Statement of Bob Everson
Chief Architect, Provider Mobility at Cisco Systems, Inc.

Before the United States Senate Committee on Commerce, Science, and Transportation
Subcommittee on Telecommunications and Media

Hearing entitled, "Intelligent Networks: Powering Artificial Intelligence and Transforming Communications"

July 30, 2026

1. Introduction.

Chairman Fischer, Ranking Member Luján, Chairman Cruz, Ranking Member Cantwell, and distinguished Senators, thank you for the invitation to testify today. My name is Bob Everson, and I am a proud resident of Dallas, Texas. I have spent more than 25 years at Cisco, where I currently serve as chief architect of our provider mobility team. In this role, I lead Cisco's platform definition and development, end-to-end architecture, ecosystem strategy, and strategic partnerships. My focus includes developing AI systems that operate across distributed environments and defining the mobile networking technologies, architectures, and products that advance enterprise mobility, IoT, and the platforms and ecosystems needed to bring these capabilities to scale.

For decades, Cisco has served as the foundational layer for the world's most critical networks, supporting a diverse customer base that spans service providers, global enterprises, government agencies, and small-to-medium businesses. Today, Cisco is building the critical infrastructure for the AI era. Our goal is to ensure the United States remains the global leader in fixed and wireless technology by building a secure, interoperable infrastructure ecosystem where licensed, unlicensed, and satellite networks work in concert and where networks are secure, observable, and highly scalable. While Cisco develops the infrastructure, our customers—whether fixed or mobile service providers, enterprises, governments, or hyperscalers—operate their own networks. For service providers, Cisco provides the core infrastructure behind the tower and the Radio Access Network (RAN).

Beyond hardware, we are integrating AI into the network itself. Cisco's deep expertise in network analytics and security makes it possible for us to automate operations, predict and prevent anomalies in real time, and enable customers to bring AI models into the enterprise with visibility and trust by leveraging the network. We recognize that AI workloads—characterized by massive data movement, high-density traffic, and some anticipated use cases requiring sub-millisecond latency—place unprecedented demands on the network. By combining our industry-leading networking, security, and observability portfolios, Cisco is ensuring that as organizations scale their AI initiatives, they do so on a foundation of resilience, security, and performance.

I want to first thank this Committee, and Congress, for its tireless work to make 800 MHz of spectrum available for commercial use and for recognizing the importance of the 6 GHz band to enterprise networks powering schools, manufacturing facilities, hospitals, and campuses of all sizes. A balanced spectrum policy is critical to enhancing service for consumers and businesses in both rural and urban areas and across different modes and use cases of technology. I would also like to thank Federal Communications Commission (FCC) Chairman Brendan Carr and National Telecommunications and Information Administration (NTIA) Administrator Arielle Roth for their leadership in bringing these efforts to life. As I will explain, network traffic is only going to increase in volume and complexity, and these policies will help enable continued U.S. technological leadership.

At Cisco, our approach toward AI in communications networks can be distilled into two questions: How is AI reshaping our networks, and how can networks leverage the power of AI? Today, I will address these questions and share how Cisco is building the infrastructure for more

resilient, secure, and capable networks. My testimony draws on three recently released reports that draw from real-world traffic analysis, third-party data, Cisco-controlled lab tests with AI agents, and interviews with more than 6,000 organizations, all of which I have appended to my written testimony.

2. AI Usage Is Reshaping the Network.

AI traffic has led to increased demand—we expect a 209% increase in campus and branch traffic volume in the next three years. However, the more consequential shift is in the nature of that traffic. Today, I will highlight what drives that change and how networks and policy must evolve to support its deployment and adoption.[1] AI workloads are not like traditional web traffic: they will change traffic shape, symmetry, duration, and criticality.[2] AI workloads are also twice as long in duration and significantly more upstream-heavy due to large, context-rich prompts. The use of new agentic AI, AI inferencing, and applications built on the connectivity fabric that AI-native networks will make possible are largely driving these workloads.

Agentic AI: AI agents are supercharging demand on the network, generating up to 450% more traffic per task than humans in a Cisco-controlled test. Unlike consumer-driven traffic, which produces minor, episodic spikes throughout the day, AI agents are always on.[3] Agents may continually perform automated network configuration checks, analyze telemetry data, or interact with other agents to perform business functions.

Additionally, survey data shows 67% of respondents reporting an increase in "east-west" traffic.[4,5] Modern AI models are rarely processed by a single server; they are distributed across clusters of GPUs and compute nodes. These nodes must constantly communicate with one another to share model parameters, gradients, and training data, which creates a massive amount of lateral traffic within a network. Traditional planning assumptions (e.g., burstiness, downlink dominance, human-paced interactions) will need to adapt to the new reality of AI and agentic AI traffic that lasts longer, demands more upstream capacity, and operates at software speed. Continuous measurement of AI-specific Key Performance Indicators (KPIs) and adaptive network architectures will be essential to sustain expectations for performance and user experience, particularly as network operators use agentic operations to assist with network management functions.

AI Inferencing: While AI inference traffic is currently negligible compared to dominant categories such as video streaming, observed growth rates are exceptional. Token-consumption data shows nearly 10x year-over-year growth, while some service provider measurements show

---

[1] "AI Impact on Wide Area Networks," Cisco. (June 2026). Available at: https://www.cisco.com/c/dam/en/us/solutions/collateral/artificial-intelligence/mass-scale-infrastructure/ai-network-traffic-report.pdf.
[2] *Id*.
[3] *Supra*, Note 1.
[4] "East-West" traffic refers to the movement of data between servers, devices, or nodes *within* a network, rather than moving into or out of the network.
[5] "No Time to Wait: The Accelerating Impact of AI on Campus and Branch Networks," Cisco. (June 2026). Available at: https://www.cisco.com/c/dam/m/en_us/solutions/networking/ai-impact-campus-branch-networks/documents/the-accelerating-impact-of-ai-on-campus-and-branch-networks.pdf.

around 4x growth in just the last eight months.[6] Sustained growth at these rates signals that AI traffic will become a meaningful component of overall network traffic by 2035.[7] AI inference paths will also become strategic network assets, requiring high levels of resilience, observability, and differentiated treatment (*e.g.*, Quality of Service (QoS) and path security).[8]

Enterprise Networks: In campus and branch networks—like the one that powers the Senate office building we are sitting in today—we have already seen customers report a 34% increase in traffic tied to AI workloads over the last 12 months, and they expect to see a 96% increase this coming year.[9] Half of enterprise customers report that AI demand is concentrated on their Wi-Fi networks, and 73% of organizations already face or expect to face campus and branch capacity limitations within the next 24 months.[10] This is largely because large majorities of organizations report increases in east-west traffic, latency-sensitive traffic, and continuous, automated AI traffic.[11] While the large majority of AI to date has come from foundation models running on central infrastructure, we are seeing enterprises deploy more small language models, open-source models, and specialized models—such as vision and voice models—which can be distributed throughout the network. Each of these characteristics underscores the value of the FCC's forward-thinking decision in 2020 to authorize the full 6 GHz band for unlicensed Wi-Fi use. Indeed, wireless leaders deploying 6 GHz-capable Wi-Fi networks report that the band is solving their congestion issues, enabling high-bandwidth applications, and supporting AI workloads and applications. These developments are translating into concrete economic benefits: one study reports that the U.S. 6 GHz decision will generate $1.2 trillion in economic value per year by 2027.[12]

Infrastructure: AI is driving the shift toward edge computing. Service providers must also consider "AI-native" traffic profiles for several reasons, including technical considerations, cost, and data sovereignty and security issues.

- **Technical:** Physical AI use cases such as robotics, autonomous vehicles, and industrial automation could require sub-millisecond decision-making. If an autonomous robot sends data to a central cloud and has to wait for a response, the round-trip latency could be too high for safe, real-time operation.
- **Cost:** AI operations generate massive amounts of data. For example, high-definition video analytics for public safety can generate terabytes of data daily. Backhauling that data to a central cloud is prohibitively expensive and creates massive network congestion.
- **Data Sovereignty and Security:** Enterprises and governments are increasingly concerned about data sovereignty and security. Many customers have security or regulatory concerns about moving sensitive information across the public internet to a third-party cloud.

---

[6] *Supra*, Note 1.
[7] *Id*.
[8] *Id*.
[9] *Supra*, Note 5.
[10] *Id*.
[11] *Id*.
[12] Telecom Advisory Services, "Assessing the Economic Value of Wi-Fi in the United States" (September 2024), Available at: https://www.teleadvs.com/assessing-the-economic-value-of-wi-fi-in-the-united-states/.

AI operations at the edge of the network address each of these challenges. Latency is dramatically reduced by cutting the distance data needs to travel. Processing data locally similarly allows for transmission of only the result to the cloud, reducing bandwidth and cloud-related costs. And operating at the edge keeps sensitive data local, reducing the risk of interception or unauthorized access.

3. Networks Will Leverage AI for Increased Performance and Resiliency.

While AI workloads present several challenges for network operators seeking to ensure seamless performance, reliability, and security, there is a tremendous opportunity to leverage AI to deliver new applications and better performance, infuse security into the fabric of the network, and manage the increased complexity.

Agentic Operations: Agentic AI will change the nature of traffic on the network, but it will also provide network operators new tools to operate at machine speed and deliver greater performance, efficiency, and security. For service providers, whose networks are the backbone of the entire digital economy, AgenticOps—the shift from manual, human-led network management to an autonomous, AI-driven operational model overseen by humans—is not just a luxury; it is a necessity for managing the exponential growth in traffic and the stringent latency requirements of AI-native applications. Cisco is helping service providers navigate the transition to AgenticOps by providing the infrastructure, telemetry, and intelligence required to manage the massive complexity of modern networks. Cisco's products help service providers close "visibility gaps" in complex, distributed networks. Cisco's deep visibility tools allow AI agents to map the entire path of a data flow across public and private infrastructure to identify the root cause of latency or packet loss in seconds rather than hours.

AgenticOps also allows the network to act as a self-healing system. Cisco's AI-native tools enable the network to reroute traffic, adjust capacity, or reconfigure network nodes when the system detects performance degradation or an impending hardware failure.[13] This dramatically increases uptime and reliability for mission-critical services. AI-driven networks can also help deliver greater energy efficiency by continuously optimizing network operations through intelligent automation using real-time telemetry.[14] AgenticOps enables networks to dynamically adjust power usage, traffic routing, and resource allocation, allowing operators to power down underutilized components during low-demand periods and reduce energy consumption without compromising service quality.[15]

One of the most significant challenges for network operators is the talent gap in managing increasingly complex, software-defined networks. AgenticOps allows operators to automate repetitive, low-value tasks—such as ticket resolution, configuration updates, and routine maintenance. These tools also help close the workforce talent gap by lowering the barrier

---

[13] Masum Mir, "Leading the Next Era of Intelligent Connectivity," (Oct. 28, 2025). Available at: https://blogs.cisco.com/news/leading-the-next-era-of-intelligent-connectivity?dtid=osolie001456&utm_loc=ar&utm_lang=en.

[14] "Energy Management in the AI Era: Cisco's Digital Transformation Advantage," Available at: https://blogs.cisco.com/industries/energy-management-in-the-ai-era-ciscos-digital-transformation-advantage

[15] *Id.*

to entry and allowing more junior analysts to ramp up quickly. By automating these tasks, Cisco's AI-enabled platforms can free network engineers to focus on higher-level architectural strategy and innovation, and free cybersecurity analysts to dedicate more time to strategic threat hunting and detection engineering.

Resiliency and Security: In an AgenticOps environment, the network transitions from a passive transport layer to the primary cybersecurity enforcement point. Cisco empowers service providers to implement AI-native security that operates at machine speed—a critical requirement in today's dynamic threat landscape.[16] By leveraging Cisco's purpose-built Deep Network Models, we analyze traffic patterns in real time to distinguish between legitimate AI-driven workloads and malicious anomalies.[17] This predictive capability allows our customers to identify vulnerabilities before they escalate. Should a device or service be compromised, the network can autonomously isolate the threat, halt lateral movement and safeguard the broader service provider infrastructure against sophisticated cyberattacks.[18]

Moreover, AI-native architectures enable network operators to simulate and validate cybersecurity patches on a digital twin of the network.[19] This capability eliminates the "patching paradox"—where the fear of downtime prevents critical security updates—by ensuring that every change is tested for stability and performance at machine speed. To maintain essential human oversight in an AgenticOps environment, we enable network operators to leverage a collaborative interface that allows security engineers to visualize, discuss, and refine complex policies with AI assistance.[20] This human-in-the-loop approach ensures that while our network operations run at machine speed, strategic security decisions remain guided by human expertise and transparent, shared visibility.

AI-Enabled Network Applications: In addition to changes in traffic patterns, networks are moving toward AI-native platforms that will become the fabric of intelligent connectivity rather than a simple pipe. As network operators move compute toward the network edge—such as at a cell site where a tower sits—they will be able to run applications directly from the network.

One promising application is Integrated Sensing and Communication (ISAC), which combines wireless communications and radio-frequency sensing to "see" objects' position and path using radio waves that reflect off them. Unlike optical sensors, it can detect intrusion even in low-light conditions, through smoke, or around obstructions where traditional video analytics might fail. This technology has been prototyped and demonstrated already, and it holds great promise for autonomous systems and robotics, AI-driven smart facilities, and public safety.

---

[16] "Shields Up: Cisco Live Protect Closes Vulnerability Gap with Compensating Controls," June, 2, 2026. Available at: https://blogs.cisco.com/news/shields-up-cisco-live-protect-closes-vulnerability-gap-with-compensating-controls

[17] "Cisco Hypershield – Our Vision to Combat Unknown Vulnerabilities," Available at: https://blogs.cisco.com/security/cisco-hypershield-our-vision-to-combat-unknown-vulnerabilities.

[18] *Id*.

[19] *See*, https://www.cisco.com/c/en/us/products/collateral/security/hypershield/hypershield-so.html#Solvingrealcustomerchallenges.

[20] *See*, https://blogs.cisco.com/ai/ai-canvas-controlled-availability.

Robotics: For example, ISAC enables the network to provide collision avoidance and path optimization, offloading some of the computational burden from the robot's own onboard sensors and ensuring safer operations in complex industrial environments.

AI-Driven Smart Facilities: By sensing occupancy and environmental changes through ISAC, the network can dynamically adjust energy, HVAC, and lighting systems in enterprise campuses. This creates "self-aware" buildings that optimize for both human comfort and energy efficiency.

Public Safety: Because the network can act as a pervasive radar system to track objects in restricted areas, ISAC could support Counter-Unmanned Aircraft Systems missions leveraging commercial networks around a stadium or critical infrastructure.

Cisco and its partners are already exploring this technology, having demonstrated an early version at NVIDIA's GTC DC Conference last year and at Mobile World Congress Barcelona this spring.

The network will also enable 'Physical AI', which is the integration of digital intelligence with physical action that allows machines to operate independently in environments where conditions change continuously.[21] While traditional AI is confined to software-based analysis, physical AI is designed to interact with the physical world. Physical AI applications—whether a warehouse robot rerouting around an unexpected obstacle, robotics in a manufacturing facility, or intelligent transportation infrastructure—will drive significant demand and constraints on the network while also opening new business monetization opportunities for network operators. Demand is driven by several factors: 1) data collection from many types of sensors such as cameras for visual detail, lidar for depth, and radar for motion detection; 2) decision-making and reinforcement learning (where the model improves by attempting physical actions and receiving feedback to improve future action); and, 3) performing physical action through mechanical components like motors, conveyors, or robotic arms.[22]

Each of these applications represents how AI is changing the network design: they require compute at the edge. They require an AI-native network core capable of integrating terrestrial and non-terrestrial networks and they require additional capacity to manage east-west traffic as agents, devices, and networks work in concert. Ultimately, Cisco envisions a combination of interoperable technologies to ensure ubiquitous, seamless, secure, and reliable connectivity for all end users, regardless of whether communications travel on wired, licensed, unlicensed, or satellite networks.

4. Policy Considerations.

As this Committee considers the changing landscape, I offer three suggestions that will help enable continued U.S. leadership and better outcomes for American consumers and businesses.

---

[21] *See*, https://www.cisco.com/site/us/en/learn/topics/artificial-intelligence/what-is-physical-ai.html.
[22] *Id*.

First, we must work together to protect American leadership in AI and compute and remain focused on building an American AI-native tech stack for wireless networks. However, we cannot wait for 6G standards to begin deployment; and we aren't. Through the AI-WIN partnership, Cisco, NVIDIA, MITRE, Booz Allen Hamilton, ODC, and T-Mobile are collaborating to develop America's first AI-native wireless telecommunications stack, while also developing innovative technologies to support new use cases in 5G-Advanced networks. We seek to leverage American leadership in AI and compute to demonstrate technical success through real-world deployments, ultimately shaping international standards to meet American companies where the technology is moving. I encourage Congress to lean in on areas where the United States has a strategic leadership role, such as compute, core networking, and applications.

Second, as compute infrastructure becomes more distributed and moves closer to the network edge, service providers may face increased infrastructure demands. Policies that enable the responsible and efficient permitting and deployment of this infrastructure will be essential to helping service providers offer new services as the network becomes an AI-native platform. Additionally, as Congress considers the future of the Universal Service Fund, we encourage you to consider how these changes in network traffic will affect the operational costs borne by rural health care facilities, service providers, and schools, regardless of whether they are rural or urban.

Finally, I encourage this Committee to continue its bipartisan approach to a balanced spectrum policy. As WRC-27 approaches, the United States needs to maintain its global leadership and expand its support for unlicensed operations in the 6 GHz band, including by authorizing higher-power operations and allowing expanded use in key enterprise environments. Likewise, the 800 MHz of spectrum made available for mobile use will be essential to enabling the connectivity needed to power new applications on AI-native networks, including the volume of uplink traffic we expect to see.

5. Conclusion.

We are at an exciting inflection point in how AI is transforming our communications networks. As Cisco works with service providers, critical infrastructure operators, enterprises, and governments to deploy the foundational infrastructure of the AI era, we are proud to offer our expertise to this Committee and to Congress to help shape a technology future that connects consumers, businesses, and communities. Thank you for the invitation to testify, and I look forward to answering your questions.

AI Impact on Wide Area Networks

Cisco Report 2026

About the report

The AI Impact on Wide Area Networks report represents the key findings and takeaways from a research initiative led by Cisco Systems, Inc. The report combines real-world traffic analysis (using the assurance tool Cisco Crosswork Assurance User Experience), third-party industry data, and Cisco controlled lab tests of AI agents. It explores how AI-driven traffic differs from traditional non-AI web traffic, with a focus on the impact of agentic AI traffic and the implications for network design, capacity planning, and assurance.

What sets the report apart is it's based on real-world traffic data including an early lens on agentic AI traffic (currently small but growing fast) that lets us start to see and measure a new class of AI network traffic and understand the implications. The report includes:

- Direct measurement of live AI inference traffic across service provider networks
- Empirical tests of AI traffic characteristics to train models that identify and track AI flows with precision
- A repeatable measurement framework and baseline to track AI traffic evolution and forecasting on an annual basis

Executive summary

The rapid adoption of artificial intelligence (AI) —particularly **AI inference** and **agentic AI systems**—is beginning to reshape network traffic patterns in measurable and structurally important ways. By 2027, 80% of executives believe their company's competitive survival will depend on agentic AI,[1] and consumer usage of AI is already prevalent and accelerating. This is driving a fundamental shift in how traffic is generated, distributed, and experienced, with implications for service providers and enterprises that manage large-scale networks.

Key Findings

The future of connectivity is no longer just about faster downloads—it's about the vital "spinal cord" of intelligence between agents and AI models—a critical dependency that's forging new paths of communications that impacts how we manage networks. While the industry has spent decades optimizing the network for human-paced, bursty video streams, the rise of agentic AI is changing network traffic profiles and behavior. By 2035, one-quarter of network traffic is projected to be AI inference (Source: Cisco model). These flows don't behave like the web. They live longer, demand more upstream capacity, and operate at software speed, not human speed. The connectivity between agent logic and AI models effectively becomes the agent's "spinal cord"—a critical dependency whereby any network degradation directly impairs agent functionality.

[1] Cisco report The Race to Agentic AI: Why Infrastructure will make or break workforce integration https://www.cisco.com/c/dam/en/us/solutions/collateral/artificial-intelligence/race-to-agentic-ai-report.pdf

Token consumption data

▲ **10x**

year-over-year growth

AI inference traffic is small today—but is growing extremely fast.

While AI inference traffic remains negligible compared to dominant categories like video streaming, observed growth rates are exceptional. **Token-consumption data shows nearly 10x year-over-year growth**, while in some service provider measurements we are seeing around 4x growth in just eight months. Sustained growth at these rates means AI traffic will become a meaningful component of overall network traffic by 2035.

