


An abstract graphic composed of numerous thin, flowing lines in shades of purple, magenta, and blue, creating a sense of dynamic movement and connectivity. The lines are concentrated in the upper left and right areas, with a central void.

The Impact of Agentic AI on Network Operations

A 2026 Cisco and Omdia research report

Contents

01	Executive summary
02	IT can't hire its way out anymore
03	Operations has more blind spots than you think
04	The autonomy consensus already happened
05	The conditions for trust
06	The organizational gap that could slow it down
07	Early evidence shows it works
08	The decisive twelve months
09	About this research



01 Executive summary

The debate over whether AI should help run the network has effectively ended.

Bring up agentic AI in network operations, and the conversation almost always circles back to one thing: trust. Will enterprises actually hand over control? An independent study conducted by Omdia, backed by a Cisco-led analysis of aggregated network telemetry data, suggests that the question is already outdated.

According to 1,000 IT and network operations leaders, the shift to AI-led operations is well underway. Three-quarters of organizations have already deployed AI for network operations. Over half are running agentic AI systems that act, not just advise, in production today. 84% expect to reach an AI-led operating model within twelve months, shifting from legacy AIOps to agent-powered operations, otherwise known as AgenticOps.

Among AI users

75%

have already deployed AI for network operations

51%

run agentic AI systems that act in production today

84%


expect an AI-led operating model within twelve months

Source: Omdia

This is not aspiration outrunning reality. It is a rational response to a structural problem: the modern network has outgrown the human capacity to operate it. Alert volumes, change velocity, and cross-domain complexity have reached a scale that no amount of hiring can address. Enterprises have looked at the math and concluded that autonomous operations are not a risk to be managed but a necessity to be engineered.

Major learnings from the study include:

- **IT can't hire its way out anymore.** Network demand is at an all-time high, yet 95% of respondents say that existing, non-agentic tools (AIOps) are failing to keep pace in one or more areas. Organizations that do not use AgenticOps would need a team of ~100 specialists to clear the networking alert backlog by hand.
- **Operations has more blind spots than you think.** With 92% of respondents saying issues commonly span multiple domains and teams relying on ten or more separate tools for visibility, fragmentation has become a structural liability.
- **The autonomy consensus already happened.** Faced with volume that cannot be staffed and fragmentation that cannot be manually correlated, 51% of organizations are already running agentic AI in production, and 84% expect to reach a fully AI-led operating model within twelve months.
- **The conditions for trust are clear.** Enterprises are not asking for less, but for more trustworthy agent-driven automation. Transparency is non-negotiable, with 69% demanding detailed causal tracing for agentic AI actions. 86% agree that a single integrated platform, not another point tool, is the most effective path forward.
- **The biggest barrier may not be the technology; it is organizational readiness.** NetOps practitioners are nearly twice as likely as line-of-business executives to recognize the shift toward AgenticOps, creating a leadership gap that could slow AI adoption and delay operational transformation.
- **AI leaders are already pulling away.** Organizations with the highest levels of AI adoption for NetOps are close to 30% more productive per employee and 3.2X more likely to achieve full end-to-end visibility than less mature peers, suggesting AI maturity is becoming a competitive advantage, not just an operational improvement.



02 Capacity

IT can't hire its way out anymore

For two decades, the answer to rising operational complexity was to add people, add tools, and add process. This model has reached its limit and AIOps-based operations can't keep up.

The arithmetic of alerts no longer works.

The average organization generates around 4,100 monitoring alerts and events per day, with slightly more than half (51%) being network-related. A typical practitioner can review, investigate, and resolve about 21 network alerts in a day. Clearing the daily backlog by hand would require a team of roughly 100 specialists. Most network operations teams number a fraction of that.

The math of alerts

4,100

alerts and events
per day, per
organization

51%

are network
related
(~2,100/day)

21

alerts a single
practitioner can
clear in a day

~100

specialists to clear
the network alert
backlog by hand

Source: Omdia

The consequences are predictable and corrosive:

46% of network alerts are closed without investigation, set aside because of staffing constraints, recurring noise, or false positives.

45% of investigation time is spent chasing false positives—work that produces no resolution and no value.

65% of leaders report that “alert fatigue” is a meaningful source of employee dissatisfaction, and 67% say alert volumes prevent their teams from doing other critical work.

Source: Omdia

This is no longer a staffing problem that better recruiting can solve. It is a structural mismatch between the volume of information environments produce and the capacity of humans to interpret and act on it.

Change has become continuous.

Nearly three-fifths of organizations now make changes to their production network environments daily or more often. Half of those organizations are making multiple changes per day and for a meaningful share, change happens multiple times per hour. 57% say their current change processes can't match the speed of the updates that are required today. The network is no longer a stable system that is occasionally modified; it is a system in constant motion, and the operational model must keep pace with that motion.

