The Impact of Agentic AI on Network Operations

A 2026 Cisco's Research Report

cisco

Findings from a soon to be released 2026 Cisco research initiative on the impact of Agentic AI on Network Operations.

Every conversation about agentic AI in network operations begins with the same question: will enterprises trust it? The data in this report suggests the question is wrong.

The right question, the question this report sets out to answer, is different. **What did enterprises actually decide, why did they decide it now, and what does that imply for the next twelve months?**

Three findings, drawn from the study, frame the answer.

1. IT Can't Hire Its Way Out Anymore

It would take 100 dedicated specialists to keep pace with daily network alerts.

A typical practitioner closes 21 alerts per day. With organizations facing 4,000 daily IT alerts, half of them network-related, manually investigating each one would require a team of roughly 100 specialists. The result of this impossible math: **46% of alerts are closed without investigation, and 45% of investigation time is wasted on false positives.**

Humans simply cannot scale to the volume of data that modern AI-driven environments generate. This is no longer a staffing problem.

2. The Network Has More Blind Spots Than You Think

It takes nearly ten different tools to see the whole network today, and none of them sees the full picture.

Modern network failures don't respect the boundaries between cloud, security, and networking teams. **92% of performance issues span multiple domains**, driving the average fix time to **88 hours**, seven times longer than the typical 12.5-hour incident. Nearly half of all incidents take more than a day to resolve, and 1 in 10 takes longer than a week.

The reason? Organizations rely on a patchwork of nearly 10 different tools, each offering a narrow slice of data, to monitor networks today. The result is IT teams guessing at problems they can't fully see.

3. The Autonomy Consensus Already Happened

51%
of AI-using organizations have agentic AI in production for NetOps today

80%
are comfortable with high or full autonomy on production networks

84%
expect to be operating in an AI-led model within twelve months

80% of operators are on board for high or full AI autonomy on production networks right now.

The industry has stopped debating whether AI will run the network and started adopting the model that will. 51% of organizations have already deployed agentic AI in production, 24% endorse full autonomy with no human oversight, and 84% expect to be operating in an AI-led model within the next twelve months.

But IT leaders aren't handing over the keys blindly. The conditions of trust are unanimous across all 1,396 respondents: explainability, approval gates, policy limits, override mechanisms, and immutable audit trails. They aren't asking for less automation, they're asking for trustworthy automation.

Discoveries from the report include:

METRIC	REALITY
Daily alerts per organization	**~4,100**, roughly half network-related
Alerts a practitioner can close per day	**~21** out of thousands
Specialists needed to clear backlog by hand	**~100** (typical NetOps team is 5–15)
Alerts closed without investigation	**46%**
Investigation time spent on false positives	**45%**
Performance issues spanning multiple domains	**92%**
Mean time to resolve	**88 hours** — 7x the 12.5-hour median

The next twelve months will determine the operating model that runs enterprise networks for the decade to come.

The full report is coming soon

This executive summary previews Cisco's *The Impact of Agentic AI on Network Operations* report. The full report will publish shortly, expanding on each of the findings outlined here with deeper analysis, regional and segment breakdowns, and architectural implications for the decade ahead.

cisco


CISCO