



No time to wait: The accelerating impact of AI on campus and branch networks

Although IT leaders acknowledge they must accelerate their network modernization efforts, even aggressive AI adopters say their networks are not AI-ready, according to new Cisco and Foundry research.



Contents

- 03 Executive summary
- 04 AI's impact on network traffic is substantial
- 06 Capacity limitations are emerging faster than expected
- 09 AI adoption hinges on trust. Security fuels trust.
- 12 Existing network environments are compounding business risks
- 14 Aggressive AI adopters are approaching modernization differently
- 15 Workplace network strategies will be foundational to company success in the next two years

Executive summary

Much of the AI-readiness conversation has centered on GPUs, cloud platforms, and data center buildouts, but new research from Cisco and Foundry points to a different challenge: AI expansion is placing extraordinary pressure on enterprise networking, especially across workplace networks – also known as campus and branch networks – where users, devices, applications, and autonomous systems come together.

More than 3,400 IT and networking decision-makers across 15 countries shared their insights and have sounded the alarm about the rapidly growing strain on networks as AI adoption expands across generative, agentic, and physical AI environments.


The message is clear: There is a widening gap between AI ambition and network readiness. [Cisco's 2025 AI Readiness Index](#) put a number on a problem the industry already sensed: Only 15% of organizations said they have networks flexible and adaptable enough to support AI at the necessary scale.

Less than a year later, this new research finds that campus and branch environments are fundamentally mismatched with the scale, speed, and variability of AI-driven traffic now moving across enterprise networks. As organizations accelerate AI deployment, networking teams are facing three pressures converging at once: rapidly escalating traffic growth, fast-approaching capacity limitations, and rising security complexity.

For AI initiatives to scale successfully, enterprises must upgrade their networking infrastructure or risk being surpassed by competitors that are truly AI-enabled.



Average increase in campus and branch network traffic tied to AI workloads over the past 12 months



Additional growth in traffic expected within the next year



Traffic projected to reach **3x** current levels over the next three years, compounded across generative AI (genAI), agentic, and physical AI




Percentage of organizations already face or expect to face campus and branch capacity limitations within the next 24 months

AI's impact on network traffic is substantial

AI workloads are changing traffic patterns across enterprise environments in ways many existing workplace networks were never designed to support. For example, 67% of the participating respondents reported increases in east-west traffic – the lateral device-to-device or server-to-server communication required for AI agents to exchange data – tied to these workloads. Additionally, 61% noted growth in continuous automated traffic generated by AI systems.

AI adoption is already reshaping campus and branch network traffic




Those changes become even more significant as organizations move beyond genAI experimentation and deeper into agentic AI capable of autonomous action.

“Usually, networks are designed for consistent traffic, like SaaS and CRM traffic, and there aren’t a lot of unpredictable traffic patterns,” said the head of AI strategy for global IT and network engineering operations at a large U.S. technology company. “Suddenly, three AI agents are trying to talk to each other and solve a problem. That is going to be a big thing... how do we support increased east-west traffic?”

That observation reflects a broader reality emerging across the workplace. AI systems are creating new categories of workloads that are more distributed, dynamic, and often sensitive to latency and reliability than many traditional enterprise apps.


This matters especially in campus and branch environments, because these networks sit closest to employees, operational workflows, connected devices, and customer interactions. As organizations embed AI more deeply into day-to-day processes, enterprises’ operational edge increasingly becomes one of the first places where networking limitations show up.

The projected growth figures add urgency. Respondents already reported substantial increases in AI-driven traffic, and expectations for future growth remain aggressive. The expected growth in network traffic is mirrored in the expected change of AI adoption, as respondents anticipate significant increases over the next 12 to 24 months:




GenAI

↑
90%



Agentic AI

↑
85%



Physical AI

↑
70%




50%

of AI-driven demand is concentrated in wireless networks within campus locations (Wi-Fi).

Capacity limits are emerging faster than expected

Capacity concerns are no longer hypothetical. More than one-third of the respondents reported that AI is already driving increased capacity demands across their networks, and 76% said their campus and branch environments require upgrades to support current and future AI-driven workloads.

AI workloads are expected to cause significant network traffic increases



Combined figures reflect the compounded effect of the three projected AI traffic increase estimates rather than a simple average; layered growth can produce values above 100%, where 100% represents a doubling of current traffic.


The pressure appears to be especially acute among organizations furthest along in AI adoption. The research found that enterprises with broad deployment of genAI technologies are significantly more likely to report traffic growth, operational complexity, security concerns, and infrastructure readiness challenges than those organizations still operating in earlier stages of adoption. Only 30% of these aggressive AI adopters – defined in the research as organizations with broad genAI deployment across the enterprise – said they are fully prepared to support projected AI growth across the network.