How often production networks change

Change is no longer occasional; for most teams it is a daily, often hourly, event.

Change daily or more often ~59%

Multiple changes per day ~29%

57% say their current change processes can't match the speed required today

Source: Omdia

AI itself is intensifying the pressure.

In one of the study's central ironies, the technologies meant to relieve operational burden are also creating it. Two-thirds (67%) of respondents say the generative AI boom has significantly increased network complexity, and 70% say AI adoption has increased, not decreased, NetOps workloads. Notably, 95% of organizations say their existing, non-agentic tools (AIOps) fall short in significant ways, mostly by requiring too much human interpretation and lack of cross-domain visibility.

The AI story goes beyond tools and workloads. It also represents a fundamental shift in network consumption and how work happens. A Cisco analysis of direct-to-AI traffic across campus and branch networks found that average daily AI traffic is on a trajectory to double every six months. That projection is grounded in hard numbers from the same analysis: per-client direct-to-AI traffic jumped 54%. This acceleration is likely fueled by the growing complexity of AI tasks which demand more data exchange than a simple query.

Crucially, this data captures only a fraction of the AI traffic reality. It excludes traffic generated by AI embedded in platforms like assistants and copilots. It also omits the inevitable surge of continuous machine-to-machine traffic generated by agents. In fact, Cisco research found that AI agents act as network “power users,” operating at software speed rather than human speed. Empirical testing shows up to 450% more total traffic is generated per task when performed by an agent.

AI's share of network traffic

Every 6 months

AI traffic doubles—then doubles again, and again



Source: Cisco

AI isn't disruptive because of how big it is...yet. It's disruptive because it's compounding and behaves like nothing your network and operations were built for.

The picture that emerges is a legacy AIOps operating model under strain on every axis at once: more alerts, more change, and a lurking tsunami of AI traffic growth.



03 Fragmentation

Operations has more blind spots than you think


If volume is the first structural problem, fragmentation is the second.

Modern network failures do not respect organizational boundaries. 92% of organizations report that performance issues commonly span multiple domains: networks, cloud, security, applications, endpoints, etc. Problems rarely live where they first appear. A latency complaint may originate in a cloud networking misconfiguration, surface as an application slowdown, and be flagged by a security tool, all at once.


Yet the tools used to diagnose these problems remain stubbornly siloed. Organizations rely on an average of nearly ten separate tools to maintain end-to-end visibility across their networks. Each tool offers only a narrow slice of the picture, none offering the whole. The result is teams diagnosing problems they cannot fully see, correlating logs and events by hand, and handing incidents across functional boundaries.

Fragmented visibility and tools

92% of performance issues span multiple domains at once — problems rarely live where they first appear



~10 separate tools for end-to-end visibility — each offering only a slice of the picture, none the whole.



Source: Omdia

Time to resolve an incident



46%

of organizations say incidents typically take more than one day to resolve

1 in 10

say incidents typically take longer than a week to resolve

Source: Omdia

The cost of this fragmentation shows up most clearly in resolution time. The mean time to resolve a network incident is 88 hours. That's seven times longer than the 12.5-hour median. That gap means that while many organizations are able to handle issues reasonably well, a sizable number of organizations regularly struggle with complex incidents that escalate into multi-day or multi-week ordeals. Organizations are being held hostage by a few outlier incidents. These catastrophic, hard-to-solve events are acting as an anchor, dragging down resolution times when operations teams run into them, preventing NetOps from reaching the speed the business demands.



04 Autonomy

The autonomy consensus already happened

Faced with volume that cannot be staffed and fragmentation that cannot be manually correlated, enterprises have arrived at a striking consensus: the solution is not more people or more dashboards, but autonomous, agent-driven operations.

The evidence of that conclusion is unusually direct:

75% of organizations have deployed AI technologies (AIOps and AgenticOps) for network operations.


51% are running agentic AI systems act in production today.

80% are comfortable granting AI a high or fully autonomous role with 56% needing a human to approve actions.

82% are comfortable letting AI make production changes without human permission first, for at least some categories of action.

Source: Omdia

The autonomy spectrum




Most organizations (85%) believe AI agents will autonomously manage most operational tasks within five years. In fact, 80% say they are comfortable with AI taking a high or fully autonomous role in NetOps, today. 82% are comfortable letting AI make changes to their production networks without seeking human permission first for at least some categories of action, including rerouting traffic, adjusting wireless parameters, isolating suspicious endpoints, and resolving incidents end-to-end.

This is not a hypothetical posture. It is a settled direction of travel. 84% of organizations expect to reach a fully AI-led operating model within the next twelve months, and agentic adoption across specific NetOps processes is poised to roughly double in that period.


Where agents are heading by NetOps process

The next twenty-four months of planned agentic adoption, process by process (roughly 2x today's levels across everything).