AI inference flows

▲ **2x**

longer

Flow rate for web vs. AI traffic

▲ **10x**

larger

AI inference traffic flows behave differently at the transport level vs. typical web transactions.

- AI inference flows last approximately **2x longer**, mainly due to token-by-token content generation.
- Median flow rates of **regular web transactions are 10x larger than AI inference traffic flows** which have smoother sustained throughput, rather than bursty delivery.

These characteristics will require **flow-aware network and security systems**, including firewalls, intrusion detection, and Deep Packet Inspection (DPI), to handle greater scale, distribution, and state efficiency over time.

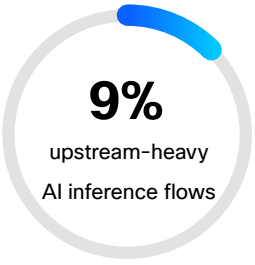

Traffic symmetry status quo is changing.

AI inference—and especially agentic AI—introduces significantly more upstream traffic due to large, context-rich prompts.

- Approximately **9% of AI inference flows carry more upstream than downstream traffic**, versus about 0.5% for typical web traffic.

This trend is expected to intensify as agents maintain and transmit growing contextual state, with direct implications for **access planning and upstream capacity modeling**.

#s ~ 5 seconds for AI inference latency vs. ~ 20–50 milliseconds for network latency

Network latency is not the dominant bottleneck—yet.

End-to-end AI inference latency is driven primarily by model processing, typically ranging from hundreds of milliseconds to several seconds. Network latency is currently small enough and does not yet justify large-scale architectural shifts toward edge inference purely for latency reasons. Inference hardware improvements will lead to network latency becoming more relevant, and we are already seeing proof points of that. Measuring and assuring the AI inference experience is becoming critical for perceived service quality.

AI agents

▲ **450%**

more total traffic

70%

of that traffic is AI inference

Agentic AI fundamentally amplifies network usage.

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.
- Approximately **70% of that traffic is AI inference**, making inference paths mission-critical.
- The connectivity between agent logic and AI models effectively becomes the agent's **"spinal cord"**—a critical dependency whereby any network degradation directly impairs agent functionality.

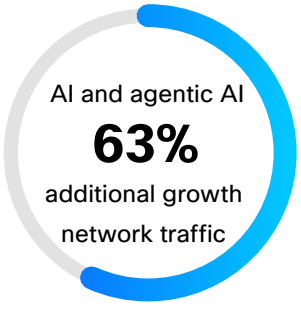

AI traffic impact by 2035*

Enterprise networks:

Without agentic AI, enterprise traffic is projected to grow by about 2.5x over the next decade. With agentic AI adoption, traffic growth could increase by about **9x**, driven by autonomous task execution and inference-heavy workflows.

Consumer networks:

Consumer AI adoption is accelerating and will reach near-universal usage by the mid-2030s. AI and agentic AI are projected to increase consumer-driven network traffic by about **6.6x**, representing about 63% additional growth compared to non-AI scenarios—making this the dominant driver of overall internet traffic expansion. AI inference traffic is projected to be 25% of total network traffic by 2035.

*Source: Cisco model

Takeaways

- AI and agentic AI will not just increase traffic volume—they will **change traffic shape**, symmetry, duration, and criticality.
- AI inference paths will become **strategic network assets**, requiring high levels of resilience, observability, and differentiated treatment, for example, Quality of Service (QoS) and path security.
- Traditional planning assumptions (burstiness, downlink dominance, human-paced interactions) will need to adapt to **the new reality of AI and agentic AI traffic** that lasts longer, demands more upstream capacity, and operates at software speed.
- Continuous measurement of **AI-specific Key Performance Indicators (KPIs)** and **adaptive network architectures** will be essential to sustain expectations for performance and user experience.

This report marks the first step in ongoing research on the impact of AI on network traffic. As AI adoption deepens across enterprises and consumers, future research work will refine projections, validate emerging patterns, and guide operators in evolving their networks for an AI-driven world.

Introduction

Adoption of AI and AI-based applications such as AI agents is growing at a very fast pace. In a Cisco report with Omdia, 80% of executives view agentic AI as critical to company survival by 2027[2]. Other industry reports confirm similar growth projections for both enterprises and consumer AI applications. In this context, service providers and organizations running large-scale networks need to evaluate how the proliferation of AI is going to impact networks in terms of traffic growth, patterns, or other characteristics.

This report is the first in an annual series where we will share key findings and implications to help network leaders make decisions on the evolution of their network infrastructure. Cisco has initiated research on AI traffic patterns and how they differ from non-AI web-based transactions, with a specific lens on agentic AI traffic and applications. The data presented in the report is based on real traffic data analysis from Cisco Crosswork Assurance User Experience, Cisco empirical tests, and third-party reports.

[2] Cisco report The Race to Agentic AI: Why Infrastructure Will Make or Break Workforce Transformation. https://www.cisco.com/c/dam/en/us/solutions/collateral/artificial-intelligence/race-to-agentic-ai-report.pdf

Data analysis and findings

Data volume growth

The first questions network operators are trying to answer are “How much AI traffic will exist?” and “How will this traffic impact our network?”

AI traffic, in this case, refers to AI inference traffic which flows across the WAN between the device requesting model inference and the data center or location where the model is hosted and the inference happens. Inside data centers, terabytes of traffic flows are happening during AI model training, but that traffic is not crossing the WAN, except for large data transfers for AI training purposes.

The actual volume of AI inference traffic is still negligible compared with other types of traffic, such as video streaming, and will likely remain very low in relative terms in the near term. However, there are very significant growth rates:

- Token consumption growth of 10x year-over-year in OpenRouter AI from December 2024 to December 2025.)[3]

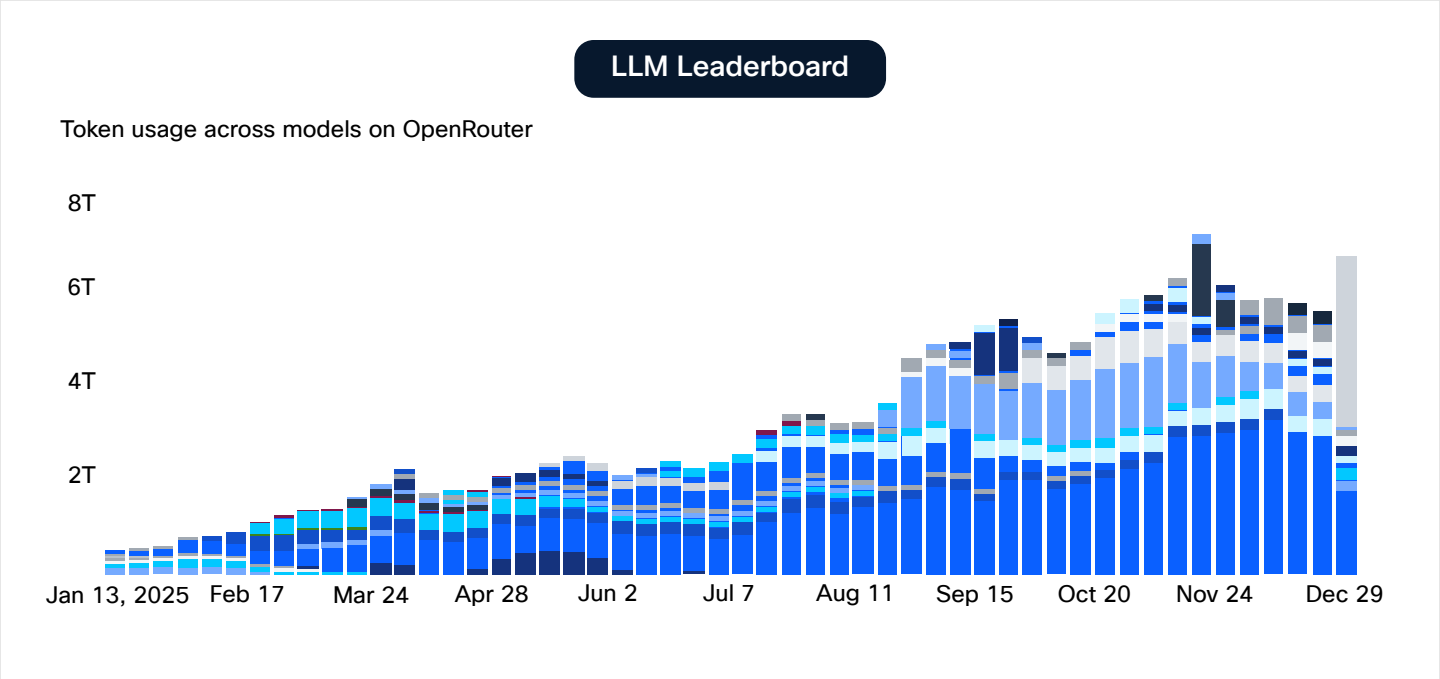

Figure 1. OpenRouter token consumption CY2026

Source: Openrouter.ai

[3] OpenRouter, LLM Leaderboard. https://openrouter.ai/rankings#leaderboard

Traffic analysis with data from two service providers shows consistent traffic growth for AI inference over a similar period: **4x traffic growth over a period of eight months** of traffic monitoring.

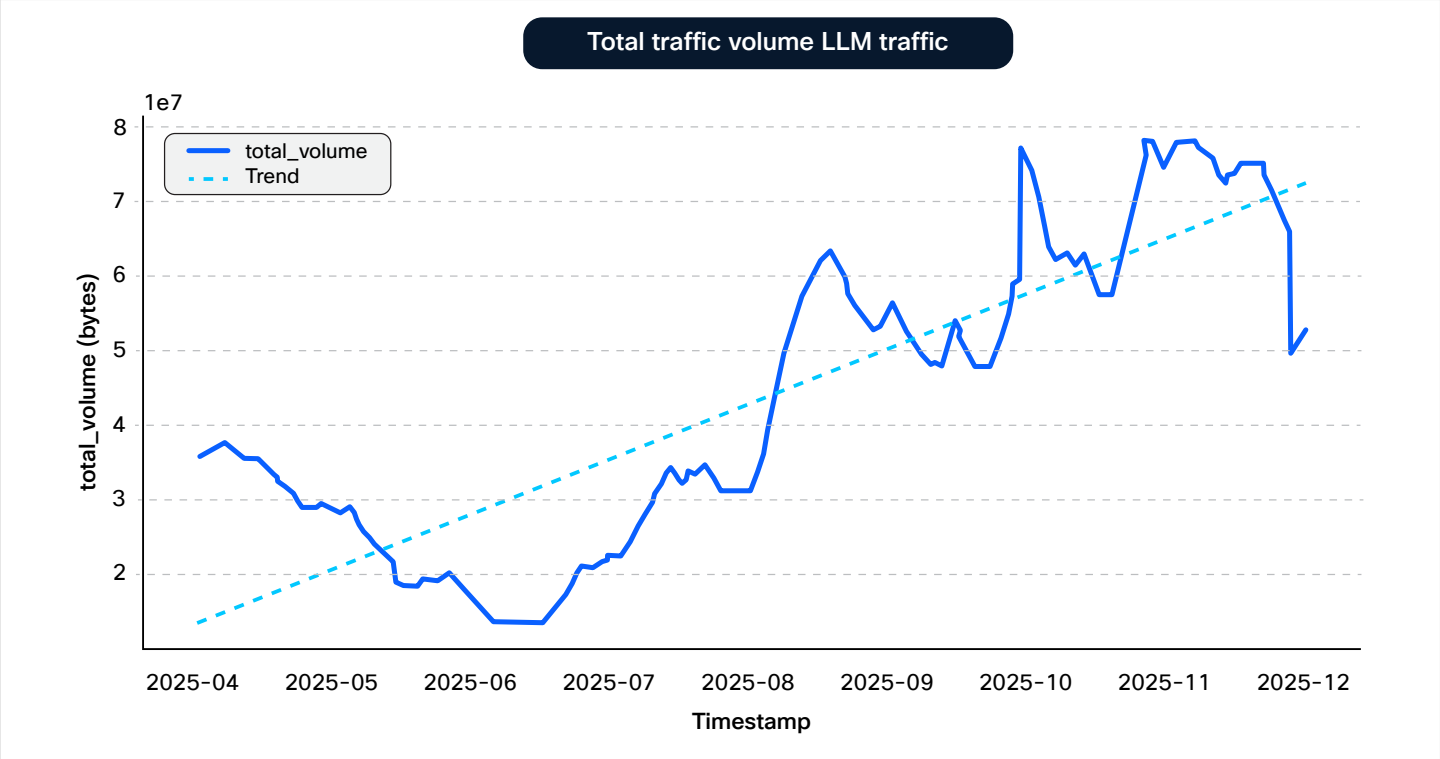

Figure 2. AI inference traffic volume evolution

Source: Cisco data

AI inference services are using both Transmission Control Protocol (TCP)- and Quick UDP Internet Connections (QUIC)-based transport. The split in terms of number of flows is almost 50%, while in terms of data volume QUIC represents 57%.

Implications: While AI inference traffic is still low in absolute terms, sustained high growth rates will result into a relevant portion of all traffic over time. Clearly, the adoption of AI by applications and users is driving fast growth. The adoption of QUIC as transport for AI inference presents encryption challenges for DPI systems to detect traffic at the application layer.

Data flow length

AI inference flows present different characteristics compared to non-AI web traffic. While not dramatic, these differences may impact capacity planning. Analysis of data flow length shows differences between AI inference flows versus regular web transactions.

Statistically, AI inference flows last 2x longer than regular web transactions.

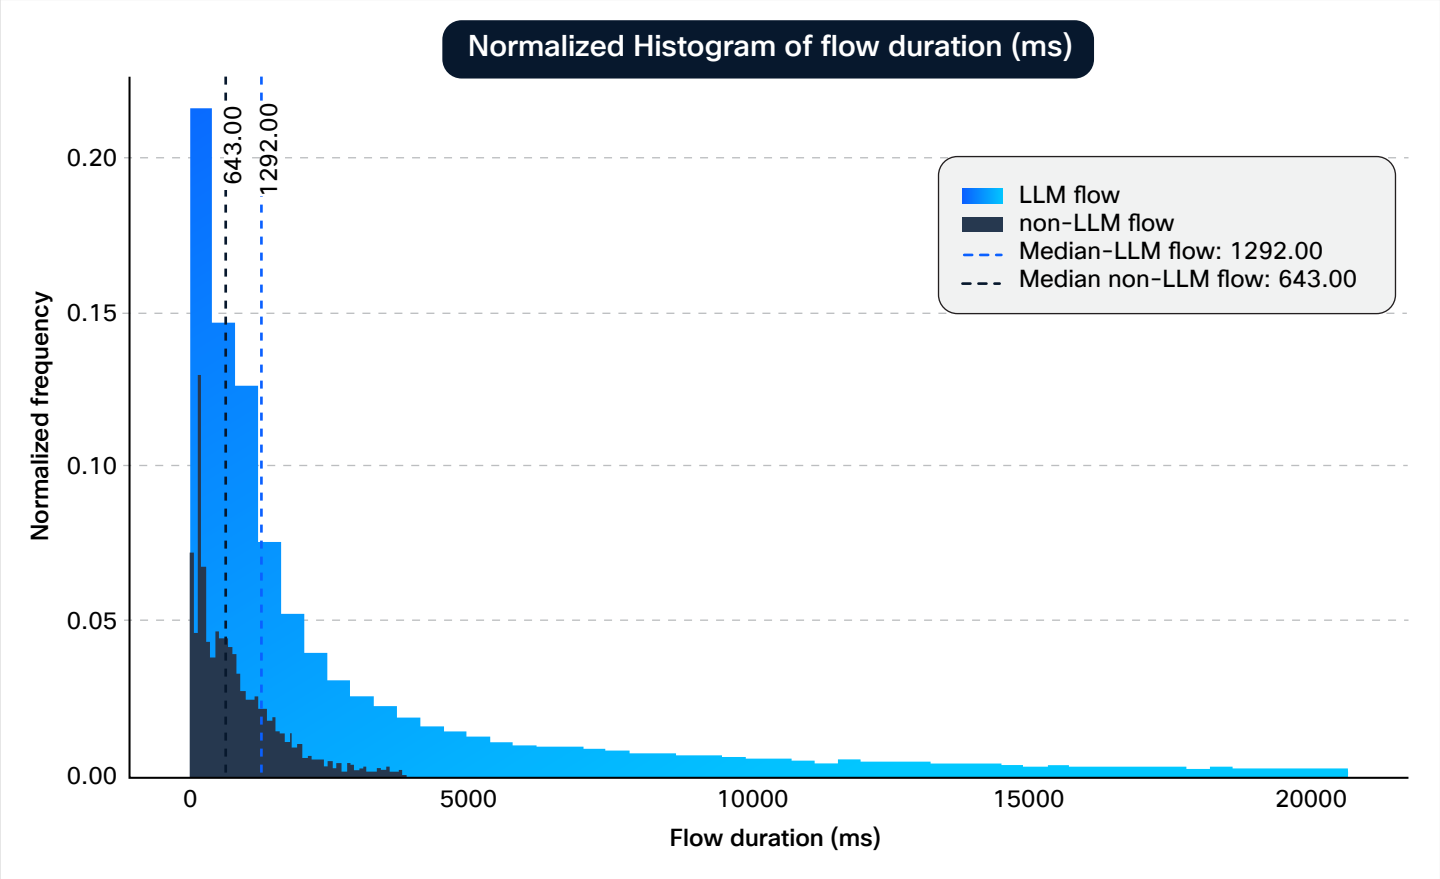

Figure 3. Flow duration distribution of AI inference flows vs. non-AI web transactions

Source: Cisco data

The main driver is the way AI inference traffic generates content, one token at a time, resulting in longer and lower rate flows compared to other flows.

Implications: For "flow-aware" network systems that must keep state for flows in tables, the proliferation of AI inference flows that last longer means growing flow tables will need to be effectively planned. Over time, security and flow-aware network systems are likely to become more distributed to cope with forwarding state growth, similar to the approach used by Cisco Hypershield™.

Data flow rate

Related directly to the flow length, the data flow rate for AI inference flows shows a different pattern compared to regular web transactions. **Median flow rate is 10x larger for regular web transactions compared to AI Inference flows**.

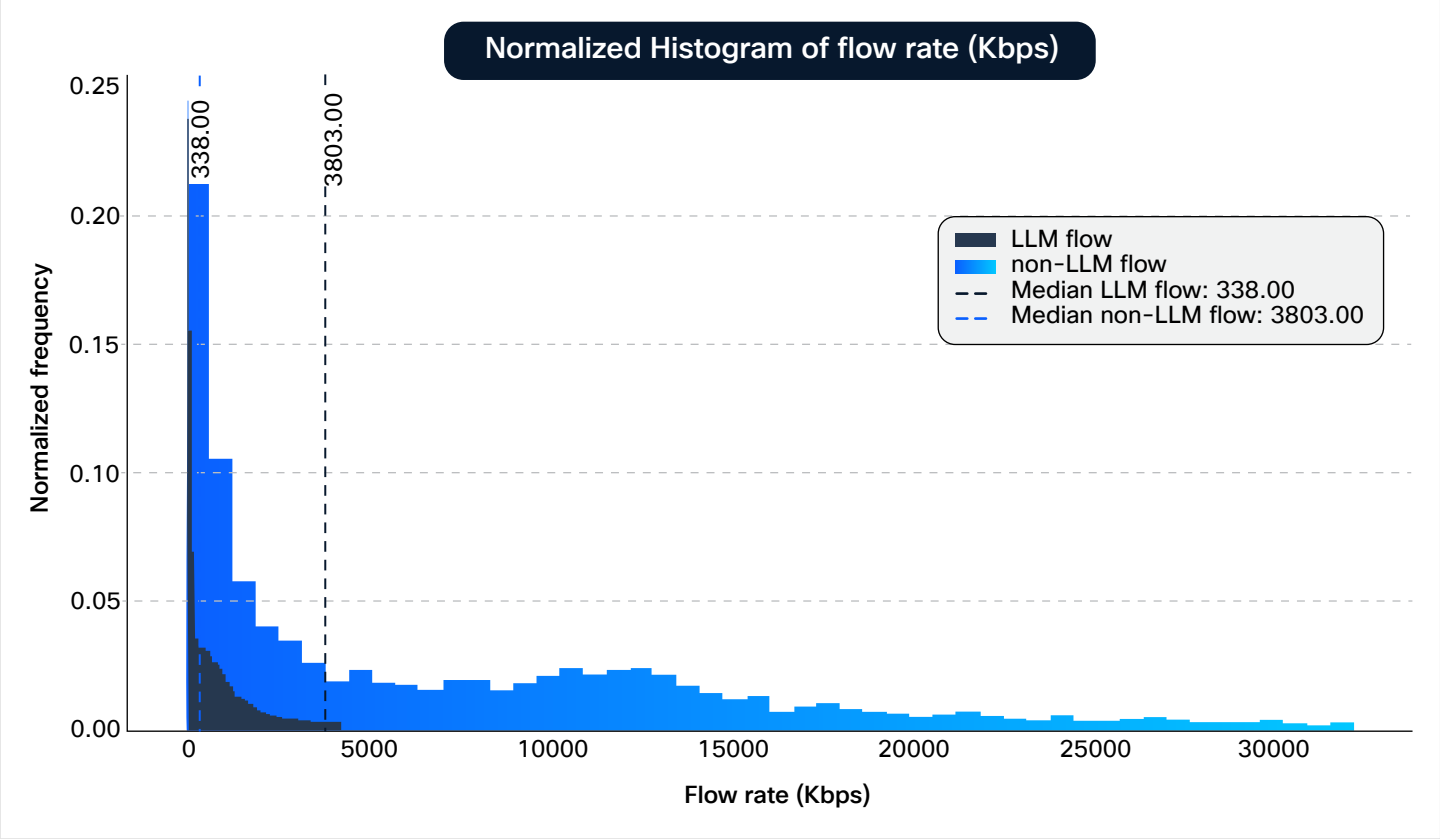

Figure 4. Flow rate distribution of AI inference flows vs. non-AI web transactions

Source: Cisco data

The main reason for this difference is the process that generates the data in AI inference. A data flow that occurs **one token at a time** is like having a "shaper" applied to the data flow. Regular web traffic flows peak much higher, as the content can be retrieved from wherever it is stored and delivered to the user.