That finding suggests that the remaining 70% of aggressive AI adopters may still be underestimating the scale of the infrastructure challenge ahead. Enterprises that have already moved beyond experimentation and into broad operational deployment are encountering the limitations of existing campus and branch environments much sooner than expected.

IT decision-makers increasingly recognize that they must rethink network priorities. Among them, 93% said they are accelerating modernization initiatives in response to AI-driven demand. The urgency reflects the reality that AI workloads often depend on low latency, continuous responsiveness, and reliable real-time interactions to function effectively at scale.

A retail executive described an AI loss prevention initiative that ultimately failed because network latency made the system operationally ineffective. The delay in the AI tool – responsible for analyzing data and determining whether a theft was occurring – was significant enough to undermine its usefulness. “There is a delay of about five seconds,” said the vice president of infrastructure, network, and end user services at a U.S.-based retail enterprise. “In those five seconds, somebody already leaves the store. So it’s pointless.”

Network latency caused “a delay of about five seconds. In those five seconds, somebody already leaves the store. So it’s pointless.” – U.S.-based retail executive



said they are already facing or expect to face campus and branch capacity limitations within the next 24 months.

“Observability is a huge gap. There is experimentation going on all over the place, and there is no way for us to really identify if somebody is deploying some kind of service on our network, whether it is a genAI solution or an agentic solution.”

— head of AI strategy, U.S.
technology company

Examples such as this illustrate why networking limitations are becoming more than a technical inconvenience. AI systems are increasingly embedded in real-world workflows, where response time matters. Even small delays can directly impact customer experiences, operational efficiency, and costs.

At the same time, many organizations lack visibility into how AI workloads are evolving inside their environments. People across business units, departments, and operational teams are experimenting with AI, often without centralized governance. This fragmentation limits operational teams' visibility into what's deployed and how systems interact across the network.

“Right now, we don't even know what the AI-driven demand is,” said an AI strategy leader. “Observability is a huge gap. There is experimentation going on all over the place, and there is no way for us to really identify if somebody is deploying some kind of service on our network, whether it is a genAI solution or an agentic solution.”

Organizations aren't just trying to add bandwidth or push more data through their networks. They're trying to scale networking as workloads shift more often, traffic becomes harder to predict, and AI systems run more independently across distributed environments. In the process, many are realizing that their campus and branch network strategies were built for conditions that no longer apply.



AI adoption hinges on trust. Security fuels trust.

Respondents cited security complexity as the most significant challenge associated with AI-driven network demand. Unlike in many previous waves of enterprise technology adoption, security is emerging not simply as a parallel concern but as a direct barrier to AI scale. If organizations cannot trust the systems, visibility, and controls surrounding AI workloads, they become far less willing to operationalize AI broadly across the enterprise. Respondents pointed to expanded attack surfaces, shadow AI activity, inconsistent policy enforcement, and limited visibility into how AI-driven traffic moves across enterprise environments as key factors driving that hesitation.

Organizations reported various challenges:



78%

expect security risks to further increase as AI adoption expands beyond generative use cases




77%

said the technology has already expanded their attack surface in the last 12 months



71%

believe that threats are evolving faster than existing controls can adapt



69%

described growing blind spots in monitoring and visibility

At the same time, IT leaders said their existing security models may not keep up with the complexity of AI environments. Although 86% said they have added security controls for AI workloads, 61% are holding back from scaling AI initiatives further until they have greater confidence in their security posture.

“The issue from a security standpoint is that it’s hard to create the guardrails for every possible AI tool that your organization must use,” explained the retail executive.

As organizations move toward more autonomous and agentic AI environments,

many leaders increasingly view the network itself as one of the most effective enforcement points for managing AI-related security risk, visibility, and policy control at scale.

Another IT leader described their current environment as one of continuous adaptation and operational uncertainty: “We’re just playing catch-up at the moment,” said this leader, a vice president of IT and digital infrastructure in the U.K. education sector. “It’s a worrying time, and I think it’ll stay like this for another 18 months or two years.”

“The issue from a security standpoint is that it’s hard to create the guardrails for every possible AI tool that your organization must use.”

– U.S.-based retail executive



The underlying issue is straightforward: AI workloads introduce a level of operational dynamism that many existing security approaches were not designed to manage.

For example:

- Agentic systems communicate constantly across environments.
- AI-enabled workflows can automatically trigger actions across applications and networking.
- Shadow AI initiatives can quickly take hold inside departments long before networking or security teams are aware of them.