The trajectory of adoption

Three different signals from the study — each a separate question — all pointing the same way.




Source: Omdia

The top drivers for adopting agent-driven AI in NetOps are enhancing security and threat detection, taming network complexity, and improving end-user experience. Enterprises are not adopting AgenticOps because it is fashionable. They are adopting it because the alternative, failing to automate, carries consequences they consider unacceptable: 64% foresee security gaps or increased enterprise risk, 62% anticipate more frequent or severe service degradations, 53% expect more failed AI initiatives, and 46% expect a diminished user experience.

The cost of not automating

Leaders don't see inaction as neutral — most expect it to cost them, on more than one front.



Source: Omdia

The industry has stopped debating whether AI will run the network. It has started deciding how.



05 Trust

The conditions for trust

Enterprises are not handing over the keys blindly. The same leaders embracing AgenticOps are explicit about the conditions under which they will grant it. And those conditions are remarkably consistent.

Transparency is non-negotiable.

While the desire for AgenticOps is strong, key conditions for trust depend on visibility and transparency into how agents reason. In fact, 69% of organizations require detailed explainability for agent-driven actions, and 36%, the largest single group, say nothing less than full observability(detailed tracing, summarized rationale, and post-action audit) is the minimum acceptable standard. The message is clear: trust will not be gained without transparency.

Trust needs a unified interface.

AgenticOps can span agents, data, and domains, but trust requires one place to see it all. A unified interface gives operators shared context, visible reasoning, and a clear record of what agents recommend and do. Without it, humans are left piecing together decisions across disconnected tools. Enterprises know it. 86% say the most effective path forward is a single integrated platform spanning compute, security, applications, and networking.

The path forward

86%

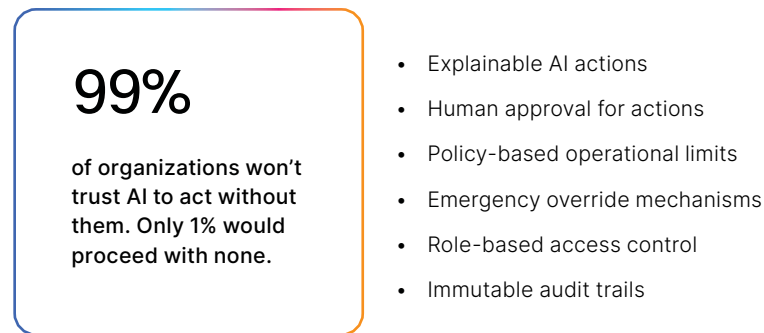
say the most effective path forward is a single integrate platform spanning compute, security, applications, and networking.

Source: Omdia

Guardrails for AI and agents are expected as a set, not a menu.

When asked which guardrails are essential for AI and agent trust, respondents selected an average of three from a list of six, with little separation between them:

Required guardrails



Source: Omdia

Organizations do not see these as alternatives to choose among, but as a baseline they expect collectively.

Taken together, these findings reframe the trust question entirely. Enterprises are not asking for less. They are asking for AI they can trust. The kind that comes through a single, unified interface with agent-driven autonomy bounded by explainability, approval gates, policy limits, override mechanisms, and audit trails.




06 Readiness

The organizational gap that could slow it down

The largest obstacle to agentic AI may not be technical at all. It may be organizational. The study reveals a significant perception gap between those closest to the day-to-day work and those who control the budgets. NetOps practitioners, who feel the operational strain daily, are roughly twice as likely as line-of-business executives to recognize the shift toward AgenticOps. In fact, line-of-business executives with budget influence are more likely to view AI in operations as a matter of insights and recommendations — the older “AIOps” paradigm, rather than agent-driven action.

How each group views AI’s role in IT operations


The closer a role sits to the day-to-day work, the more it recognizes AI’s job is shifting from advising to acting.



This gap matters because the people who most clearly understand the need are not always the people who authorize the investment. It creates an imperative for technical leaders to educate their financial and business peers on the transformational, not merely incremental, nature of AgenticOps.

Other barriers for AI in NetOps are what one might expect. Leading concerns include data security and privacy (38%), AI implementation and integration complexity (34%), and risk and compliance (30%). Tellingly, the barriers cited least often are “unclear goals” and “lack of leadership support,” indicating that organizations are not short on motivation or direction. The obstacles are about execution and trust, not intent.

Biggest barriers to using AI in NetOps





07 Results

Early evidence shows it works

Is the bet paying off?

Organizations that have invested most deeply in AI for network operations, including deploying AgenticOps, are demonstrably better off than their less-mature peers on the dimensions that matter most:

AI in NetOps maturity payoff

~30%

more productive per practitioner. The most AI-mature organizations, with AI and AgenticOps, resolve more alerts per person per day.

3.2X

more likely to have full end-to-end visibility with AI and AgenticOps than the least AI-mature peers.