Implications: Different median average-traffic and peak-to-average rates may require different QoS settings in the network to manage AI inference traffic versus non-AI inference web transactions.

Traffic asymmetry

Potential changes in traffic asymmetry are expected with AI inference traffic. Network traffic asymmetry varies depending on access type (mobile vs. wireline networks) and services. Mobile network traffic tends to be more upstream heavy due to social networking, with more content sent upstream compared to wireline networks.

Analysis of AI inference flows versus non-AI web transactions shows clear differences in traffic symmetry. As Figure 5 shows, in AI inference traffic, **9% of flows have more upstream traffic than downstream traffic** compared to other HTTP transactions, where this occurs in only about 0.5% of the flows.

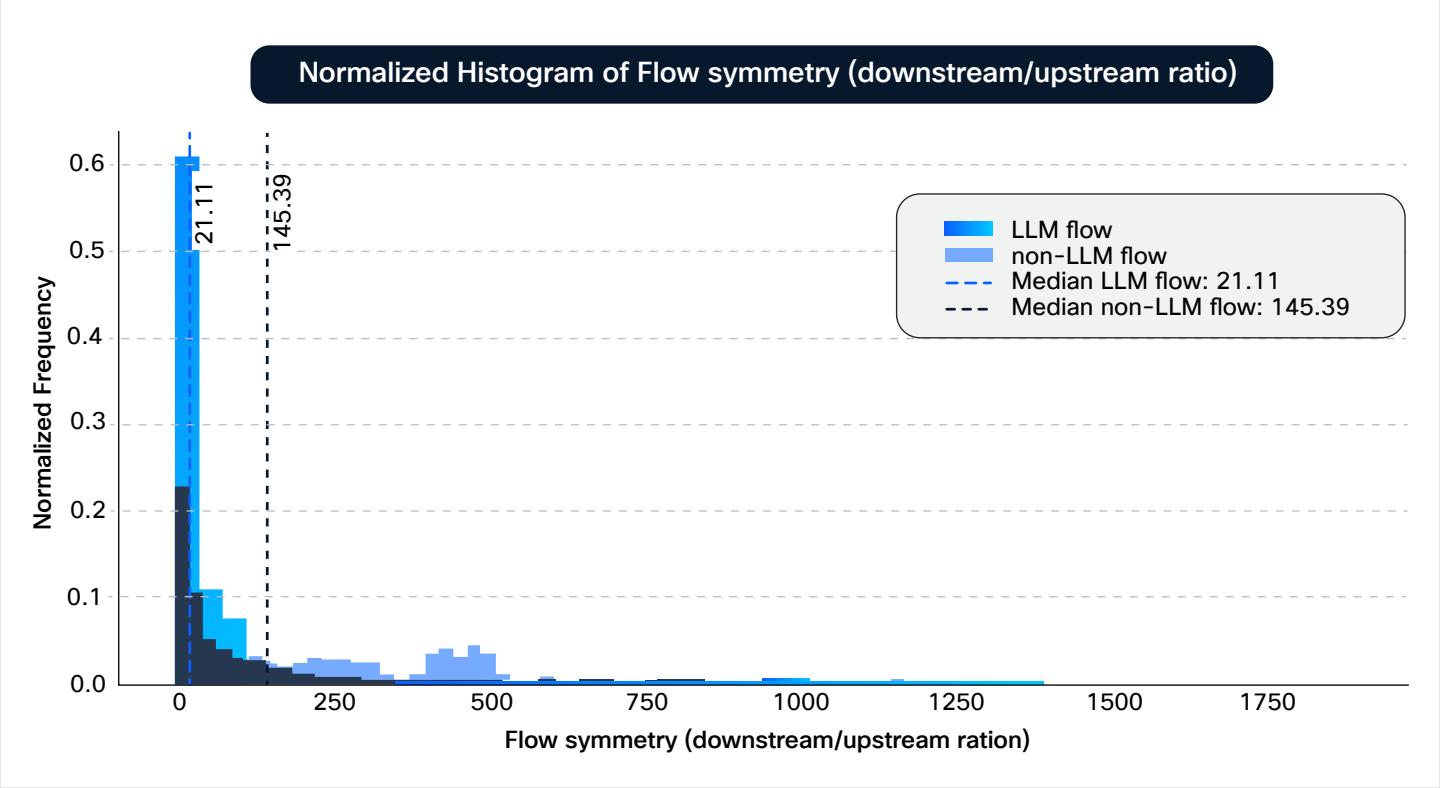

Figure 5. Traffic flow symmetry of AI inference flows vs. non-AI web transactions

Source: Cisco data

We foresee this trend further **consolidating and growing**. The reason for this difference is the fact that AI inference requests often capture a growing amount of context in the prompt in order to produce a relevant response. As adoption of agentic AI grows, more significant changes to network traffic symmetry patterns can be expected. Agents will capture continuous context as tasks are executed, and this state needs to be inserted into the prompts.

Implications: Particularly for radio capacity planning, traffic symmetry assumptions are a relevant factor. As adoption of AI inference and agentic AI grows, it will be important to track the evolution of traffic symmetry, as it will continue decreasing over time.

Network latency impact

Large Language Model (LLM) inference requests tend to have significantly higher latency than typical web API calls, and the response times can be more variable. Traditional web application REST APIs often strive for subsecond or even sub-100-millisecond (ms) response times. For instance, **a good API response time is typically below 100 ms** in many web scenarios. By contrast, even short LLM queries incur response times of hundreds of milliseconds just to begin producing an output, and full responses often take **seconds**. LLM latency can be modeled as follows:

Latency = Time To First Token (TTFT) + Time Per Output Token (TPOT) × Number of output tokens

TTFT is the initial overhead to process the prompt, and TPOT is the generation time per token. This means latency grows with output length. For example, OpenAI GPT-5.2[4] might take on the order of 9 ms per token, and Anthropic Claude Opus 4.5 about 20 ms per token. A **100-token GPT-5.2 answer could take about 1 second**, and a maximum-length 600-token answer could approach 5.4 seconds. There is wide variability among LLMs, model providers, implementations, and the hardware used for inference.

Implications: AI inference latency varies widely from a few hundred milliseconds to multiple seconds. Network latency is still at this point a very small component of the total end-to-end inference latency. Inference hardware is, however, evolving and becoming much faster. As inference accelerates, the network latency component plays a more relevant role on the end-to-end perceived experience. Network latency will become a key factor for inference distribution, combined with others like scale, data sovereignty, and security. Additionally Service Providers will need to monitor the actual AI inference latency that the customer is experiencing, as it's a key factor in perceived user experience. With Cisco® Crosswork Assurance, it is possible to measure key AI inference KPIs across different network endpoints, along with network and service KPIs.

[4] LLM Leaderboard—Comparison of over 100 AI models from OpenAI, Google, DeepSeek & Others. https://artificialanalysis.ai/leaderboards/models

Agentic AI traffic patterns

In addition to AI Inference flows, agentic AI as an AI application deserves a special focus due to its transformational impact and the implications for networks.

AI agents are software processes that use an **AI model as a brain** to make decisions. Then, depending on the type of agent, multiple tools are used to complete specific tasks or goals.

There are three main sources or destinations of data in an agent (see Figure 6):

- The agent itself (where the agent code runs)
- The AI model (where AI inference happens)
- The tools and/or sources of data for the agent

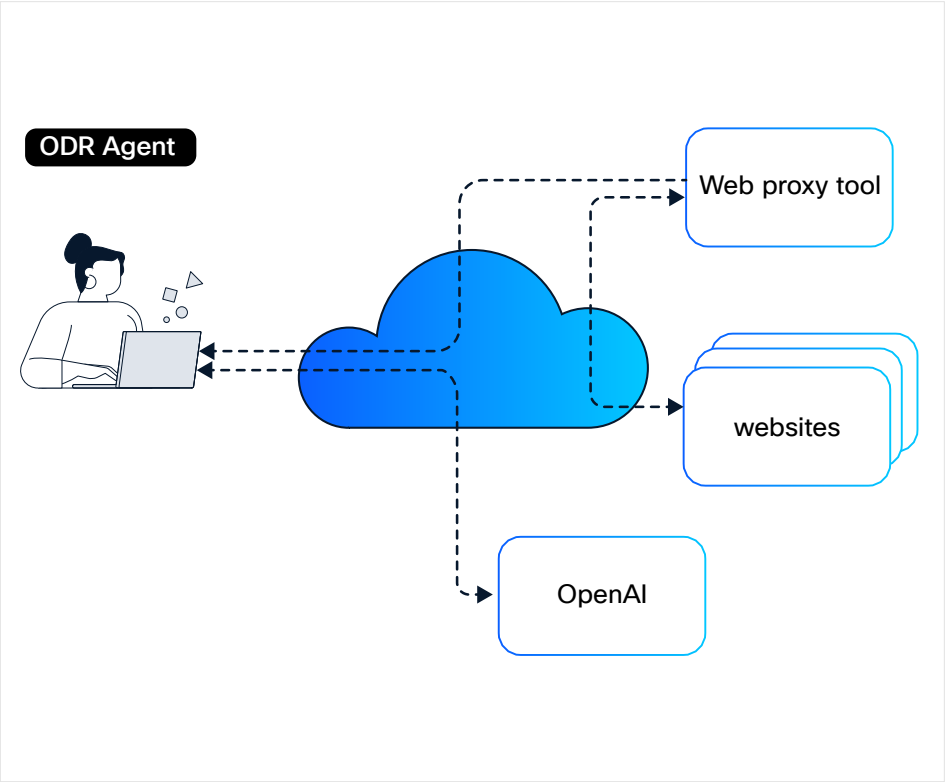

Figure 6. AI Agent testing environment

Source: Cisco

For example, we are using an Open Deep Research agent from Langchain5 (other agents will behave similarly). Let's analyze what happens when a given task is requested to the agent (see Figure 7).

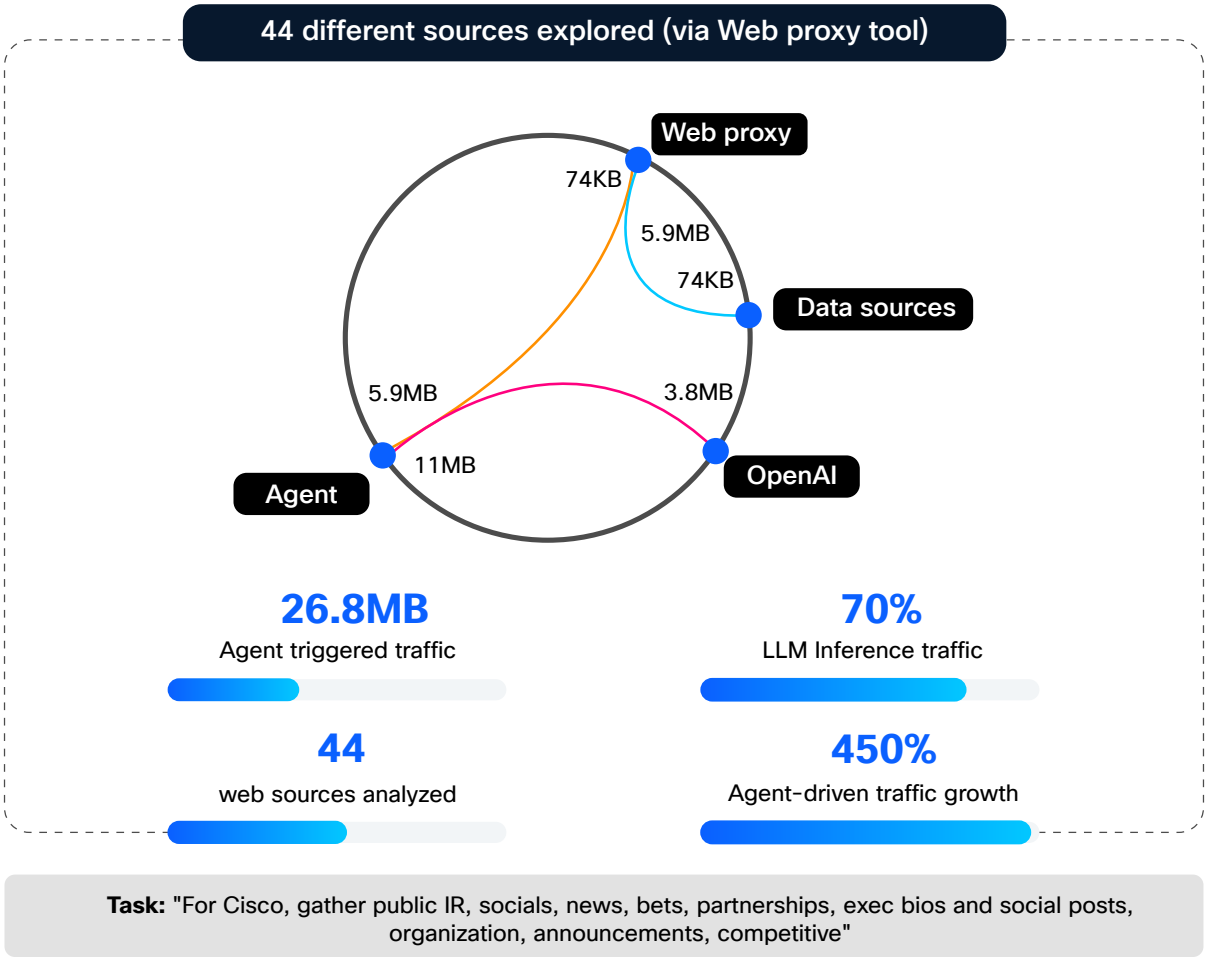

Figure 7. AI agent traffic flows

Source: Cisco

A request to an agent generates a significant amount of network traffic between the different components of the agent. Overall, the network traffic generated or triggered by the agent is much larger than the network traffic generated if the same task was performed manually by a human. In our example there is a **450% increase in traffic** associated with using an agent, and notably, **70% of that new incremental traffic is AI inference**.

The following are key considerations related to AI agent behavior and its interaction with and dependency on the network:

- AI agents act as network **power users**. Agents consume cloud and network resources at **software speed**, compared to when these tasks are performed manually at **human speed**. This results in a significant increase in network traffic.
- AI inference traffic is not only a significant portion, but the critical portion of traffic that is necessary for the agent to operate. The network path between the agent code and the AI model can be considered the agent's **'spinal cord'**. Any failure or constraint in that communications path will significantly impact the agent's ability to operate.

[5] Open Deep Research, https://github.com/langchain-ai/open_deep_research/tree/main

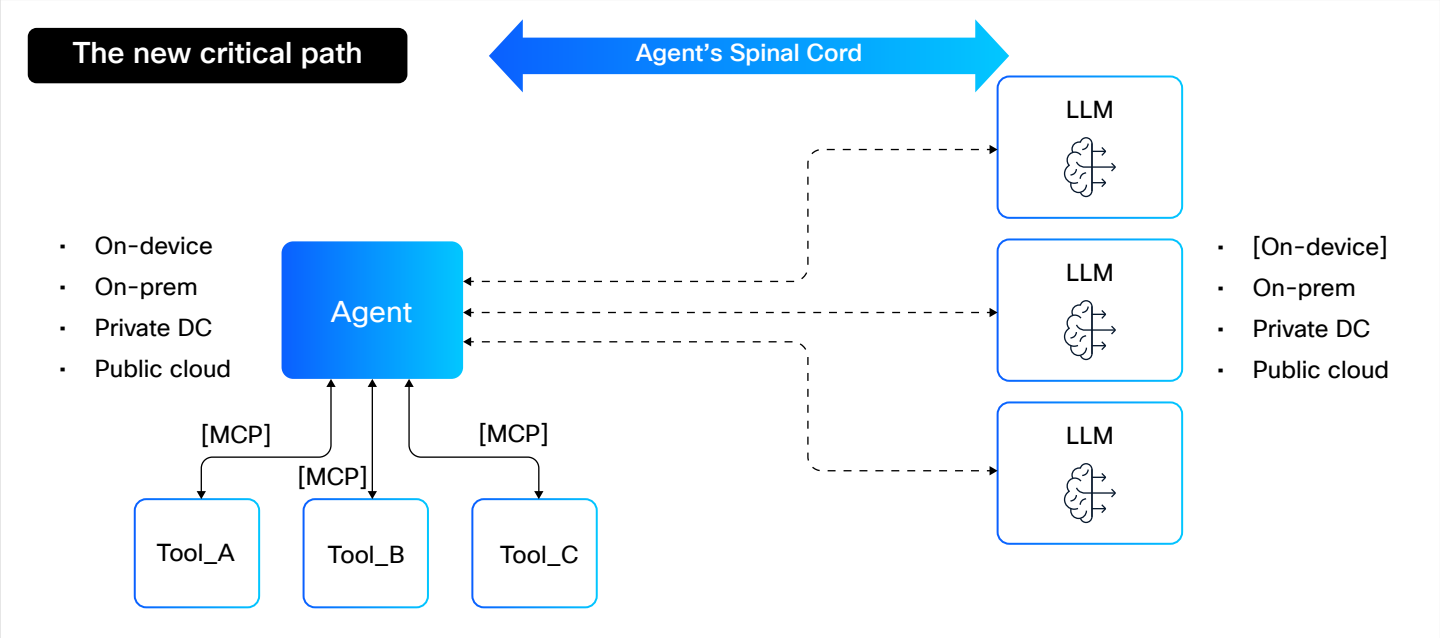

Figure 8. AI Agent traffic pattern and new critical path

Source: Cisco

Implications: As agentic AI adoption grows in both enterprise and consumer applications, network traffic will grow accordingly, primarily driven by the adoption of agents in user tasks. In that process, AI inference traffic will grow significantly as it is a large percentage of AI agent-driven traffic. Resiliency at the network level for communication paths is critical for AI inference across the agent 'spinal cord' between the agent code and the model.