These issues have led networking leaders to see network architecture, observability, and security posture as deeply connected, with a secure network serving as the most effective enforcement point for the unique security challenges AI brings. Gaps that once caused minor operational issues can now create significant governance and security exposure inside environments where AI systems continuously generate activity across distributed networks.



Existing network environments are compounding business risks

These networking pressures extend well beyond technical operations. Organizations are starting to understand that network readiness will directly affect their ability to compete, scale AI initiatives, and innovate into real outcomes over the next several years.


More than 90% are aware of financial and competitive risks if they fail to adapt campus and branch networking for AI-driven demand. Respondents also cited risks such as operational disruption, degraded user experiences, longer response times, rising costs, and reputational damage from outages or inconsistent policy enforcement across distributed environments.



of IT leaders agree
there are higher
long-term costs due
to reactive upgrades
or remediation.



IT leaders cited the top business risks of delaying or failing to modernize networks for AI*




Higher long-term costs due to reactive upgrades or remediation




Inability to meet customer expectations




Missed business opportunities



Increased security risks or an expanded attack surface



Decreased operational efficiency



Business reputation risks



Falling behind competitors



Many IT leaders are aware of the financial and competitive risks if they fail to adapt campus and branch networking for AI-driven demand.

* Percentages reflect those respondents who agree or strongly agree with these risks.

Aggressive AI adopters are approaching modernization differently

Among the surveyed enterprise-wide AI adopters, 96% said the technology has somewhat or significantly increased their network modernization plans, versus 88% of those organizations that have not deployed AI enterprise-wide. Mature AI adopters are more aggressively upgrading their campus and branch environments than those that have not deployed AI enterprise-wide to:

	Mature AI adopters		Early-stage AI adopters
Support AI capacity demands	55%	vs.	26%
Meet compliance requirements	53%	vs.	32%
Keep ahead of the competition	51%	vs.	26%

The aggressive AI adopters are approaching campus and branch networking strategically rather than tactically. They recognize that AI is fundamentally changing the operational behavior of enterprise environments, not just handling a temporary surge in traffic.

That strategic shift is shaping decisions about segmentation, wireless performance, automation, integrated security, network intelligence, and long-term architectural planning. Even with significant progress in these areas, aggressive AI adopter respondents are concerned about moving quickly enough to address the rapid pace of genAI deployments along with agentic and physical AI adoption.

This uncertainty is why modernization timelines are compressing. IT leaders know that decisions they make today will determine whether their organization can successfully support the next wave of AI and other emerging technologies such as quantum computing.

Workplace network strategies will be foundational to company success in the next two years

Organizations that modernize now will be better positioned to support increasingly dynamic AI workloads, maintain operational performance, strengthen visibility and security, and scale AI initiatives with confidence.

The research findings ultimately point to a broader conclusion: AI adoption is reshaping enterprise operations much faster than organizations anticipated, with its effects extending beyond centralized networking to the edge of the enterprise.

Network modernization efforts must move beyond the data center to also encompass workplace networks, where AI workloads increasingly intersect directly with employees, workflows, devices, customer interactions, and operational systems. As AI adoption expands across generative, agentic, and physical systems, network demands will continue to become more complex, distributed, and performance-sensitive.

Now is the time to act. The research shows that 85% of organizations expect moderate or major expansion in agentic AI deployment within the next 24 months, and 73% said they are already facing or expect to face campus and branch network capacity limitations within that same timeframe. Together, those findings point to a rapidly narrowing window for infrastructure modernization.

Organizations that modernize now will be better positioned to support increasingly dynamic AI workloads, maintain operational performance, strengthen visibility and security, and scale AI initiatives with confidence. Those that delay modernization risk allowing infrastructure limitations to become a direct constraint on innovation, execution, and competitiveness.

AI is creating fundamentally different operational environments, defined by autonomous systems, real-time interactions, and continuously shifting traffic patterns across distributed enterprise environments. In such environments, network modernization is no longer simply an infrastructure initiative. It is becoming a prerequisite for operating and competing effectively in an AI-powered economy.

Learn how to get ahead of the AI-driven surge the next two years is likely to bring with an [AI-ready secure network architecture](#).

About the research

Foundry conducted a quantitative survey, sponsored by Cisco, of 3,472 CIOs as well as networking, end user computing, and technology leaders in Asia-Pacific, Europe, the Middle East, Latin America, and North America. The respondents work at organizations with 500+ employees that have an average of 3,292 campus/branch locations. In addition, Foundry conducted six in-depth interviews with executives in Asia-Pacific, Europe, and the United States. All research was conducted between March and April 2026.