrusion nobody has ruled out, while a delayed application case can become the next production outage.

Unresolved work puts security at risk

► In the past 12 months, which of the following has your team experienced as a result of operational capacity or tooling limits?



*Select up to three

Teams that treat capacity as operational risk track unresolved and uninvestigated work alongside completed cases. They identify where analysts spend time on repeated tasks, then test whether automation or an agent can reduce that effort without lowering the quality of investigation. Measure the result before expanding its role.

Where agents can help investigations

Detection has been improving as security and network tools surface problems more quickly. The response work, however, still moves at the speed of the slowest manual step: gathering evidence, finding the owner, getting approval to act, applying the fix, and confirming that the remediating action worked.


60% of teams say detection effectiveness has improved over the past two years, compared with 52% for investigation, 49% for response and resolution, and 44% for preventing recurrence. For incidents requiring cross-team investigation, 31% of organizations take at least a day to contain.

Scripted automation has been introduced to reduce manual workload for common, repeatable issues. It works best when the trigger, steps, and response are known in advance. While a playbook can enrich a familiar alert, open a ticket, and route it to a fixed owner, an investigation often changes paths as evidence arrives. A suspicious login may first require identity context, then network or data-transfer evidence, and finally a decision about whether to restrict access. Each answer determines the next source, owner, or action.

An agent, on the other hand, can treat the alert as a starting hypothesis, weigh what it finds, choose the next evidence source or tool, and then update the case as results arrive. It follows the evidence without requiring every investigative path to be coded in advance, and it can show the analyst which evidence drove each step. However, this same discretion is what makes the agent's authority difficult to govern.

Detection is improving faster than resolution

► Over the past two years, how has your team's overall effectiveness improved in each of these areas?



31% report containment takes at least a day for incidents requiring more than one team to investigate.

To close the resolution gap, teams need integrated case records that include evidence, owner, approval status, next action, and outcome from detection through verification. Test an agent's ability to gather evidence, update the case, and recommend a response so that analysts can apply the necessary expert judgment.

One event becomes several disconnected cases

Different tools can record the same event under different names and owners. Each may hold part of the evidence under a different identifier or owner across security, network, cloud, identity, and ticketing systems. Only 9% of teams describe their security and networking tools as highly integrated. Analysts must often connect those records themselves before they can establish what happened and who should respond.


Consider this scenario: Users report a cloud application slowing down. The network operations center (NOC) sees a degraded network path; the security operations center (SOC) sees unusual outbound traffic from the same workload. The identity system holds the service account details, while the cloud console records the data movement. A data loss prevention (DLP) tool flags the same transfer under a separate ticket. Each console can be correct individually, but none shows the whole event. The NOC treats the symptom as a performance problem while the SOC investigates a possible intrusion. Until an analyst connects the records, no one can reliably determine priority, whether the problems are related, and which response is appropriate.

The survey shows the scale of the problem: 51% name scattered context and signals as a leading obstacle to faster resolution, ahead of insufficient evidence at 38% and manual handoffs or approvals at 35%. 46% spend more than a quarter of their time moving information between tools, dashboards, tickets, and reports. Only 9% report effective correlation across users, devices, applications, and network paths to find the root cause.

An agent with access to those records can group the related alerts into one case. It can connect the identity, workload, and transfer evidence and prepare a recommendation. The case should show the sources, identify the response owner, and flag missing or conflicting information.

Scattered context slows resolution

► What most prevents your team from resolving security or network operations issues quickly and consistently?



*Select up to three

Only 9%
call their tools
highly integrated,
and 46% spend
more than a quarter
of their working
time moving
information between
tools, dashboards,
tickets, and reports.

Test whether the agent reduces the time analysts spend assembling and checking evidence. If they must reconstruct its work before trusting it, the expected capacity gain may disappear.

Operational pressure is driving agent adoption

Capacity pressure is the leading driver of interest in agents at 56%, followed by threats moving faster than manual response at 46%, and difficulty hiring or retaining skilled staff at 35%. Despite board-level interest in AI innovations, executive or board expectations rank last at 18%.

Teams have stopped treating headcount as the first answer to case volume. Asked how they will address capacity over the coming year, 47% expect to adopt AI agents for high-friction workflows with the most manual steps. That is ahead of improving integration across existing tools at 44% and adding headcount at 32%. These plans suggest teams are looking for ways to gain capacity by deploying agents and better connecting their existing tools.