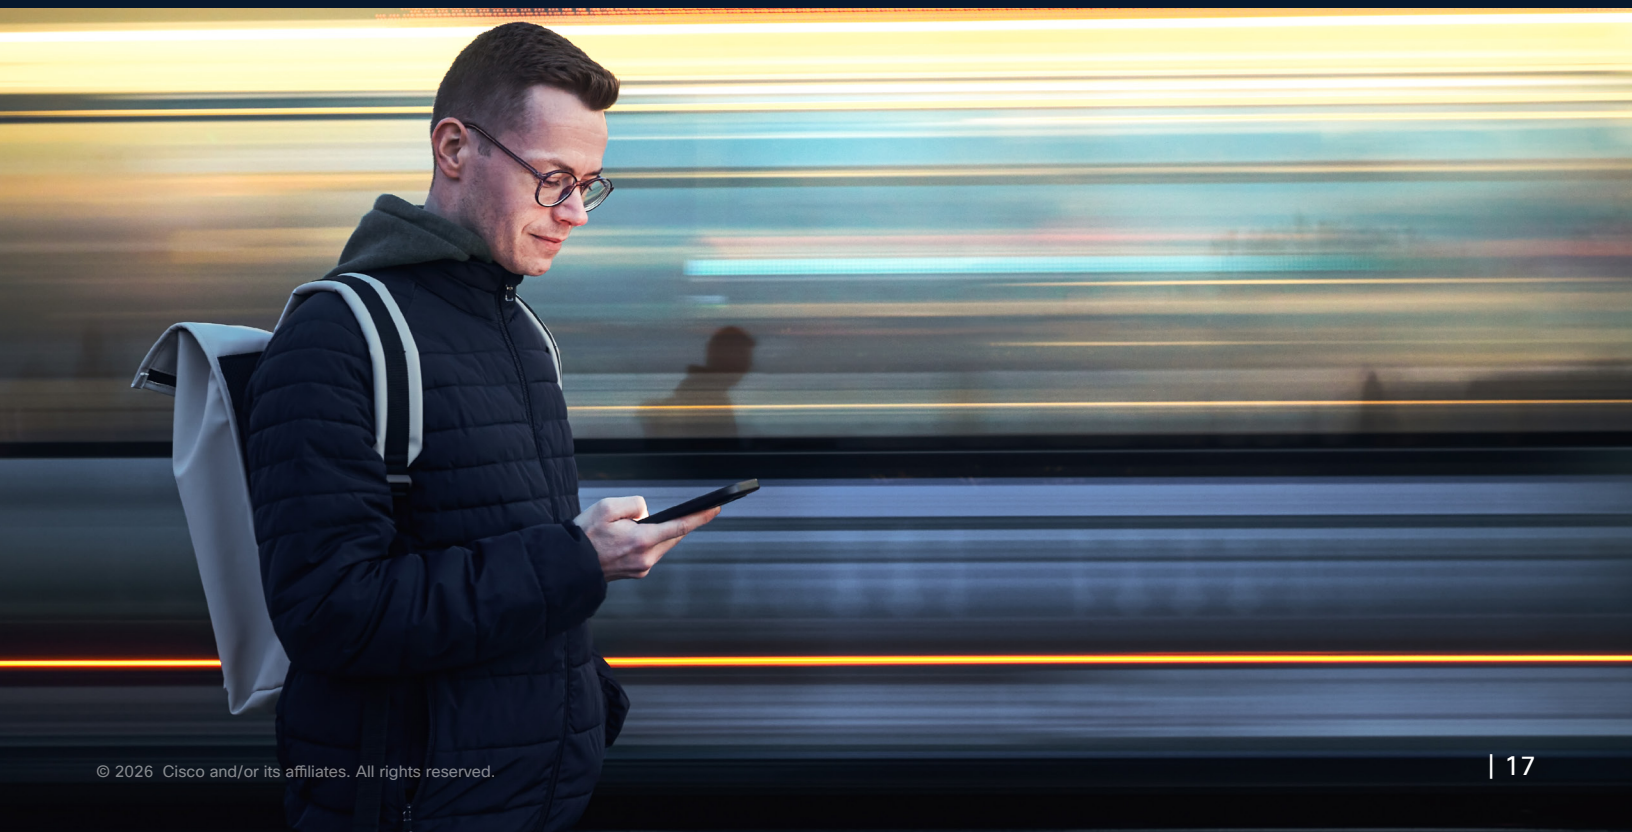

Looking ahead

This report explores the behavioral differences and special characteristics of AI inference traffic and agentic AI to understand the impact the traffic can have on the network.

Our next step is to envision what the future of network traffic may look like as consumer and enterprise adoption increases for these technologies.

This report provides projections through 2035. While long-term modeling involves inherent uncertainty, the objective is to support service providers and network operators with extended planning horizons. As additional data becomes available and models are refined, updated projections will be issued in future reports.

Enterprise usage of agentic AI and network impact: 2026–2035

Agentic AI, with autonomous, action-taking AI agents, is only at the early stages of enterprise usage, but momentum is rapidly building. Analysts note that adoption in 2024–2025 was nascent (fewer than 5% of enterprise applications had AI agents in 2025[6]), yet organizations across industries are already planning for widespread implementation. Key projections for the mid-2020s include:

By 2026

40%

of enterprise applications will include integrated, task-specific AI agents

- **By 2026:** Gartner forecasts that **40% of enterprise applications will include integrated, task-specific AI agents**, up from less than 5% in 2025. This reflects a rapid shift from simple chatbots toward more capable agents embedded in software tools (for example, AI-driven assistants in Customer Relationship Management [CRM], Enterprise Resource Planning [ERP], and IT service management applications).

By 2027

80%

believe their company's survival will depend on agentic AI.

- **By 2027:** Surveys show adoption accelerating year-over-year. A Cisco Report, shows that by 2027, 80% of executives believe their company's competitive survival will depend on agentic AI.[7] In IBM's 2025 global executive survey, 24% of business leaders said they already have AI agents taking independent action in their operations—and **67% expect to have AI agents autonomously making decisions in workflows by 2027.**[8] Even in regulated domains, executives anticipate that agents will automate significant work (for example, an average of **29% say they will automate risk and compliance tasks by 2027** in leading firms).

---

[6] Gartner press release, Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025.

[7] Source: Cisco report The Race to Agentic AI: Why Infrastructure will make or break workforce integration.

[8] IBM, Agentic AI's Strategic Ascent.

These early projections underscore that **agentic AI usage is moving from experimental to strategic.** Organizations are laying the groundwork through investments in generative AI platforms and automation, which will pave the way for deploying autonomous agents.[9] Cost reduction and productivity gains are key motivators for early adopters (39% cite these as primary drivers). By 2027, many enterprises will have progressed from planning and pilots into active use of AI agents in day-to-day operations.

The use of AI agents will become commonplace in business processes. Gartner analysts predict that by 2028, **"one-third of all user interactions with software will shift from traditional apps to agentic AI-driven interfaces,"**[10] illustrating how pervasive these agents will be in everyday workflows. In customer operations, Gartner expects that **80% of routine customer service issues will be handled autonomously by AI agents by 2029**.[11]

Looking further ahead, industry experts anticipate **continued acceleration of agentic AI adoption into the 2030s**, with these AI "co-workers" becoming ubiquitous across business domains.

Key forecasts and expectations for the mid-2030s include:

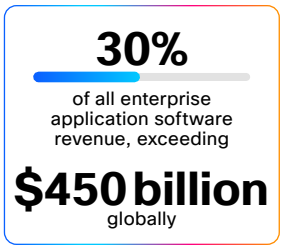

- **Enterprise Software Revolution:** Gartner projects that by 2035, agentic AI will drive approximately **30% of all enterprise application software revenue**, exceeding **$450 billion** globally (up from just 2% of software revenue in 2025).[12] In essence, nearly one-third of the enterprise software market might be attributable to AI agent capabilities by 2035 – a radical shift in a decade. At that stage, most enterprise software is expected to have AI agents deeply embedded, and new software business models will revolve around autonomous functionality.
- **New Normal of Work:** By 2035, the enterprise environment will likely feature humans and AI agents working side by side in most business processes. Gartner describes a democratized future state where employees routinely create or customize AI agents on-demand, and autonomous agents orchestrate complex workflows across departments.

[9] Omdia, New Omdia Analysis Shows Agentic AI Outpacing Growth Rates of Traditional Generative AI, September 2025.

[10] Gartner press release

[11] Gartner press release

[12] Gartner press release

By the mid-2030s **agentic AI is no longer viewed as an emerging technology, but rather as a standard component of enterprise operations**. Companies slow to embrace autonomous AI may find themselves at a competitive disadvantage. Overall, the consensus is that **the 2020s are the adoption decade for agentic AI, and the 2030s will be the era of its maturation and full integration** into the fabric of business.

Based on the data captured from industry reports, the following captures a projection model for agentic AI penetration in enterprise tasks from 2026 to 2035.

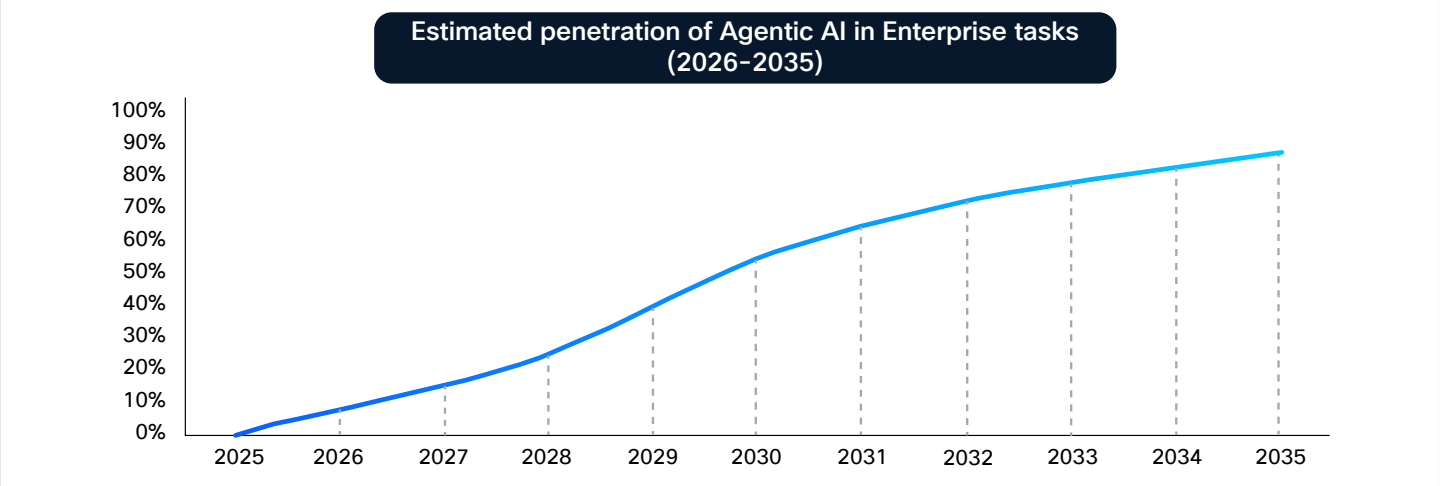

Figure 9. Projected adoption of agentic AI in enterprise tasks

Source: Cisco model

Considering the network implications of using AI agents to automate tasks and the forecasted agentic AI penetration, Figure 10 shows the overall impact on network traffic driven by enterprise services and agentic AI adoption from 2026-2035.

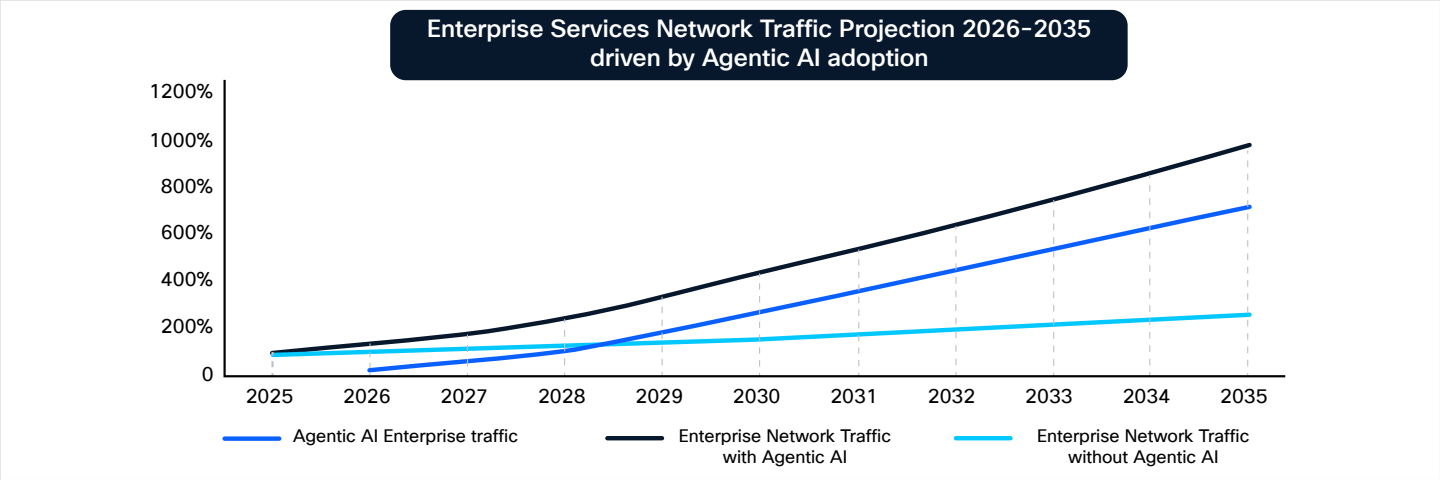

Figure 10. Projected impact of agentic AI in enterprise services network traffic (% increase)

Source: Cisco model

Without agentic AI, enterprise traffic would be expected to grow by about 250% over the period 2026-2035. However, **agentic AI will bring significant disruption and impact to enterprise network usage, driving 9x traffic growth.**

Consumer Usage of AI and Agentic AI and network impact: 2026–2035

AI is rapidly becoming embedded in everyday consumer life and is poised for an even more dramatic surge in adoption. Key drivers include **generative AI chatbots and assistants** (such as ChatGPT)**, AI in gaming and entertainment, AI-powered image and video generation, smart apps for shopping and banking,** and the rise of **augmented reality (AR)** experiences that rely on cloud AI inference. Agentic AI with autonomous AI agents that can act on a user's behalf is also emerging in personal applications.

61%

of adults in the U.S. had used AI in the past six months

Consumer adoption of AI reached a tipping point in the mid-2020s. Surveys in mid-2025 found that **61% of adults in the U.S. had used AI in the past six months,** with about **19% using AI daily.**[13] Scaled globally, this equates to roughly **1.7 to 1.8 billion people** worldwide who have accessed AI tools, including **500 to 600 million daily users**. This explosive growth has come just a few years after the advent of mainstream generative AI (ChatGPT launched in late 2022), indicating an unprecedented pace of consumer technology adoption.

46%

of all consumer transactions in the U.S. by 2030

By 2030, many analysts project that AI will be as common in daily life as smartphones are today. Gartner suggests that AI will be deeply integrated into daily activities; for instance, by 2028 a large share of business purchases may be handled by AI agents,[14] and the consumer realm is on a similar trajectory. Cognizant and Oxford Economics forecast that **AI-driven purchasing will account for about 46% of all consumer transactions in the U.S. by 2030**.[15] In other words, nearly half of consumer spending could be influenced or executed by AI (with similar figures of 46% in Germany and 55% in Australia).

By the early **2030s**, we can expect most consumers to

regularly rely on AI

By the early 2030s, we can expect **most consumers to regularly rely on AI**—either consciously (such as asking a chatbot or using an AI assistant) or behind the scenes (AI-powered features in apps, appliances, and vehicles). As AI becomes more trusted and integrated, occasional users today are likely to become daily users by 2030. Menlo Ventures analysts expect consumer AI adoption to ultimately **catch up to enterprise AI** in both usage rates and economic value, given the broader reach into the general population.

By **2035**, projections tend to converge on the assumption that AI will be nearly ubiquitous for anyone using connected technology. Although precise penetration percentages for 2035 are scarce, it's reasonable to expect **global consumer AI usage to approach close to 90% to 100% of internet users by the mid-2030s,** effectively meaning **billions of people interacting with AI daily**. AI will likely become a default layer in devices and services—much like broadband or electricity—largely invisible, yet indispensable.

---

[13] Menlo Ventures, 2025: The State of Consumer AI.

[14] Networkworld, AI's Dark Side Shows in Gartner's Top Predictions for IT Orgs, October 2025.

[15] Cognizant, Cognizant Study Shows Consumers Who Embrace AI Could Drive $4.4 Trillion in Spending Over Five Years, January 2025.

By 2030

23%

of Americans had made a purchase using AI in the past month

A major trend in this period is the rise of **agentic AI**, where autonomous AI agents can proactively perform tasks for users. These range from **personal shopping bots** that find and buy products for you to AI assistants that manage schedules, finances, or home devices with minimal oversight. Industry futurists predict that by 2030, personal AI assistants will act as consumer proxies—screening offers, negotiating deals, and automating decisions. Morgan Stanley reports that as of late 2025, roughly **23% of Americans had made a purchase using AI in the past month**.[16] The report expects **"agentic shoppers" to account for 10% to 20% of all e-commerce spending in the U.S. by 2030.**

By 2030

$3 to $5 trillion

of global retail revenue could be orchestrated by AI agents

Globally, the impact of agentic consumer AI could be enormous. McKinsey research forecasts that by 2030, **$3 to $5 trillion of global retail revenue could be orchestrated by AI agents** that anticipate needs and execute purchases autonomously.[17] This would be a seismic shift in how consumers find and buy products—a change on the order of the original e-commerce revolution but potentially occurring at a much faster rate. Agentic AI isn't limited to shopping; similar autonomous assistance includes personal finance (AI bots optimizing investments or budgets), travel (AI agents planning trips end to end), and even health (AI assistants scheduling doctor visits or managing medications).

By 2030

a majority of consumers worldwide will likely use AI assistants and tools as routinely as they use the internet

If even a fraction of that agent-mediated paradigm extends to consumers, by 2030, a typical individual might routinely rely on an AI agent to handle a host of daily tasks, from grocery replenishment to filtering incoming information. In short, **agentic AI adoption**, while nascent today, is slated to become a mainstream aspect of consumer life in the 2026–2035 timeframe, essentially giving people a "digital personal assistant" that is continuously working on their behalf.

The device landscape for consumer AI is also broadening. The smartphone remains king in sheer numbers, but AI is permeating our homes, wearables, cars, and emerging AR interfaces. In 2026–2035, this multidevice proliferation ensures that AI will be accessible almost anywhere, anytime. The net effect is that adoption rates will climb not just because of personal preference, but because the infrastructure and devices around us will be inherently AI-enabled.

Between 2026 and 2035, AI is set to transition from a novel technology to a universal utility in the consumer world. Both general AI usage and agentic AI adoption are projected to grow exponentially, reaching billions of people and touching all demographics globally. By 2030, **a majority of consumers worldwide will likely use AI assistants and tools as routinely as they use the internet**—with younger generations leading the way and older generations steadily following. Agentic AI, in the form of personal autonomous agents, is expected to handle an increasing share of our digital and physical tasks, from shopping to scheduling, heralding a new era of convenience (and raising new questions around security, trust, and control as well as bias in models).

[16] Morgan Stanley, Here Come the Shopping Bots, December 2025. https://www.morganstanley.com/insights/articles/agentic-commerce-market-impact-outlook

[17] QuantumBlack, AI by Mckinsey, The Agentic Commerce Opportunity: How AI Agents Are Ushering in a New Era for Consumers and Merchants, October 2025.

Physical AI will have a major impact on network infrastructure as robotics technology advances. Broader commercial adoption across industrial and consumer segments will increase reliance on AI inference. Many systems will combine embedded, real-time inference for immediate actions with cloud-based reasoning capabilities. For this analysis, robots can be viewed as physical agents equipped with tangible tools rather than purely digital ones. From a network standpoint, we expect the aggregate impact to be comparable to that of digital agents. This is a critical area to monitor closely in future reports.

Crucially, these adoption trends are not happening in isolation; they coincide with advancements in devices, such as AR wearables, and deeper integration of AI into every app and service. Therefore, consumer AI penetration can be visualized not as a simple linear increase, but as an accelerating curve approaching saturation. Reports and forecasts consistently point to **high double-digit growth rates in AI usage and market size through the 2020s.**[18] If those trends hold, by 2035 we will live in a world where **nearly everyone, across all age groups and regions, regularly relies on AI**—Whether through direct interactions or invisible assistance embedded in their environment. In effect, AI will be ambient: an ever-present copilot for the average consumer.

Based on the data from industry reports, Figure 11 shows a projected model for AI and agentic AI adoption in consumer tasks from 2026-2035.