Source: Omdia

Productivity

Greater individual productivity. Practitioners at the most AI-mature organizations using AgenticOps are close to 30% more productive than those at the least AI-mature, investigating and resolving meaningfully more alerts per person per day. All else equal, that productivity gain translates into dozens of fewer specialists needed to work through the same backlog, freeing scarce talent for higher-value work.

Visibility

Better visibility. The most mature agentic-driven organizations are 3.2 times more likely to have full end-to-end visibility into edge infrastructure bottlenecks, and far more likely to recognize the cross-domain nature of performance issues, a prerequisite to resolving them.

Capacity

Greater confidence and capacity. Mature organizations are markedly more likely to report that their teams are keeping up with their workloads, rather than barely keeping pace or falling behind.



08 Outlook

The decisive twelve months

The findings in this report describe an inflection point that has already occurred. Enterprises have concluded that networks are now too voluminous, fast-changing, and fragmented for human-scale operations. Autonomous, agent-driven operations are their answer. They are already deploying them in production, defining the conditions for trust, and expecting an AI-led operating model within a year.

The open questions are no longer about direction. They are about execution: Can organizations turn their requirements for trust into an operating model that works at scale? Can they consolidate fragmented tooling into the integrated platforms they say they need? And can technical leaders bring the rest of the business along fast enough?

The next twelve months will not decide whether AgenticOps runs the enterprise network. The enterprises in this study have already decided that. The next twelve months will decide how well, and which organizations emerge with an operating model built to run the network for the decade to come.

From operational chaos to autonomous efficiency.

To move from the “current reality” of operational chaos to the “future state” of AgenticOps and agent-first efficiency, leadership must take decisive, strategic action:

01

Pivot from “hiring” to “architecting”

Stop trying to solve a math problem with human headcount. You cannot hire enough people to clear a 4,100-alert-per-day backlog, and even if you could, it's not a sustainable business model. Shift your focus and budget from “adding more bodies” to “architecting more autonomy.” Equip your teams to direct, validate, and collaborate with agents, freeing their expertise for the complex decisions, architecture, and optimization that require human judgment. The goal is to move your best people from the front lines of the alert flood and system change churn to the high-value work of strategic network design.

02

Bridge the “executive-practitioner” divide

There is a dangerous disconnect between those who run the network and those who fund it. While your NetOps team knows the current human-based model is broken, your line-of-business executives may still view agent-driven automation as a “nice-to-have” rather than a survival strategy. Your job is to translate the technical reality—the 4,100 daily alerts and problematic resolution times—into business risk. Frame the move to AgenticOps not as an IT upgrade, but as a risk-mitigation strategy that protects the business from outages and security breaches.

03

Build a “trust-first” governance framework

Agent action and agent autonomy are only as good as the trust you can have in them. Deploy solutions that provide clear agent security, transparency, auditability, and human-in-the-loop oversight before you scale. Reinforce the tools with clear governance that defines exactly which actions AI can take, making sure that those decisions are logged and explained. Only then can you remove the fear of the unknown. Make “explainable AI” a non-negotiable requirement for every vendor and system you bring into your environment.



09 Appendix

Glossary of AI terms

As AI in network operations matures, a shared vocabulary is taking shape. The definitions below reflect how the industry distinguishes between types of AI as they apply to IT operations and NetOps.

Artificial Intelligence (AI) is the broad capability of machines to perform tasks that normally require human intelligence, such as understanding language, recognizing patterns, solving problems, and learning from new information. It is the umbrella term that includes everything below.

Generative AI is a type of AI that creates new content, such as text, code, images, or summaries, in response to a prompt. Rather than only classifying or predicting, it produces original output, which is what powers tools like chat assistants and coding copilots.

AIOps grew out of machine learning and analytics, with generative AI now extending it. It detects issues, correlates events across systems, and suggests fixes, but generally a human still has to interpret the findings and act.

Agentic AI goes a step further by acting on its own, within defined guardrails. It makes decisions, takes action, and adapts to changing conditions to achieve a goal, with a human directing agents and setting their objectives rather than executing each step. Humans stay the architect while the agent handles the doing.

Agentic Operations, or AgenticOps, is a human-led, agent-driven operating model for IT where people and AI agents partner to solve complex operational issues. Agents sense, reason across domains, and act within defined guardrails, while humans set intent, guide decisions, and retain oversight.

Information in this report includes data from an independent 2026 study conducted by Omdia, which surveyed over 1,000 IT and network operations leaders at organizations with 500 or more employees across North America, Western Europe, and Asia-Pacific. Additional sources: Cisco analysis of aggregated telemetry from Meraki campus and branch networks to measure observable traffic directed to identified AI services, Cisco's 2026 report AI Impact on Wide Area Networks, and Cisco IT.