Adoption is still early: 53% are evaluating or piloting agents, while 29% have reached production, including 7% who have scaled agents across multiple workflows.

When asked to assess the most critical operational functions considered for agentic automation, triage is the workflow where current manual effort, comfort with agent support, and expected value align most closely: 46% identify it as a leading source of manual effort, 46% are comfortable handing triage to an agent, and 48% expect it to deliver significant value. Remediation runs a step behind: 29% expect real value from it, and 20% are already comfortable letting an agent act. This makes triage a stronger candidate for an initial pilot than allowing an agent to change systems.

Triage has the strongest case for agent automation

	Manual effort today	Comfort with delegation to agent	Expected value
Alert triage and prioritization	46%	46%	48%
Cyberattack investigation and case correlation	39%	35%	41%
Case documentation and reporting	32%	32%	22%
Incident remediation and response actions	28%	20%	29%

56% say capacity pressure drives their interest in agents.

18% say executive or board expectations, the lowest-ranked driver.

A successful pilot should remove a defined unit of work so the analyst stops doing it. Teams should measure manual steps removed, analyst time reduced, issues investigated, and resolution time improved against a documented baseline, then expand scope only when the agent proves safe and useful.


Permissions determine what an agent can change

An agent can investigate a case without being allowed to change the systems involved. The permissions attached to its service identity or API token determine whether it can disable an account or isolate a workload. The discretion that makes an agent useful in triage can become operational risk the moment the agent has permission to act.

Teams carefully assess agent authority before granting it: 38% would not yet accept agents executing operational changes. Another 30% require explicit human approval first. 23% already let agents complete low-risk actions within defined guardrails, and 5% allow broader autonomy. A practical place to draw the line is between actions a team can reverse and actions it cannot.

Return to the cloud-workload scenario. A response agent treats the unusual outbound transfer as data theft and isolates the workload. Its identity permits the change, and no human approval is required. However, the workload supports a customer-facing service, which now fails. A plausible containment step became an outage because the agent had more authority than the workflow could safely support. The API enforced the permissions correctly, but the process lacked a check on the business impact before the agent acted.

How much authority would teams allow AI agents?



TOP CONCERNS ABOUT AGENTS THAT CAN ACT

48% Unsafe or unintended actions. **40%** Excessive access and permissions. **35%** Acting on faulty or manipulated information.

Give each agent a distinct identity and permissions tied to its assigned work. Enforce those limits at the tool or service level, and require approval for high-impact changes. Before expanding authority, demonstrate that the team can stop the agent, revoke access, and recover from mistakes.

Agents need clear ownership

Security and network operations teams have a clear understanding of what controls they need before trusting an agent to act: 54% prioritize human oversight and approval, followed by least privilege and separation of duties at 49%. Transparency and monitoring follow (45%), then the ability to pause, restrict, or roll back an action (42%).


Yet only 9% are very confident that controls and audit trails are applied consistently in their organization. 45% have limited or no confidence, even as 29% report agents already running in production within the SOC and NOC.

The survey results do not point to one single model for governance ownership. 26% share governance across functions, while 21% assign it to SOC leadership. No other answer exceeds 12%. Shared governance can work, provided each decision has a named owner. That includes approving exceptions and accepting risk, as well as stopping the agent and leading recovery.

Return to the isolated cloud workload scenario. The SOC has the agent's decision record and transfer evidence. The cloud team owns the workload and sees the isolation in its change log. The NOC sees the service outage and opens another ticket. Each team can explain its part, but no one is responsible for deciding whether the transfer was malicious or when service can safely resume. The workload stays isolated while three teams compare notes, with no shared case record or named recovery owner.

Wide agreement on AI controls ... low confidence on enforcement

► Which controls would be most important before allowing AI agents to support security or network operations workflows?



*Select up to three

► How confident are you that your organization can govern AI agents used in security and network operations?

9% are very confident in governing AI agents.

45% have low or no confidence they can govern AI agents.

Mature programs assign a named owner to every agent and define who approves exceptions, can suspend operation, accepts risk, and leads recovery. They preserve the evidence, authorization, action, and outcome in one record, then test containment and recovery before expanding authority.