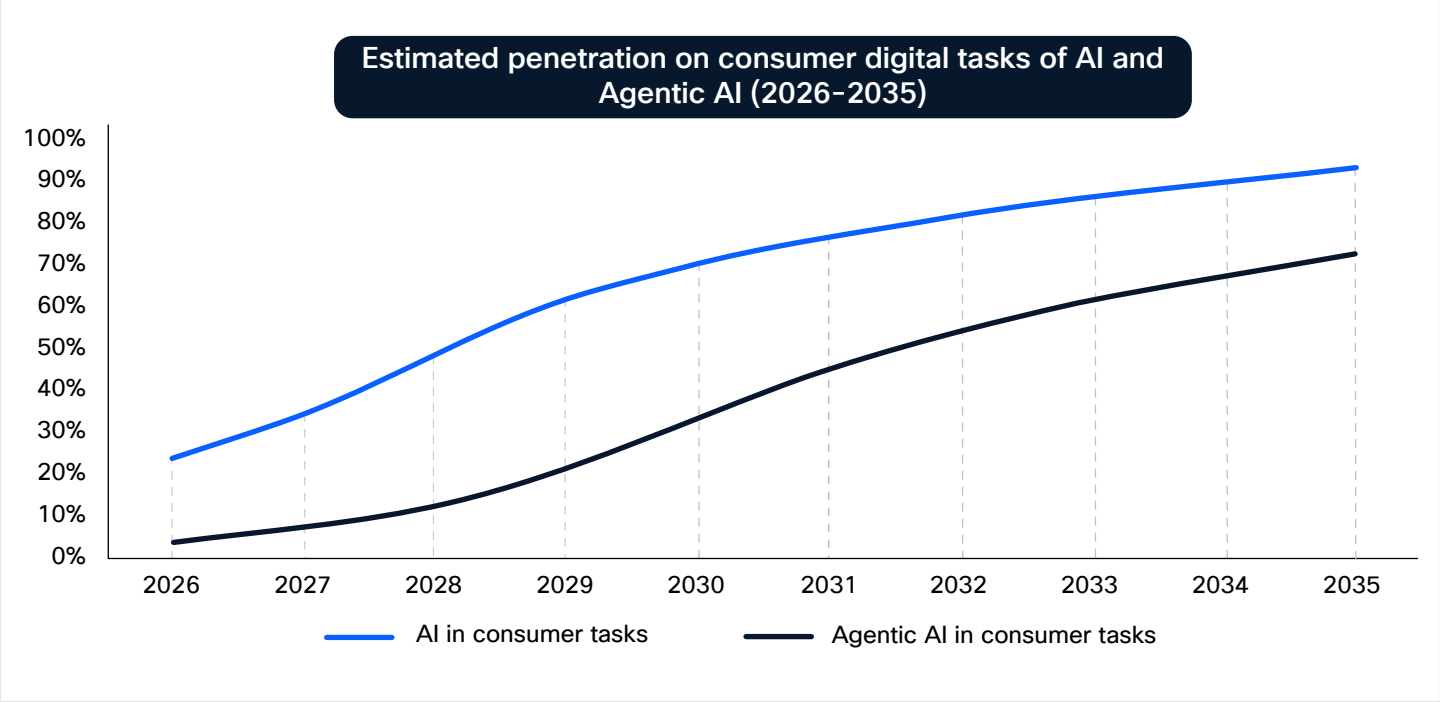

Figure 11. Projected adoption of AI and agentic AI in consumer digital tasks

Source: Cisco Model

[18] GlobeNewswire, Voice Assistant Application Market Report 2025-2035: Industry Revenues to Grow from $8.1 Billion in 2025 to $153.5 Billion by 2035, December 2025.

Not all consumer tasks or digital services have an equal impact on the network, and not all of them will be equally impacted by AI or agentic AI. The model in Figure 12 shows the aggregate impact on overall network traffic driven by the adoption of AI and agentic AI over the period 2026–2035.

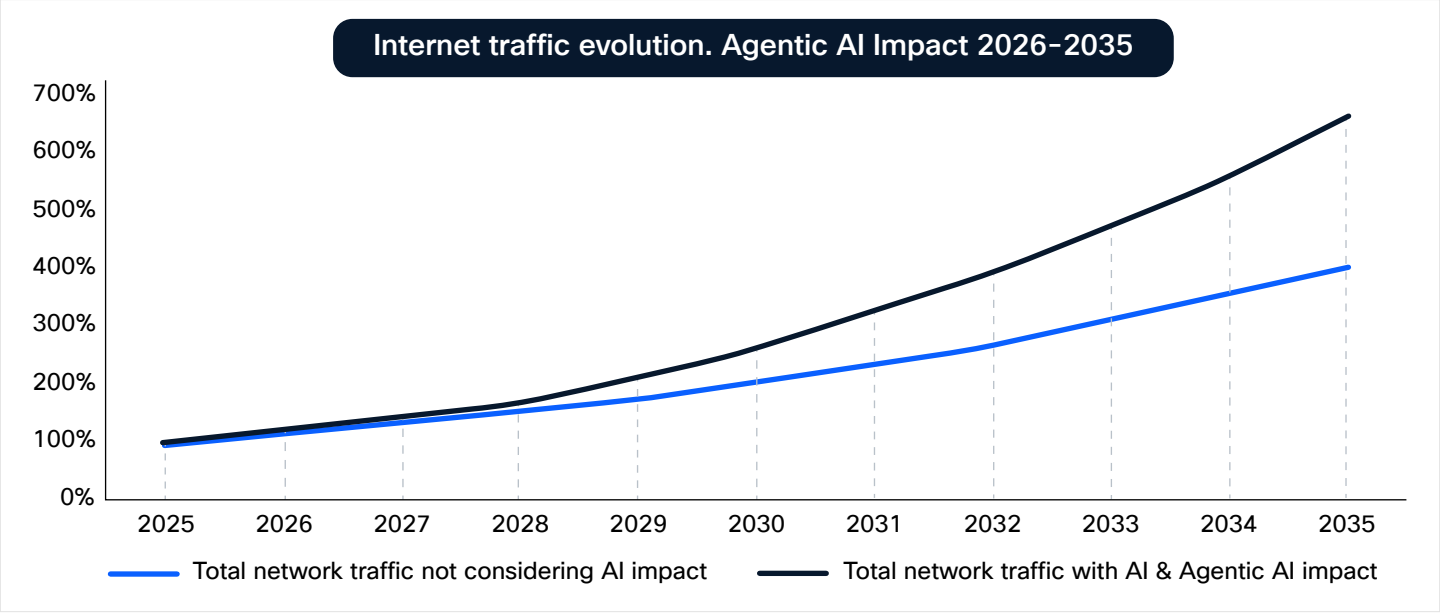

Figure 12. Projected impact of AI and agentic AI in internet network traffic

Source: Cisco model

The implications are significant, as consumer traffic is the primary contributor to overall internet and service provider network traffic. Without the impact of AI or agentic AI, network traffic is expected to grow 4x over the period 2025–2035, driven by the adoption of digital services. However, with the adoption of AI and agentic AI across many of those digital services, **traffic is projected to grow 6.6x from today's levels, or a 63% additional increment compared with the non-AI-impacted projection.** A relevant portion of that traffic will be **AI inference, which is expected to account for 25% of total traffic by 2035.**

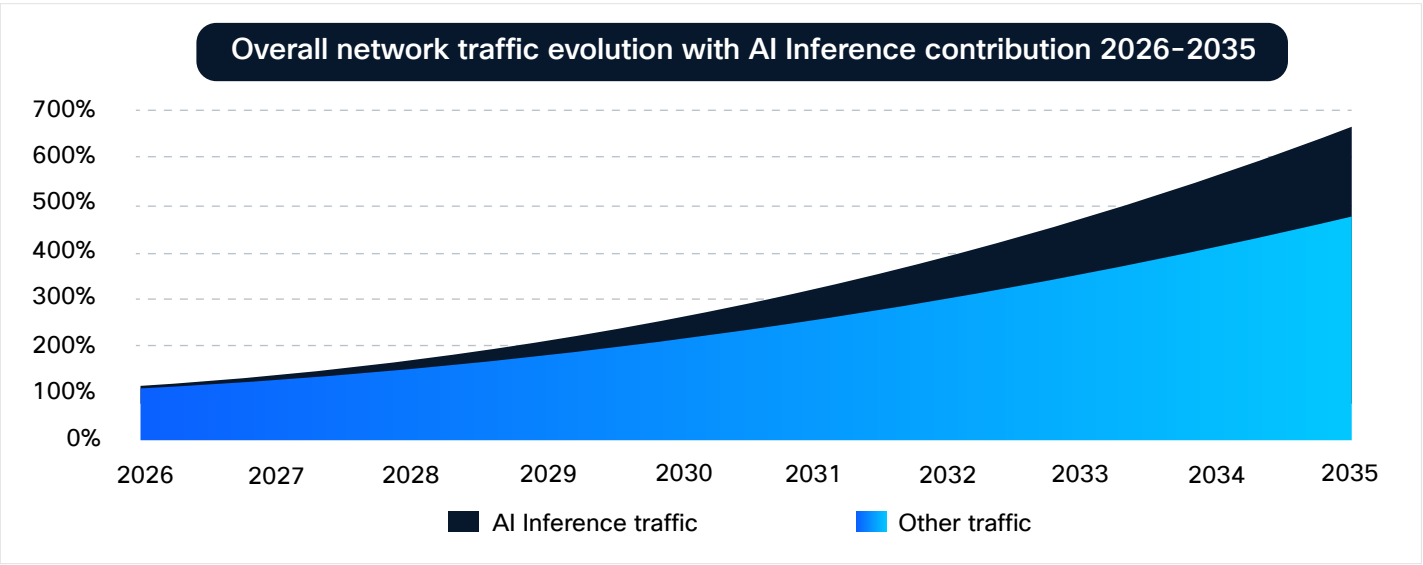

Figure 13. Projected network traffic growth, AI inference traffic versus other traffic

Source: Cisco model

Total **network traffic will experience significant growth** over the period 2029–2032, when adoption of agentic AI will experience a more pronounced increase, with a compound annual growth rate **(CAGR) of around 25% in AI inference traffic**.

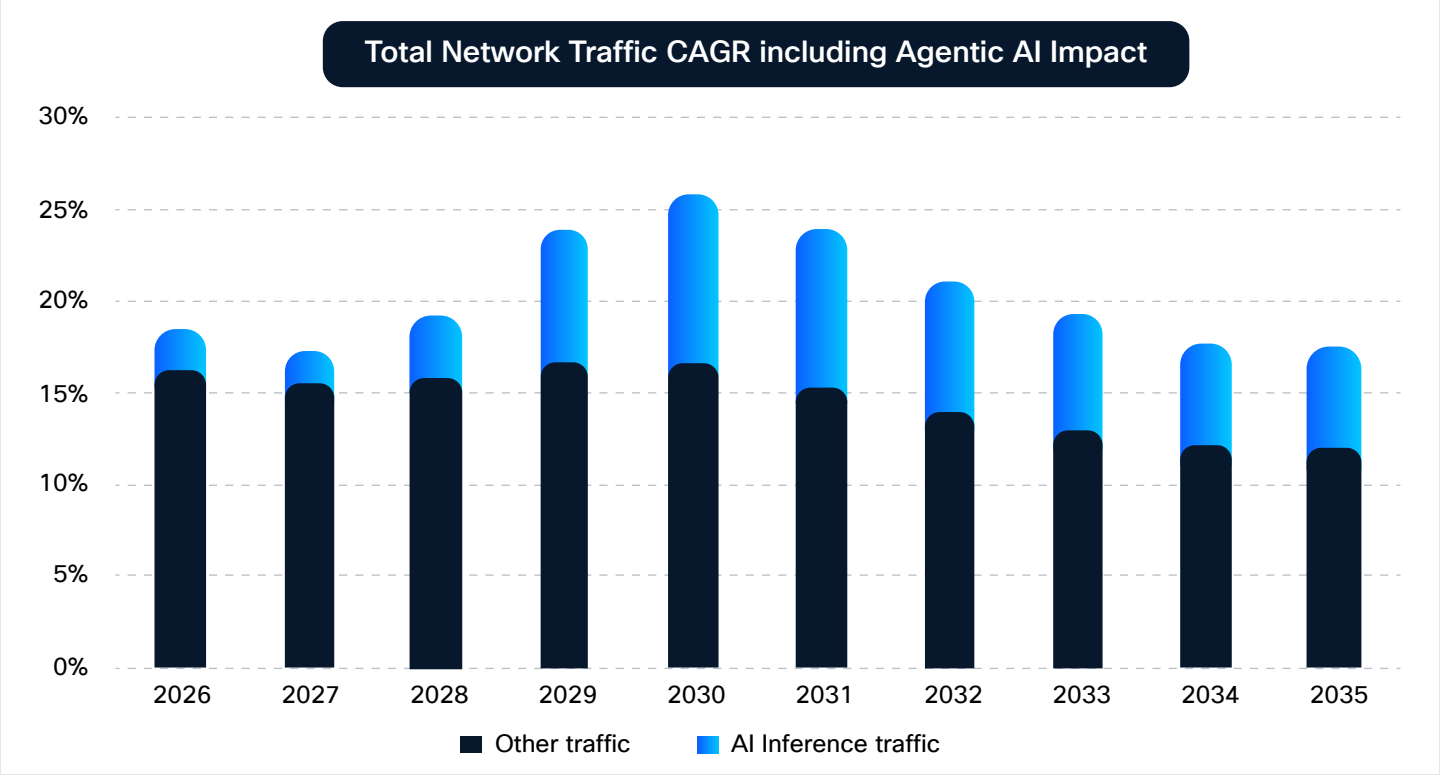

Figure 14. CAGR for both AI inference traffic and other traffic over the target period

Source: Cisco modeling

As adoption of AI and agentic AI continues in both enterprise and consumer digital services, **the transformation ahead is significant in terms of network traffic volume, with new implications for traffic symmetry, flow duration, and resiliency for critical paths.**

Further reading: Cisco report 2026 The Race to Agentic AI: Why Infrastructure Will Make or Break Workforce Transformation https://www.cisco.com/c/dam/en/us/solutions/collateral/artificial-intelligence/race-to-agentic-ai-report.pdf.

This research has been conducted in close collaboration with Opanga Networks.

Report Authors:
Johan Gustawsson (Cisco)
Javier Antich (Cisco)
Waris Sagheer (Cisco)

Contributors:
Ben Hadorn (Opanga Networks)
John Burnette (Opanga Networks)
Ryan Johnston (Opanga Networks)

C99-5828901-00 05/26

State of Wireless 2026

Unlocking the Multiplier Effect: How Strategic Wireless Investments Drive Enterprise Growth in the AI Era

Contents

Executive summary

Wi-Fi is a strategic growth engine — delivering measurable returns across multiple business dimensions simultaneously. This "multiplier effect" distinguishes wireless from other IT investments that typically optimize for a single outcome. In 2026, as organizations face the inflection point of connectivity demand and AI-driven transformation, strategic wireless investments are proving to be a catalyst for enterprise-wide success.

The evidence is compelling: more than three-quarters of organizations report operational efficiency gains (78%), see employee productivity improvements (75%), and report enhanced customer engagement (75%), while almost 7 in 10 (68%) experience positive revenue impacts from wireless investments. When organizations prioritize wireless strategically, they achieve returns that compound across the enterprise — a true multiplier effect, with one network delivering multiple business outcomes.

However, this opportunity emerges alongside mounting complexity. Organizations must adapt to diverse connectivity needs and support a growing spectrum of users and devices — from employees and contractors to autonomous robots, smart sensors, and AI-powered applications. New use cases appear daily, and IT teams are challenged to keep pace. This inaugural Cisco report explores the global state and future trends of wireless enterprise networking, revealing a "wireless AI paradox" with AI as both the leading driver for wireless return on investment (ROI) and the primary source of escalating challenges. The report is grounded in Wi-Fi as the primary enterprise connectivity layer, while examining the broader wireless ecosystem it enables, including AI-driven applications, IoT and OT environments, and emerging enterprise use cases.

We surveyed 6,098 wireless professionals from across the world who agree that AI simplifies operations and can free up over 850 hours per IT person, per year. But simultaneously, AI is challenging their infrastructure, the number one security threat, and the biggest talent risk. Legacy infrastructure plus three key interconnected barriers must be addressed to resolve this paradox and unlock the full multiplier effect of wireless:

1. Solving operational complexity.
98% of organizations report increased complexity, driven by factors like organizational mergers and the proliferation of IoT workloads. Most IT teams remain trapped in reactive ticket management. While 81% of wireless leaders prefer AI to simplify operations, only 29% have achieved high-level implementation.

2. Mitigating intensifying security threats.
AI-generated attacks are the leading driver of increased wireless security risk. Over half of organizations (58%) have experienced financial losses from wireless security incidents, with 50% of them exceeding US$1 million annually. Over a third (36%) of affected organizations report compromised IoT or OT devices as the culprit, representing a substantial threat to Wi-Fi as the most common connectivity technology for IoT.

3. Addressing talent shortages.
A shortfall of qualified staff amplifies complexity and threat vectors. Nearly 9 in 10 (86%) wireless leaders struggle to hire professionals, driving 70% higher security incident costs and trapping teams in reactive operations. While IT talent gravitates toward AI and cybersecurity, wireless teams need advanced AI expertise to fully realize the benefits of modern wireless.

These barriers are intertwined, compounding risk and making holistic action essential.

Organizations that successfully address legacy infrastructure and the three barriers together can unlock significant opportunities, becoming four times more likely to achieve strong wireless ROI (4:1 or higher). Strategic wireless investments yield measurable, compounding returns across multiple dimensions. This explains why wireless investment momentum continues to accelerate, especially as AI usage expands and innovations advance.

Our research identifies four key recommendations to capture the multiplier effect:

1. Accelerate your Wi-Fi refresh timeline.

2. Get on the path to AI and automation.

3. Prioritize holistic security modernization.

4. Build and enhance wireless talent.

Now is the time to drive competitive advantage. Organizations that act decisively and holistically – by addressing legacy infrastructure, simplifying operations, mitigating wireless security threats, and augmenting talent with technology – will position Wi-Fi as a strategic growth engine for the next decade.

Organizations investing holistically in AI, automation, modern security, and certified expertise have an edge on those that do not:

+4x

more likely to achieve strong returns on wireless investments

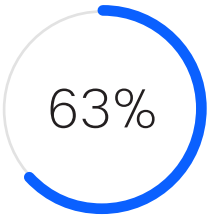

higher average ROI on wireless investments

The opportunity:
Wi-Fi as a strategic growth engine

Investment momentum continues to accelerate as benefits are realized

The business case for Wi-Fi investment is now undeniable. Organizations are seeing measurable ROI and the rise of business-critical AI deployments amplify the need for more resilient, higher performing Wi-Fi. Wireless budgets are growing accordingly, a trend that will likely intensify in the years ahead.

Eight in 10 organizations increased wireless investment over the past five years, with 29% implementing significant budgetary increases of 50% or more. What's more, 82% forecast that wireless budgets will continue to rise, with over a third (35%) expecting investment increases of 50% or more over the next four to five years. This is a clear signal that wireless will play an increasingly critical role in enterprise strategy.

This acceleration extends beyond aggregate spending. Three times as many organizations (32%) expect high returns from wireless tech over the next two years compared to the previous two. This trajectory is driven by use cases that are increasingly wireless-first (see page 10).

The pace of wireless change is increasing across industries and markets

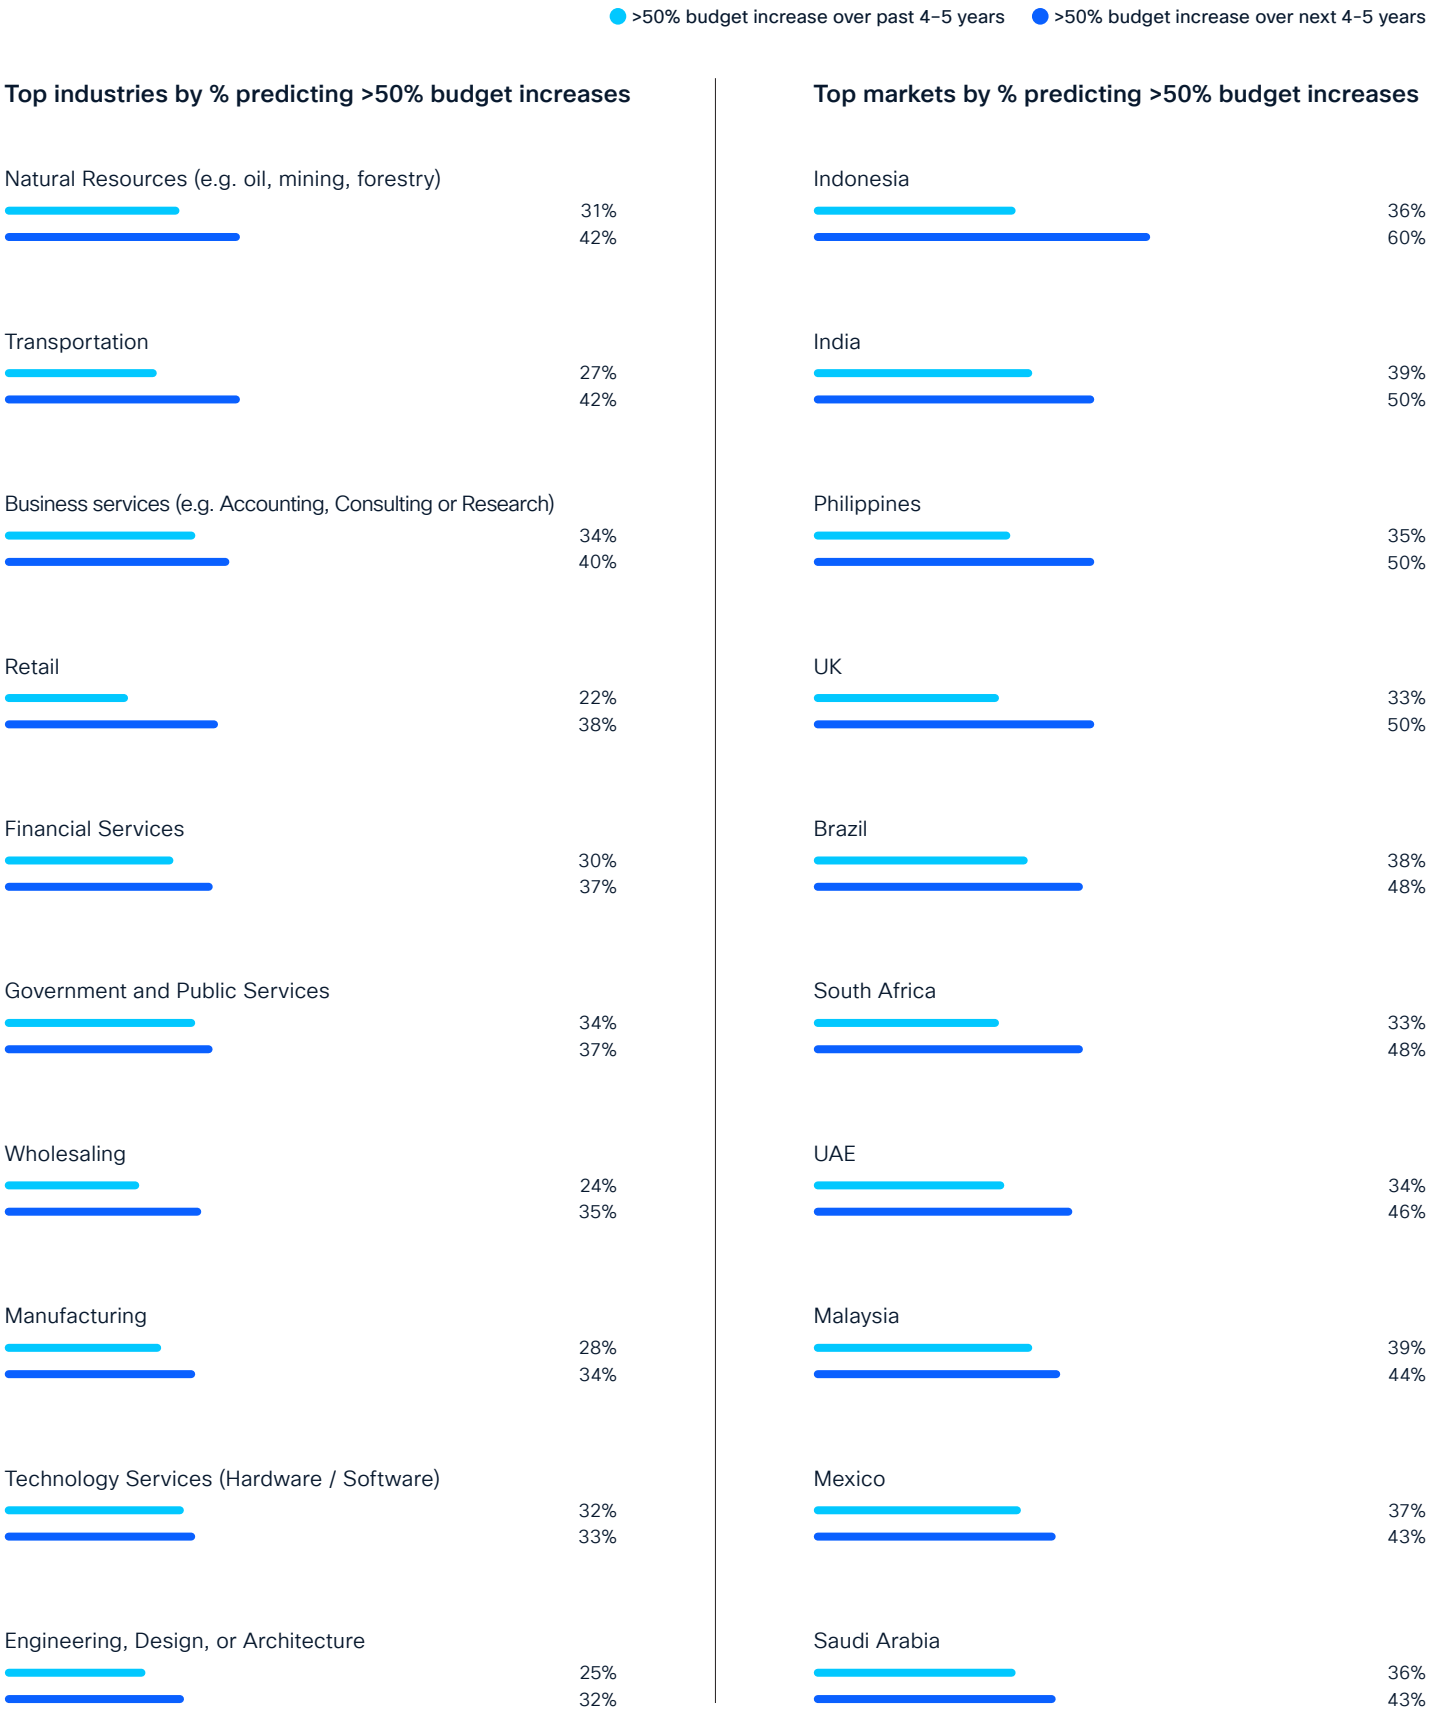

Capital expenditure typically accounts for 61% of wireless budgets on average, with the remainder allocated to operational expenditure.

The composition of wireless budgets further reveals organizational priorities. Organizations cite reducing operational risk as their number one budget priority, followed by security enhancement. The relatively even split across these priorities signals that organizations are weighing the opportunities and risks of wireless in almost equal measure when making investment decisions.

This increased investment is directly enabling organizations to deploy Wi-Fi to meet changing business needs and address emerging AI-powered use cases.

Top investment drivers for next wireless investment

High-ROI expectations for wireless are accelerating

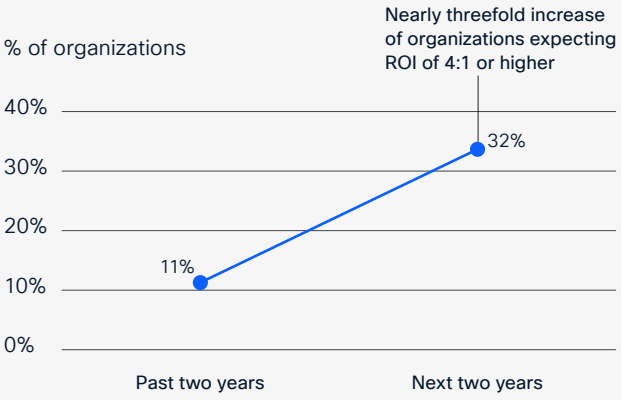

Average annual wireless spend, by organization size

250–499	$6.15M
500–999	$9.63M
1,000–4,999	$12.74M
5,000–9,999	$19.94M
10,000+	$39.01M

Expanding use cases drive new infrastructure requirements

The breadth and diversity of wireless use cases now extend across entire organizations, further illustrating how wireless has become foundational to how modern enterprises operate and compete.

No longer relegated to employee connectivity use cases, Wi-Fi enables everything from critical AI workloads to physical security, customer experience, and operational intelligence. This expansion of wireless "purpose" creates new performance and reliability demands that traditional infrastructure cannot meet.

More than a quarter of organizations (28%) have already deployed AI workloads. This figure is estimated to climb to over three quarters (79%) by 2027, with an additional 29% in pilot stage and 22% planning deployment over the next 12 months. This rapid trajectory further highlights that AI deployment is indeed the defining use case for next-generation wireless infrastructure.

While core use cases such as wireless for physical security are already widely deployed, the next phase of wireless growth is being driven by emerging applications that depend on high-performing, resilient networks. Organizations are increasingly piloting or planning wireless investments to support autonomous systems and robotics, smart facilities and energy management, space analytics, and immersive collaboration.

Organizations are leveraging wireless across an array of use cases

	Currently deployed (%)	Planning to deploy (%)
Autonomous systems & robotics	45%	50%
Space analytics & optimization (footfall, etc.)	50%	44%
Smart facilities & energy management	51%	43%
High-definition streaming	54%	43%
Immersive collaboration & training	54%	43%
Indoor wayfinding	53%	42%
Supply chain & inventory intelligence	55%	42%
AI applications & workloads	57%	41%
Operational visibility & flow analytics	57%	41%
Real-time asset & equipment tracking	57%	41%
Internet of Things	59%	38%
Guest wireless	61%	38%
Customer & user experience enhancement	60%	37%
Physical security (CCTV)	62%	36%

Current footprint = Deployed + Pilot stage
Future expansion = Planned next year + Planned in next 2–5 years

Wireless leaders cite five primary drivers of increased wireless dependence for these use cases. Supporting the growth of the Internet of Things ("IoT") ranks as the leading driver, followed by increased user mobility, bring-your-own-device adoption, and high-bandwidth application adoption, including real-time collaboration, AI workloads, and operational technology requirements for point-of-sale systems and manufacturing.

These drivers reflect a concerted industry transformation affecting the entire enterprise landscape simultaneously. IoT proliferation, remote work normalization, high-bandwidth application adoption, and AI deployment are occurring across every geography and industry. This common pressure further explains why organizations are increasing wireless investments, and they must continue to invest to take advantage of the multitude of opportunities presented by Wi-Fi amid AI advancements and transformation and maintain a competitive edge.

Key drivers of growing wireless dependence

Increase in IoT and Smart Devices (e.g. sensors, cameras, smart building controls)	63%
Increase in user mobility and BYOD (e.g. supporting hot desking)	55%
Adoption of high-bandwidth apps (4K/8K streaming, AR/VR, real time collaboration (e.g. Teams/Zoom/Webex)	52%
AI workloads	44%
Supporting OT requirements (wireless point of sale, warehouse scanners, and manufacturing robots)	23%

The multiplier effect: Wi-Fi drives measurable business value across multiple areas

Wi-Fi has evolved into a strategic growth engine, delivering tangible returns across multiple dimensions. This “multiplier effect” distinguishes it from other IT investments that typically optimize for a single outcome.

The evidence is compelling: More than three quarters of organizations report operational efficiency gains (78%), see employee productivity improvements (75%), and report enhanced customer engagement (75%), while almost 7 in 10 (68%) are experiencing positive revenue impacts from wireless investments. When organizations strategically prioritize wireless in line with business goals, they achieve returns that compound across the enterprise.

This multiplier effect — one network delivering multiple business outcomes — explains why strategic wireless investments correlate directly to strong ROI and will amplify as innovations advance. This measurable impact is the reason why wireless investment momentum continues to accelerate.

However, this opportunity arrives hand-in-hand with a convergence of challenges. Organizations find themselves navigating an increasingly fragmented connectivity landscape, forced to accommodate an exploding taxonomy of users and devices—from traditional employees and contractors to robots, IoT, and AI-driven applications that fundamentally redefine network demands. As new use cases proliferate at an accelerating rate, IT teams confront a widening capability gap.

This inaugural Cisco report examines the global state and evolutionary trajectory of wireless enterprise networking, exposing a striking “wireless AI paradox”: AI simultaneously serves as the most compelling driver of wireless ROI while simultaneously functioning as the principal catalyst of operational complexity and risk.

Most organizations see positive operational and commercial impacts from wireless investments

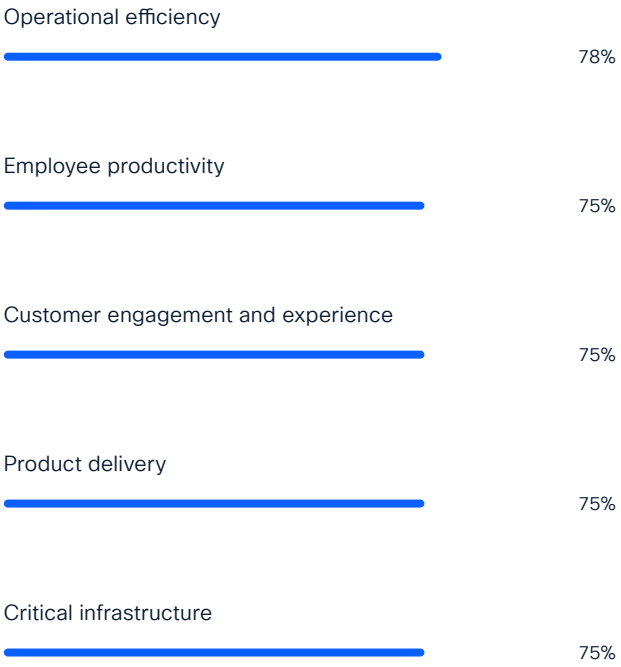

Wireless strategy in a perfect storm:
Navigating the AI paradox and the barriers limiting ROI realization

Defining the wireless AI paradox and why it matters

The wireless AI paradox lays out the central strategic challenge for enterprise leaders in 2026 and the opportunity for first movers. AI is simultaneously the leading driver of wireless ROI and the source of its greatest challenges. Organizations deploying AI view wireless as strategically critical and are able to achieve substantially stronger returns when they integrate wireless optimization into AI deployment strategies. Yet this same AI is introducing new security threats and intensifying talent competition.

The Wireless AI Paradox
AI is both the solution and the challenge

Solution

- AI-driven operations simplify wireless complexity
- Automation frees IT teams to focus on strategy
- Streamlined ticket resolution accelerates workflows

Challenge

- AI-generated cyberattacks are a top security threat
- Talent shortages in advanced wireless/AI skills hinder adoption
- IT talent is prioritizing AI over wireless

AI is the leading path to ROI in wireless, but also the biggest source of risk.

Organizations deploying AI workloads recognize wireless criticality differently than others. Among wireless leaders in organizations deploying AI workloads, three in five (62%) view wireless as strategically critical, compared to 46% overall. Wireless leaders in organizations deploying AI are more likely to have experienced 6 GHz adoption drivers first-hand, including solving congestion issues (72% vs. 46%) and enabling high-bandwidth applications (70% vs. 32%).

The reason for this heightened strategic importance is straightforward: AI workloads demand higher-performing and more resilient wireless networks. Those who integrate wireless optimization into their AI deployment strategies realize substantially greater returns. Around 8 in 10 organizations deploying AI workloads report positive impacts from wireless investments in the areas of operational efficiency, customer engagement, employee productivity, and revenue gains, which was at least seven percentage points higher than average. The conclusion here is that organizations see higher ROI from wireless modernization when paired with investments in AI.

How are these opportunities, challenges, and risks around AI advancements connected?

While AI is seen as the path to simplify wireless operations and resolve complexity issues, AI-generated or automated cyberattacks are also the top driver of increased wireless security threats. In addition, there is a shifting trend in prioritization of talent from wireless to AI-oriented roles.

Organizations deploying AI workloads recognize the criticality of wireless and potential of 6 GHz use cases

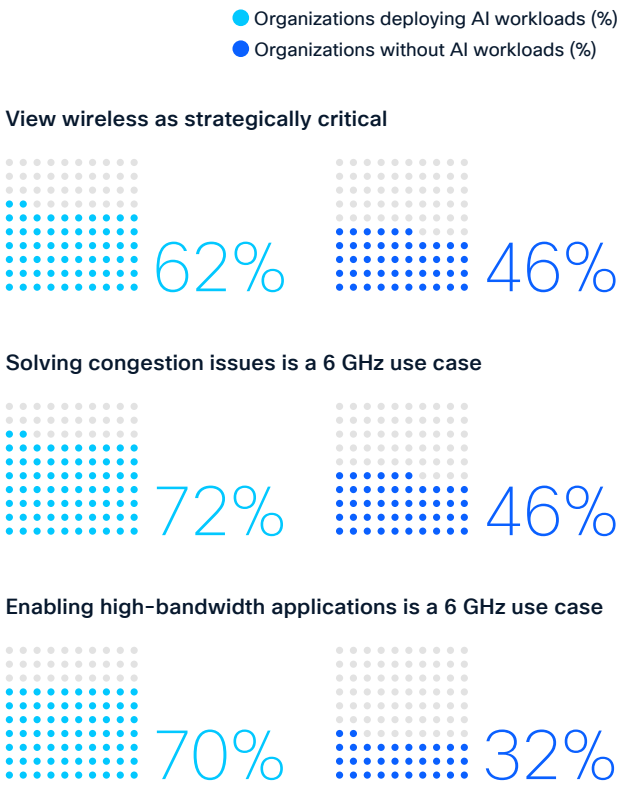

Organizations see greater benefits from wireless modernization when paired with AI deployments

AI poses multifaceted challenges to wireless teams

Top drivers of security threats	Top domains attracting IT talent away from wireless	Top barriers to hiring wireless talent
#1 AI-generated or automated cyberattacks / automated intrusion tools	**#1** AI / Machine Learning	**#1** Shortage of candidates with advanced wireless or AI-integrated skills
#2 Increased use of IoT and connected devices	**#2** Cybersecurity	**#2** Internal budget constraints or hiring freezes
#3 Lack of skilled personnel or bandwidth to monitor and respond to threats	**#3** Software engineering / app development	**#3** Geographic limitations or remote work challenges

This convergence of factors is giving rise to the wireless AI paradox that organizations must resolve. To do this, organizations must overcome legacy infrastructure and three key operational barriers: operational complexity, wireless security challenges, and major talent gaps.

Organizations that are already actively resolving this paradox – by modernizing infrastructure, simplifying wireless operations, mitigating security challenges, and addressing talent gaps – are benefiting with 63% higher average wireless ROI than those that have not.

Infrastructure at a crossroads: bridging the Wi-Fi performance gap

To capture the opportunities outlined above, including emerging AI workloads, organizations need to reconsider whether their current infrastructure can support their business needs. The performance gap between legacy and modern Wi-Fi is clear, but adoption is still slow.

Wi-Fi 5 is still the most widely used Wi-Fi standard (43%), but since it is an aging standard, it's questionable whether it's ready to scale many of the above use cases. A small yet growing number of organizations (19%) are presently deploying Wi-Fi 6E or 7, while three in five organizations (59%) plan to deploy Wi-Fi 6E or 7 in the next year, suggesting that organizations are beginning to react to enterprise needs.

Legacy Wi-Fi simply can't scale high-bandwidth and low-latency needs, especially in device-dense environments. This explains why there has been a 23% increase in Cisco Wi-Fi 6E and 7 access point deployments across its five million networks in the second half of 2025 (June–December).

Each successive Wi-Fi generation builds on the previous one, adding features, spectrum availability, and functionality. Built in response to business needs, they determine which use cases organizations can support and which competitive advantages they can realize.

Organizations increasingly factor Wi-Fi 7 into deployment plans

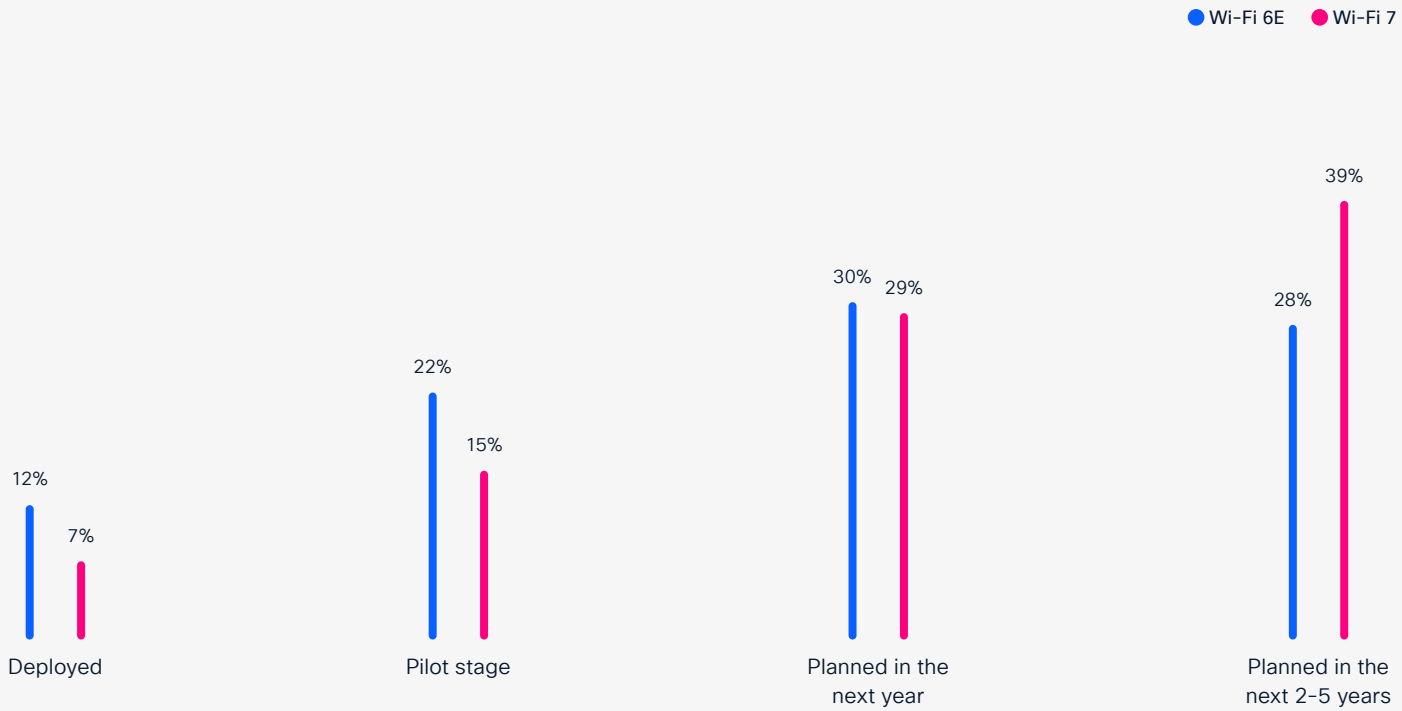

One such advance is the addition of the 6 GHz band added by Wi-Fi 6E, and utilization further improved with Wi-Fi 7. Organizations are using it to solve capacity and congestion issues (46%), enable high-bandwidth applications (32%), and support AI workloads (31%). And the research shows that those adopting this added spectrum are seeing strong benefits.