Agents must work together

Teams need clear ownership of agent actions, but they also need a record of how those actions fit together. A single investigation may involve multiple agents built by different teams or providers, but their records need to connect as the work moves between them.

For example, a vendor-supplied agent records the security alert it triaged. An internally developed agent records the access change it made. A managed-service agent logs the ticket it updated. Each agent records only its own work, but unless those records connect, an analyst must manually reconstruct the sequence, repeating much of the work the agents were introduced to reduce.

Mixed agent estates will be common. A hybrid model combining vendor-provided, internally developed, and managed-service agents is the single most preferred approach at 28%, ahead of any one source on its own. A platform can reduce that burden by running its native agents on one governed foundation, applying the same security and audit controls to each, and extending those controls to third-party integrations with other vendors' data.


Teams have clear preferences for that platform: 55% want agents working from shared context, followed by open integration with existing tools (51%). Another 46% want one governance and policy model across agents, regardless of who built them.

The forensic value of linked records became clear during the July 2026 Hugging Face* incident. An autonomous OpenAI agent escaped containment and compromised parts of Hugging Face's production infrastructure. Hugging Face reconstructed approximately 17,600 actions by mapping recovered agent logs to shell commands and correlating them with its own platform records. No single log contained the full sequence.

*Source: Hugging Face, "[Anatomy of a Frontier Lab Agent Intrusion: A Technical Timeline of the July 2026 Incident](#)," July 27, 2026.

Teams want agents connected

► What platform capabilities would matter most to keep AI agents from creating another layer of operational fragmentation?



48%
expect agents
to replace some
existing tools
and supplement
others.

*Select up to three

A well-designed program maintains one shared record across all agents, then links each agent's tool calls and actions to the same case ID and timeline. Preserve original logs, with links to the evidence and approvals behind each action. The shared case should show what changed and whether the intended outcome was achieved.

How ready are your agent workflows?

Before agents take on more work, teams need to know which workflows are ready. This matrix translates the report's findings into six capabilities and three maturity levels. Match your current workflow to a description in each row, then prioritize the gaps to close next before expanding agent use.

CAPABILITY	REACTIVE	MANAGED	ADAPTIVE
Strategy & Value Common gap: Pilots advance without clear proof of value	Pilots follow visible issues; owners, baselines, expected outcomes, and stop-or-expand criteria vary by team.	Priority workflows have accountable owners, baseline measures, target outcomes, risk limits, and expansion criteria.	Teams add agents to workflows that show measurable gains and stop those that do not. Broader use requires both operating results and safety evidence.
Visibility & Trusted Context Common gap: Scattered, untrusted case context and no agent inventory	No inventory of the agents, connections, and non-human identities already operating. Analysts rebuild cases across consoles; agents inherit whichever identity, asset, policy, or evidence a tool exposes.	Agents, connections, and identities are discovered and tracked. Priority workflows connect verified identities, owners, policy state, evidence, and action history from defined sources.	Context stays connected and current, with agent inventory, context integrity, and risk reevaluated before consequential decisions or actions.
Workflow Orchestration Common gap: Work stalls across tool and team handoffs	Cases move through copied evidence, tickets, chat, and ad hoc approvals; automation stops at boundaries.	One case record preserves evidence, ownership, approvals, agent and tool activity, action status, and outcome as work crosses tools and teams.	Governed human and agent steps coordinate across replaceable tools while evidence, exceptions, and case history persist.
Access & Execution Control Common gap: Agent actions lack consistent limits and run on shared or standing credentials	Authority is withheld or configured separately by pilot; agent identities, permissions, and action limits remain inconsistent.	Production agents have distinct identities, least-privilege permissions, approved targets, action tiers, and explicit approval gates.	Identity, authority, and permissions tighten, expand, pause, or revoke according to current context, action risk, and demonstrated performance.
Governance & Accountability Common gap: Ownership and decision rights vary	Pilot teams set local rules; ownership, risk acceptance, exception handling, and suspension authority depend on who deployed the agent.	Named owners, decision rights, approval thresholds, exception paths, and lifecycle policies cover every production agent.	One governance model follows agents across providers and workflows while accountable humans retain traceable decision authority.
Assurance & Resilience Common gap: Actions are hard to verify, contain, or reverse	Logs and checks vary by pilot; teams reconstruct actions manually and improvise containment, rollback, or recovery.	Production agents preserve evidence and outcomes, track performance and safety, and maintain monitored containment and tested rollback or recovery paths.	Continuous monitoring detects unsafe behavior, stops downstream actions, verifies recovery, and updates permissions, tests, or approval rules before the agent returns.