Organizations already deploying 6 GHz show almost double the rate of AI applications and workloads (45%) compared to non-adopters (26%). The 6 GHz band supplies the bandwidth required by AI-powered applications and correlates with improved scalability — 6 GHz users report higher readiness to scale Wi-Fi across devices, sites, capacity, and bandwidth-intensive apps.

It's no surprise that Cisco telemetry highlights a 60% increase in 6 GHz clients going live in 2025. Wireless professionals are seeing the 6 GHz opportunity and are helping their organizations deliver on that.

Wireless modernization is a driver of the value creation shared earlier in the report, but alone it is not enough. The research reveals that overcoming the wireless AI paradox requires organizations to tackle legacy infrastructure together with complex operations, intensifying security threats, and talent gaps obstruct IT from resolving it.

Next-generation Wi-Fi adoption gap signals upgrade pressure

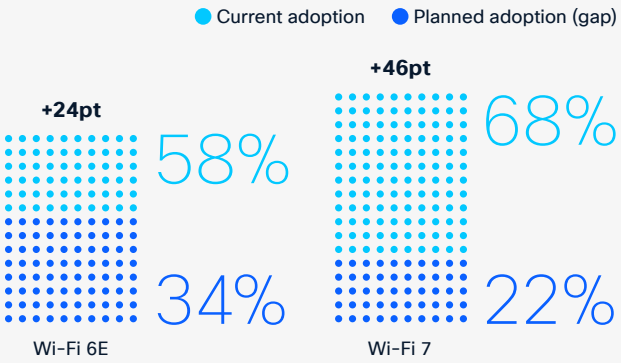

Top cited 6 GHz use cases

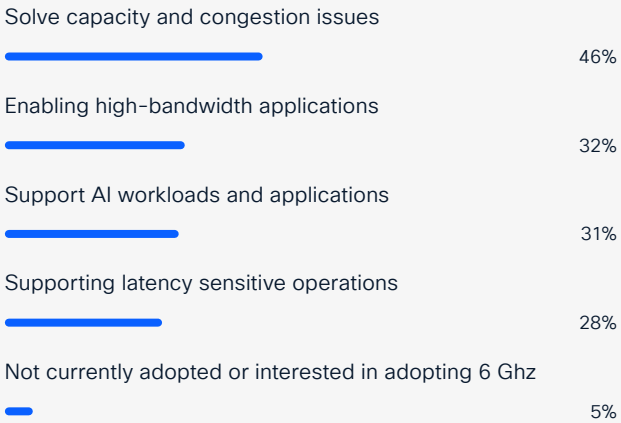

AI deployment and scalability differs by 6 GHz deployment status

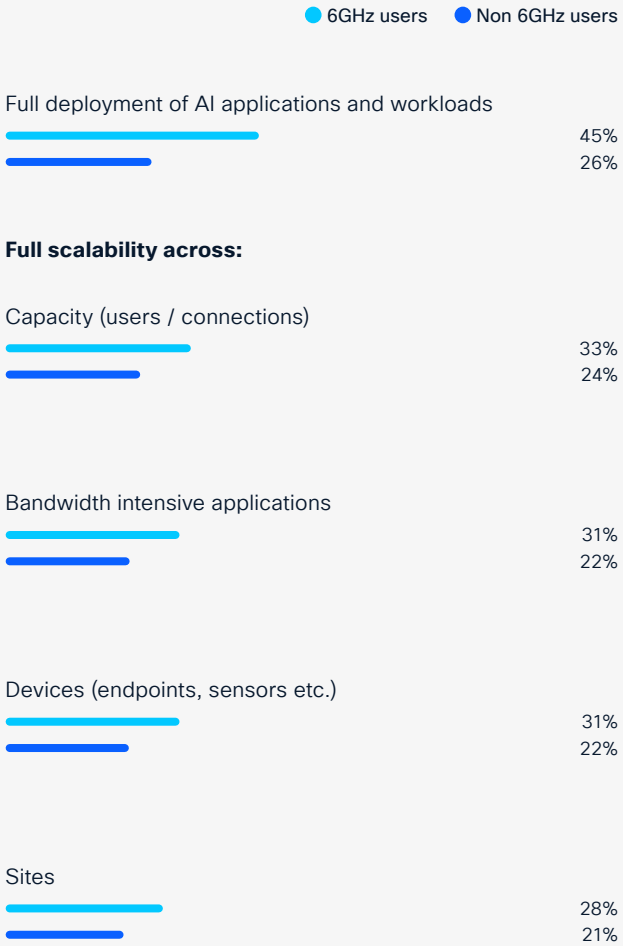

Barrier 1: Operational complexity overwhelms current capabilities

The first barrier limiting organizations from resolving the AI paradox is mounting operational complexity. Nearly every wireless leader (98%) reports that wireless operations are becoming more complex, creating a reactive posture that drains resources, prevents strategic work, and directly undermines the AI initiatives that help reduce complexity. This creates a vicious cycle: complexity drives reactive work, reactive work limits modernization, and delayed modernization perpetuates complexity.

Organizations cite three primary drivers of this growing complexity: mission-critical workloads, increasingly including AI-driven applications (43%), the need to mitigate new security risks (42%), and rising bandwidth demands from new use cases (38%).

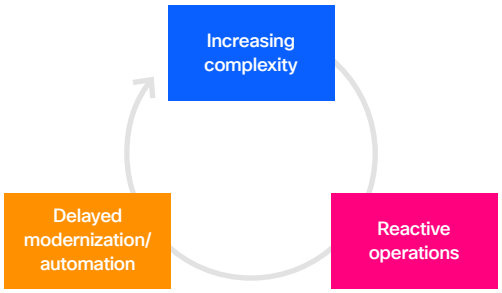

This complexity translates into tangible operational strain. Nearly half of respondents (43%) report that their team receives at least 50 wireless support tickets a week, with the average standing at 68. This means IT teams are spending hundreds of hours per month managing wireless tickets.

Wireless leaders expect this burden to compound. Nearly two-thirds (64%) expect ticket resolution times to increase over the next two to three years, suggesting a growing need to address this complexity.

An area of concern is the reactive posture typically stemming from this complexity. A majority of wireless teams (55%) spend most of their time on reactive troubleshooting and incident management. This reactive burden is most pronounced in K-12 education (64%), wholesaling (62%), and business services (57%).

Organizations often cite AI, security risks, and rising bandwidth demand among the top three drivers of growing wireless complexity

Supporting mission-critical IT/IoT/OT workloads (including AI)	43%
Security risks	42%
Bandwidth demand increase or unpredictability	38%
Client unpredictability (usage patterns, mobility)	36%
Proliferation of devices	25%

Industries with respondents most likely to spend the majority of time on reactive tasks

Education (primary and secondary)	64%
Wholesaling	62%
Business services	57%
Technology Services	57%
Engineering, Design, or Architecture	57%

This reactive operational posture directly undermines modernization efforts. Teams constrained to reactive troubleshooting may divert resources and attention from other pursuits, such as strategic wireless planning and modernization, training and certifications, or implementing automation. The result is a reinforcing cycle: complexity contributes to reactive tasks, reactive tasks redirect teams from implementing solutions that would otherwise help reduce complexity, and unresolved complexity contributes to reactive work.

A critical factor compounding this operational challenge is the lack of visibility. Nearly 9 in 10 organizations (87%) report visibility gaps that impair their ability to troubleshoot Wi-Fi issues effectively. The top three frequently reported challenges are poor client visibility, application and cloud visibility, and packet visibility.

Without end-to-end visibility, teams cannot rapidly isolate problems. This creates a particularly dangerous dynamic: wireless networks become scapegoats for problems originating elsewhere. A quarter of complaints (25%) are inaccurately attributed to wireless, with each misattributed incident wasting an average of 18 hours across teams. The true culprits are most commonly application problems (25%), client or endpoint issues (22%), and cloud or external service failures (18%).

The operational challenges directly impact organizations' ability to realize ROI.

Amid rising AI-driven organizational transformation, wireless leaders overwhelmingly believe AI represents the most promising path to overcome these rising complexity challenges. Over 8 in 10 (81%) would prefer AI with automation to fully or mostly handle routine wireless operational tasks.

87% face visibility gaps, including:

Poor client visibility	43%
Poor application/cloud visibility	41%
Poor packet visibility	38%

This demand for AI with autonomous actions suggests a need beyond simple AIOps. It indicates that wireless professionals are looking for AI to do more than assist with alerts and recommendations. Their preference is for AI-driven autonomous networking to be reasoned through and acted upon at machine speed, coupled with an appropriate level of human oversight.

Fast movers are already acting, as we observed a 64% increase in usage of wireless AI capabilities in the second half of 2025—a significant uptake that helps bring immediate value to organizations.

The benefits are substantial and measurable. Almost all organizations that have implemented high-level AI with autonomous actions (98%) report dramatic time savings, freeing up an average of 3 hours and 20 minutes per day, per person. These organizations benefit from operational improvements: they are four times more likely to rate their network operations as very simple with 12% faster ticket resolution times than manually operating counterparts.

However, a significant gap exists between preference and reality. Only 29% of organizations have AI with autonomous actions supporting wireless ticket management, and 29% supporting security monitoring and incident response, while 23% have automation supporting capacity planning and provisioning. This gap reflects multiple constraints. Two-thirds (66%) do not expect to achieve a high-level of AI with autonomous actions within five years, which may signal technical, organizational, or resource barriers to deployment, despite strong organizational preference.

The AI gap: Desire versus reality

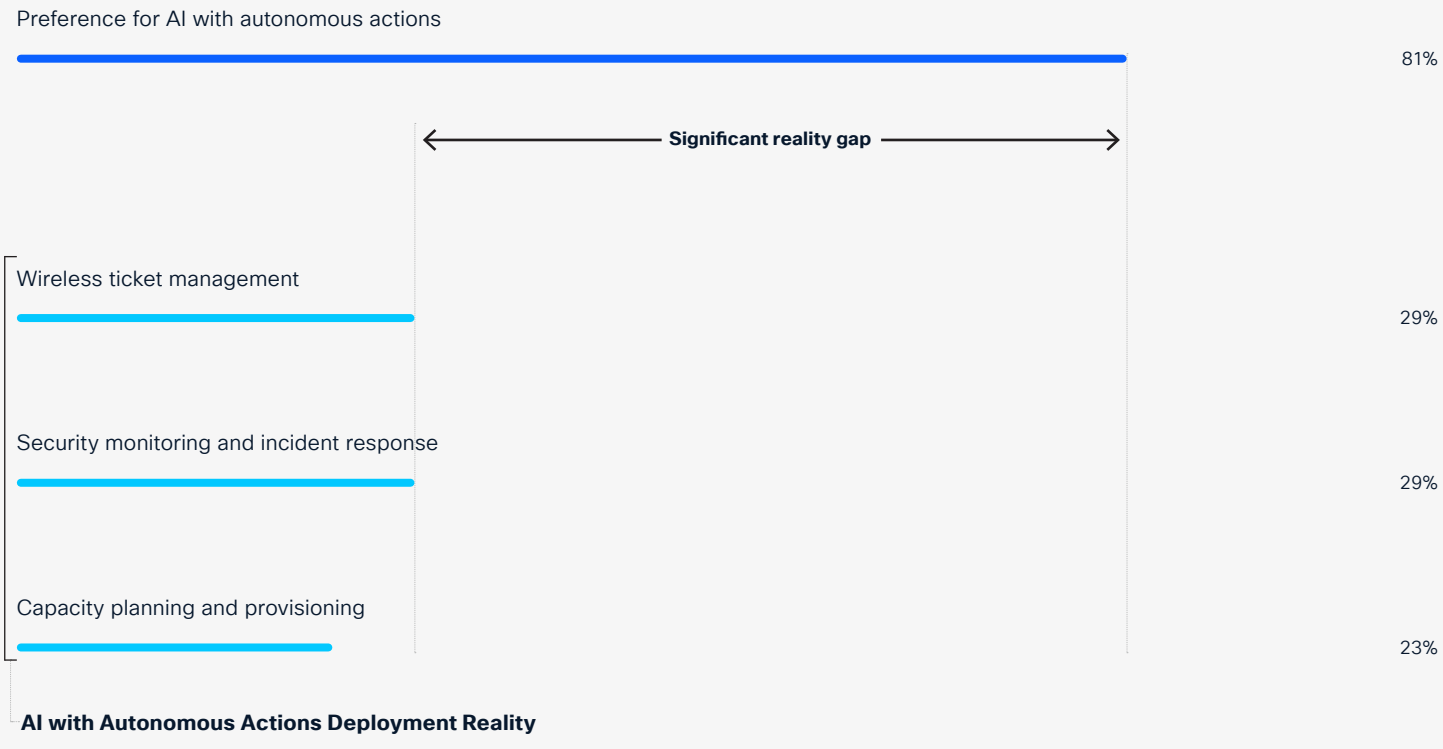

Organizations that have fully deployed AI with autonomous actions substantially outperform those with mostly manual operations. Wireless leaders from these organizations are three times more likely to expect reduced ticket resolution times over the next two years, four times more likely to describe network management as very simple, and nearly twice as likely to describe infrastructure as fully scalable.

Organizations with high levels of AI with autonomous actions show far greater confidence in infrastructure scalability. They are almost twice as likely to describe their infrastructure as fully scalable compared to those with mostly manual operations. Organizations reporting strong ROI show substantially higher deployment of AI with autonomous actions across ticket management (31%), capacity planning (28%), security monitoring (31%), vendor coordination (28%), and network optimization (30%).

Operational complexity alone represents a significant challenge to resolving the wireless AI paradox and thus improving wireless ROI. When combined with escalating security threats, the second barrier — the impact on organizational resilience and financial performance — becomes even more severe.

Organizations with fully automated operations and AI-driven operations outperform those with mostly manual operations in several key areas:

3x more likely to forecast significant reduction in ticket resolution times over the next two years

4x more likely to describe management of their networks as 'very simple'

Almost **2x** more likely to describe device infrastructure as fully scalable

12% lower average ticket resolution time

Area of automation and AIOps deployment

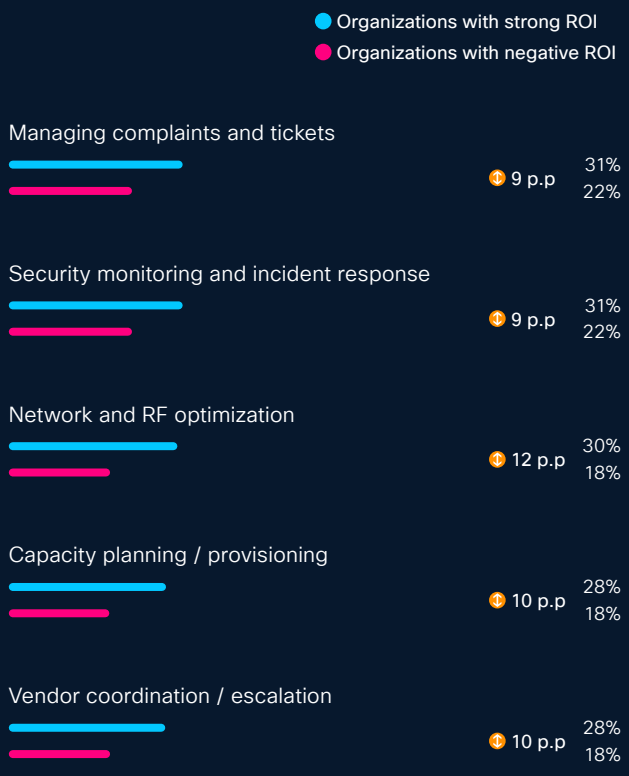

Barrier 2: Wireless security under siege – IoT sprawl meets AI-powered attacks

Wireless security represents the second critical barrier preventing organizations from resolving the AI paradox and realizing strong wireless ROI. Organizations cannot confidently deploy Wi-Fi as a platform for business-critical workloads while facing escalating security threats and mounting financial losses.

Eighty-five percent of organizations have experienced at least one wireless security incident in the last 12 months. Further to this, more than one-third report experiencing escalating wireless threats over the past two years (38%). Wireless leaders state that security threats have become more frequent and damaging, while also being more difficult to detect and remedy.

AI has emerged as the most significant new security threat vector. Wireless leaders are most likely (35%) to cite AI-generated or automated cyberattacks among the top three drivers of increased wireless security threats. These attacks can identify network vulnerabilities, adapt attack strategies based on defensive responses, and operate at a scale and speed far exceeding human attacker capabilities. Additionally, AI-generated or automated cyberattacks may lower the barrier to entry for attacking Wi-Fi networks, allowing threat actors to operate with a sophistication and speed that previously required more significant resources.

AI-driven attacks on wireless networks confront underprepared defenses

Feature	Traditional Wi-Fi Attack	AI-Powered Wi-Fi Attack
Speed	Limited by hardware and wordlist size.	Optimized by AI models that ‘guess’ faster.
Detection	Easier to spot via repeated failed attempts.	Can mimic human patterns to stay under the radar.
Automation	Requires manual setup and monitoring.	Can autonomously ‘pivot’ from one device to another.
Social Engineering	Generic phishing emails.	Hyper-personalized lures using Deepfakes/GenAI.

The nature of Wi-Fi cybersecurity threats is also shifting from passive reconnaissance and credential theft toward active compromise and operational disruption. Rogue access points, compromised credentials, and denial-of-service attacks represent the top attack vectors.

The attack surface continues to expand as well. Over a third of affected organizations (36%) report disruption from compromised IoT or OT devices, representing a substantial threat to Wi-Fi since it's the most common connectivity technology for IoT. The proliferation of IoT devices, especially when unmanaged, creates an aggregated vulnerability as individual weaknesses compound into network-wide exposure.

Wireless threats are accelerating over time

report increases in wireless security incidents in the past two years

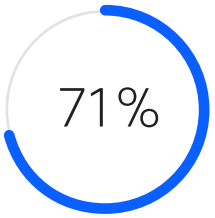

expect such incidents to increase in the next two years

Key contributors to increased threat level for wireless networks

AI-generated or automated cyberattacks / automated intrusion tools	35%
Remote and hybrid work models expanding attack surface / unmanaged endpoints	31%
Increased use of IoT and connected devices (rapid device growth)	30%
Lack of skilled personnel or bandwidth to monitor and respond to threats	27%
Difficulty managing multiple security layers and segmentation	25%
Budget or resource constraints limiting security improvements	25%
Poor user practices, human error, or insider risks	24%
Limited network and application visibility	21%
Legacy infrastructure and protocols	20%
Difficulty managing security patching	20%
Limited client visibility	16%

Losses incurred over the past 12 months from wireless security incidents

Less than US$1 million	42%
US$1-10 million	25%
US$10-50 million	19%
Over US$50 million	7%

Organizations most likely to lose over US$ 1m from wireless security incidents

Industries:

K-12 Education 62%

Wholesaling 55%

Construction 54%

Company size:

500-999 employees 63%

Outlook on wireless:

Not of strategic importance 60%

The financial impact of these security incidents is substantial. Fifty-eight percent of organizations have experienced financial loss due to wireless security incidents. Half (50%) report losses exceeding US$1 million in the past year, representing a sizable financial impact that alone justifies Wi-Fi security investment.

Organizations are losing more than money due to wireless security incidents. In addition to financial losses, one in three (36%) experienced loss of customer trust, and a similar number (34%) faced regulatory penalties or compliance consequences. The impact of reputational damage and regulatory exposure often extends well beyond direct incident costs.