Six steps to reliable agent operations

The matrix shows where a workflow needs improvement. These six steps turn that assessment into a plan for production. Revisit the checks when tools, data, or permissions change.

1

Prove the agent removes work

Start with one workflow that requires the agent to choose what happens next and has clear inputs, outputs, and a stopping point. Name the owner and document the workflow's current manual steps, analyst effort, and delays involved today. Define the expected result and risk limits, then compare the pilot with that baseline using the measures teams rank first: time to resolve and issues investigated.

2

Give the agent reliable context

Map the identities, tools, connections, and data sources the agent will use. Verify where each tool and input comes from, who owns it, and whether it should be trusted. Connect the affected user or service, device or workload, application, network path, and owner to the same case before the agent recommends or changes anything. 51% say scattered context and signals are a leading obstacle to faster resolution.

3

Carry one case through closure

Decide where the official case lives. Use one case ID for evidence, ownership, approvals, actions, and outcomes across tools and teams. Keep that record intact from detection through verification and automate the handoffs across tools and teams.

4

Define exactly what the agent may do

Run each agent under its own verified identity, never a shared or standing credential, and grant only the access its task requires. Prefer short-lived credentials over standing ones, and define what agents may recommend, what they may perform after approval, and what they may complete automatically once their track record supports it. Require approval for destructive, disruptive, and high-impact changes that delete data, disable access, alter policy, or interrupt service.

5

Assign agent ownership

Identify who owns the agent, approves high-impact actions and exceptions, accepts risk, can suspend it, and leads recovery. Let the platform that ships your agents enforce those rules for the agents it runs, then hold every agent you build to the same rules.

6

Test controls before expanding authority

Before production, test manipulated inputs, faulty tool responses, permission failures, and actions that do not complete as expected. Monitor actions and outcomes, preserve the decision and action record, and feed analyst verdicts back so the agent learns and stops surfacing what the team keeps dismissing. Expand authority only when operating results and recovery tests support it.

Measure the payoff across three areas:

Operational:

Measure faster resolution time and cases completed per analyst hour. Include review effort and track errors.

Financial:

Compare labor and processing time savings with agent, integration, and oversight costs.

Strategic:


Use recovered analyst time for proactive work, including threat hunting, architecture, and readiness exercises.

Methodology and Demographics


This report is based on a survey of 921 security, network, and IT professionals conducted in 2026. Respondents include practitioners, architects, managers, and technology leaders responsible for security operations, network operations, cloud security, identity, data protection, and AI governance across a wide range of industries and organization sizes.

The research examines NOC and SOC AI agent adoption, operational needs, and governance. Findings reflect respondents' reported experiences and expectations. For the full sample of 921 respondents, the margin of sampling error is approximately ± 3.2 percentage points at the 95% confidence level.


RESPONDENT ROLE



FUNCTIONAL AREA



INDUSTRY



COMPANY SIZE



©2026 Cybersecurity Insiders. All rights reserved.

Limited editorial citation (up to 100 words and one unaltered chart) is permitted with clear attribution to "2026 State of AI Agents in Security and Network Operations" and a visible link to cybersecurity-insiders.com.

The report sponsor may reference the findings and use individual charts or data points in presentations and marketing materials with proper attribution. The full report, underlying dataset, and research methodology remain the intellectual property of Cybersecurity Insiders and may not be reproduced, redistributed, or incorporated into derivative research without written permission.

This report was produced by Cybersecurity Insiders with the support of Netskope. Permissions: info@cybersecurity-insiders.com



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 how at

Netskope.com

Cybersecurity

I N S I D E R S

BENCHMARK YOUR SECURITY MATURITY

Independent cybersecurity research revealing the gaps
that shape cybersecurity strategy

Cybersecurity Insiders produces independent research through
surveys of security leaders and practitioners worldwide.

Our reports examine how organizations protect their systems,
data, and people, helping readers compare capabilities,
identify gaps, and decide what to improve next.

For more information, visit

cybersecurity-insiders.com