Most organizations still maintain confidence in their wireless security. Eighty-three percent report that their organization is doing enough to protect wireless networks despite over seven in 10 (71%) expecting wireless security failures to increase in the next two years. Executive-level respondents show far greater confidence in wireless defenses than frontline staff, suggesting that organizational leadership may

The hidden costs of wireless security incidents

Unauthorized access to internal systems via wireless (credential misuse, rogue AP)	36%
Compromise of IoT/OT devices through wireless leading to process disruption	36%
Increased incident response / forensic costs due to wireless security incidents	36%
Customer data exposure or loss of customer trust specifically traced to a wireless breach	36%
Regulatory, contractual, or compliance penalties directly linked to a wireless security breach	34%
Data exfiltration or confirmed data breach tied to wireless access	32%
Financial loss from fraud or transactional compromise originating via wireless	32%
Targeted lateral movement after initial wireless compromise (escalated attack)	32%

Top cited barriers to improving wireless security

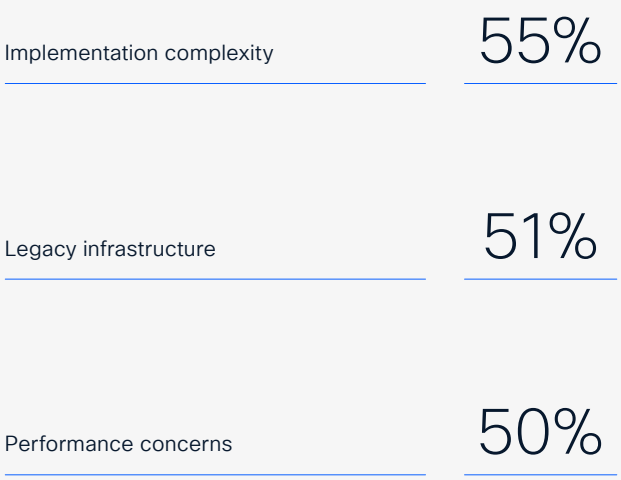

WPA3 implementation drives safety and ROI for wireless operations

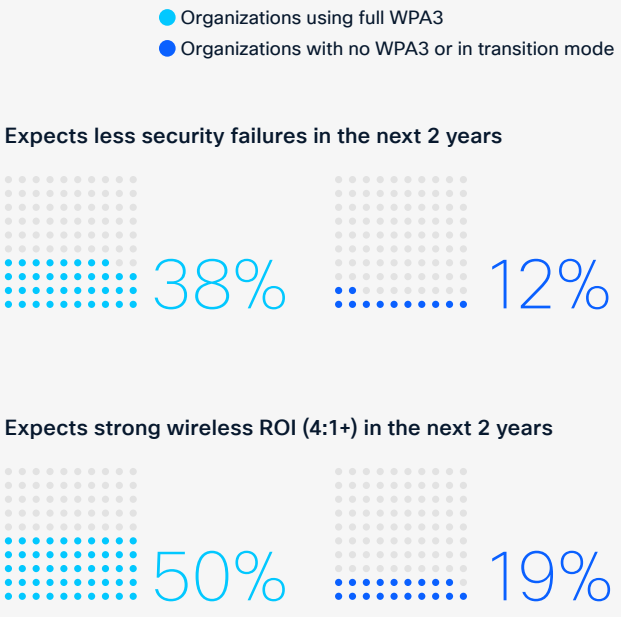

underestimate security challenges compared to technical staff who directly manage networks.

Organizations cite three primary barriers to improving wireless security: (1) implementation complexity, (2) legacy infrastructure, and (3) performance concerns. These barriers do not exist in isolation but reflect the broader wireless challenges of talent shortages, visibility gaps, and rising operational strain that restrict organizations' ability to modernize security.

The result is a widening vulnerability gap: even as risks escalate, organizations remain constrained by outdated systems, complexity, and performance concerns, slowing transformation and eroding resilience.

The last point is especially important because research shows organizations that implement modern authentication and encryption protocols demonstrate better security outcomes and outsized business performance compared to those that don't. For example, those using certificate- or profile-based authentication are almost three times more likely to predict a decline in wireless security failures compared to those using only passwords, and four times more likely compared to those operating open networks. They also experience approximately 18% lower financial losses on average than those not using modern authentication protocols.

These advantages are further reinforced when organizations adopt WPA3 where adoption is directly correlated with stronger security and ROI. Over a third of respondents from organizations using full WPA3 (38%) expect fewer security failures in the next two years compared to 12% of organizations with no WPA3 or in transition mode. Half of respondents from organizations with full WPA3 (50%) expect strong ROI (4:1 or greater) in the next two years compared to 19% without WPA3.

Organizations leveraging the embedded security of next-generation Wi-Fi are around three times more likely to expect strong ROI and predict fewer security failures over the next two years. Organizations now have a clear path to reducing operational risk and financial losses while improving user and customer trust. However, implementing modern security protocols requires specialized expertise — which is increasingly difficult to find. This brings us to the third barrier: the competition for wireless talent.

Barrier 3: Wireless competition for AI skills

Talent represents the third barrier and, in tandem with operational complexity and increasing security threats, is creating a strong catalyst that inhibits organizations from scaling wireless ROI.

Talent shortages do not merely slow modernization, they directly amplify operational strain and security exposure while making it more difficult to implement AI for networking. This creates a vicious cycle: organizations lacking talent struggle to modernize, complexity and security risk escalate, costs rise, and the best talent leaves for more modern organizations.

Eighty-six percent of organizations report challenges in hiring, with IT talent shifting to technology fields such as AI and cybersecurity. Respondents have indicated that AI ranks as the number one area that attracts talent from wireless (50%), while a shortage of candidates with advanced wireless or AI-integrated skills is the number one barrier to hiring.

This creates a skills gap that translates into higher operating costs (39%), lower morale (37%), and reduced innovation (32%). Talent pressure is most evident in technology and manufacturing, where competition from AI, cybersecurity, and OT roles further constrain the ability to scale wireless expertise.

The correlation between talent and negative outcomes is clear. Organizations facing extreme difficulty in hiring have wireless teams spending far more time on reactive tasks. Half of those facing extreme recruitment difficulty (50%) spend most of their time on reactive tasks compared to 37% of those facing no difficulty.

The impact is not just operational: organizations struggling to hire wireless specialists expect an increase in wireless security failures at substantially higher rates (85% vs. 59% for those with no hiring difficulty, a 26 percentage-point gap).

AI-linked to wireless brain drain and skills shortages

Ranking among the top three domains attracting talent away from wireless

AI / Machine learning	50%
Cybersecurity	48%
Software engineering / app development	40%
Cloud infrastructure / DevOps	37%
Data science / analytics	34%
Smart building / sustainability tech	31%

Primary reasons for difficulty in hiring wireless talent

Shortage of candidates with advanced wireless or AI-integrated skills	51%
Internal budget constraints or hiring freezes	37%
Geographic limitations or remote work challenges	36%
Lengthy hiring process or internal bottlenecks	34%
Lack of interest in wireless as a career domain	34%
Competition	28%

Talent constraints create a reactive, low-automation trap

	Organizations facing no difficulty in recruitment	Organizations facing high difficulty in recruitment
Wireless team time spent on mostly reactive tasks	37%	50%
Average automation of wireless network functions	34%	23%
Expectation of increase of wireless security failures	59%	85%
Annual cost of wireless security incidents	**US$12,448,230**	**US$21,169,950**

More concerning, they already experience 70% higher costs for security incidents annually than those with no recruitment challenges. Wireless leaders at organizations facing no recruitment difficulties say they have average annual security incident costs of US$12.4 million compared to US$21.2 million for those facing extreme difficulty.

Wireless resilience starts with certified expertise. Less than half of those managing wireless operations (46%) are certified in wireless technologies. Teams with deeper wireless credentials deploy modern security protocols faster and more comprehensively. Those with at least 50% of personnel certified in wireless technologies are 17 percentage points more likely to implement full WPA3 security, reducing exposure to legacy threats. They are also 17 percentage points more likely to use certificate- or profile-based authentication, minimizing access conflicts and lowering troubleshooting volume.

The implication is clear: organizations facing recruitment challenges while lacking certified talent face compounding disadvantages, including higher operational costs, greater security exposure, lower automation, and diminished capacity to modernize. Those investing in talent and certification early gain competitive advantage as complexity increases and specialized skills become essential to operational success, particularly amid the growing competition for talent.

The talent crisis reveals the interconnected nature of the wireless AI paradox. Without bringing AI to the core of wireless operations, organizations will continue to lose talent. Without the talent, strategic projects such as security modernization become harder to realize. Without modern security, incident costs rise, making it harder to invest in both talent and technology. This compounding dynamic explains why organizations must address all three barriers simultaneously to escape the paradox.

Resolving the paradox:
Recommendations for maximizing ROI

Organizations achieving strong wireless ROI are addressing operational complexity, security threats, and talent gaps together. They recognize that solving one and not the others leaves the wireless AI paradox intact. When organizations take this holistic approach to Wi-Fi they create compounding returns that multiply technology and business outcomes.

Our research supports four recommendations that, when implemented together, can break the reactive cycle and potentially unlock four times ROI advantage.

1

Revisit your Wi-Fi refresh timeline

Modern Wi-Fi enables substantially higher infrastructure scalability, AI deployments, and emerging business use cases. Organizations cannot deliver on the opportunities discussed earlier in this report without modern infrastructure foundations.

The business case is compelling: fast movers with a focus on ROI are continually investing in technology to support their business goals; for example, those with strong wireless ROI are 24% more likely to have deployed 6 GHz.

Migration should target application-driven deployment with organizations identifying the highest-impact use cases, conducting business case analysis, and deploying infrastructure in phases that align with business value delivery. This approach ensures that infrastructure investment directly supports business outcomes rather than technology for its own sake.

Organizations should prioritize Wi-Fi 6E and 7 deployments, particularly looking at how using 6 GHz can support those high-impact use cases. Early adopters of 6 GHz have been able to realize stronger ROI and higher AI deployment rates, underscoring its role in scaling modern workloads and unlocking the wireless AI paradox opportunity.

2

Get on the path to AI and automation

Organizations should prioritize building toward AI for networking to break the reactive cycle that constrains wireless teams. Organizations must go beyond simple AIOps or AI-assisted insights (e.g. alerts, recommendations) to AI-driven autonomous actions that diagnose and resolve issues at machine speed. The operational benefits of AI with autonomous actions are substantial and immediate: wireless teams free up over three hours per day, enabling them to shift from reactive operations to a more strategic, proactive approach.

Getting on the path to AI with automation delivers benefits extending beyond time savings. It improves security posture, enables faster incident resolution, and creates a foundation for future innovation. Organizations should view this not as a cost-reduction initiative but as a strategic investment in operational excellence.

Prepare your team with training and set clear rules for the role of AI in their work. Start small; implementation should begin with bringing AI features and more autonomous actions to high-impact areas first. Ticket management and network automation are starting points for immediate value. Mature AI networking strategies can then lead to zero-touch, fully autonomous network operations capable of supporting increasingly demanding workloads. With AI itself cited as the top field drawing talent away from wireless, AI for wireless can stem the flow and ease time constraints. Organizations with fully automated AI-driven operations are three times more likely to forecast significant acceleration in ticket resolution, four times more likely to describe networks as very simple, and almost twice as likely to describe infrastructure as fully scalable. Organizations deploying Ai with autonomous actions across key functions achieve substantially higher ROI.

Investing in AI for networking helps to stitch together context from across the network providing end-to-end visibility and observability. With 87% of organizations

reporting visibility gaps that impair effective Wi-Fi troubleshooting, it's critical to address this blind spot. Organizations should explore how they can better connect visibility across IT and lines of business, achieving end-to-end insight across clients, applications, cloud services, and network packets. The operational benefit is immediate, eliminating wireless scapegoating and focusing troubleshooting on actual problem sources by the right team.

For wireless teams, the first three areas to target are: improving visibility into client device behavior, application and cloud service performance, and packet-level analytics. This three-dimensional approach enables rapid root-cause identification and faster incident resolution.

End-to-end visibility supercharges operations. AI requires rich data inputs to make effective operational decisions, so organizations should view broad and deep visibility as a prerequisite for AI for networking.

3

Prioritize holistic security modernization

Wireless security modernization should be treated as a strategic initiative aligned with broader enterprise security posture. Organizations should treat wireless security as a priority for managing business risk, which means looking at wireless security in the context of the entire business. This would enable security investment to compete for capital alongside other strategic IT priorities.

Organizations that want to minimize exposure to losses from security incidents should accelerate deployment of modern wireless security protocols, particularly full WPA3 and certificate- or profile-based authentication. Organizations with more modern security postures also tend to see a better return on their wireless investments. For example, organizations using full WPA3 have a 50% likelihood of predicting strong ROI compared to 19% among those without WPA3.

Legacy infrastructure needs to be addressed systematically rather than be allowed to persist indefinitely. Organizations should align their security lifecycle planning with new infrastructure investments.

Organizations implementing modern authentication and encryption protocols demonstrate superior outcomes: four times more likely to predict declining security failures and experience 18% lower financial losses from incidents. These findings show that modern wireless security is not merely more secure but also delivers measurable financial benefits by preventing incidents that would generate losses.

Wi-Fi security alone is insufficient. To secure the network effectively, organizations must:

1. Implement network segmentation to dynamically verify user identity, enforce security policies, and organize clients into appropriate groups.
2. Extend visibility beyond access points to monitor and protect the entire network from wireless threats (reinforcing recommendation 2).

These measures ensure comprehensive network protection, not just wireless access control.

4

Build and enhance wireless talent

Organizations should invest in developing their talent through training, certification, and career advancement opportunities. This investment creates measurable impact across security outcomes, incident response times, and job satisfaction.

Investment should target advanced skills, including optimizing Wi-Fi 6E and 7 standards, security protocol modernization (WPA3), AI implementation, and automation platform expertise.

Organizations also need to recognize wireless specialization as a strategic career path and compensate

accordingly, given that the impact of certification is quantifiable. Organizations with at least 50% of personnel certified in wireless technologies are 17% more likely to fully implement WPA3. They are also 17 percentage points more likely to use profile- or certificate-based authentication.

Training and certification investment compounds over time, with organizations that start talent development programs early gaining competitive advantage as complexity increases and specialized skills become essential to operational success. Those delaying investment until talent shortages become acute face substantially higher hiring costs and operational disruption.

Beyond direct investment in talent organizations should look at investment in AI with automation (specifically Recommendation 2) to enhance wireless talent. AI has shown to reduce time spent on reactive tasks thus improving morale. This approach can also help reverse the flow of talent leaving wireless to go to AI while improving the changes of recruiting the next generation of wireless professionals.

The multiplier effect: Why integrated implementation matters

These four recommendations compound when they are implemented together:

- Modern infrastructure without AI and automation = overwhelmed teams
- AI without security modernization = efficiently managed vulnerable networks
- Security modernization without talent = protocols that can't be properly implemented
- Talent development without modern infrastructure = skilled teams managing outdated systems

As this report clearly outlines, organizations that address all dimensions simultaneously are four times more likely to achieve strong ROI compared to those neglecting these areas.

The AI paradox will not resolve itself. Organizations that act decisively and holistically in 2026 will establish competitive advantages that compound over the years to come. Organizations that delay will find themselves trapped in reactive cycles, bleeding budget to security incidents, and unable to capitalize on AI-driven transformation while competitors pull ahead.

The choice, and the window to choose, is now.

Methodology

This research comprised interviews with 6,098 organizations in 30 markets conducted in November 2025 by Sandpiper Research and Insights.

Research Scope

Respondent Profile: Interviews were conducted with 6,098 wireless decision makers and technical specialists in organizations with at least 250 employees. Six in 10 (61%) respondents work in organizations with annual turnover of at least US$100 million.

Americas — 22%

Europe, Middle East, and Africa — 28%

Asia Pacific — 50%

Geographic Coverage: Research covered 30 markets including Australia, Brazil, Canada, Chinese Mainland, France, Germany, Hong Kong, India, Indonesia, Italy, Japan, Malaysia, Mexico, Netherlands, New Zealand, Philippines, Poland, Saudi Arabia, Singapore, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Thailand, United Arab Emirates, United Kingdom, United States, and Vietnam.

Industry Representation: Respondents worked across a range of industries including Business Services, Construction, Education, Engineering, Design and Architecture, Financial Services, Government and Public Services, Healthcare, Manufacturing, Media and Communications, Natural Resources, Real Estate, Restaurant Services, Retail, Technology Services, Transportation, Travel Services, and Wholesaling.

Timing: Research was conducted in November 2025.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**

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

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

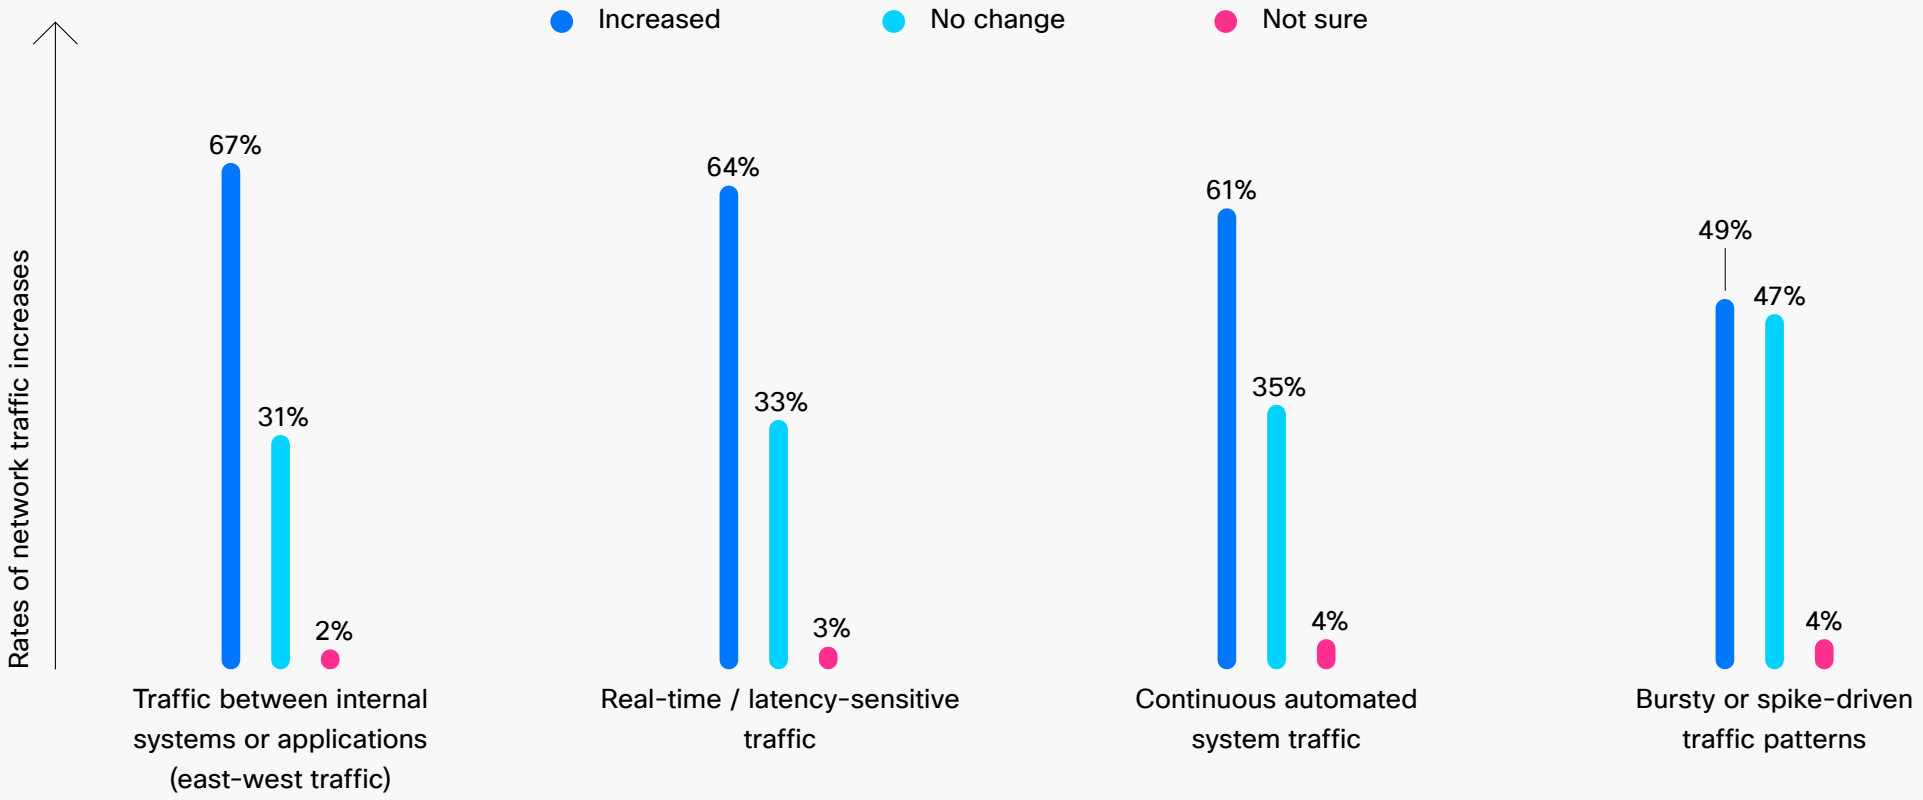

Those changes become even more significant as organizations move beyond genAI experimentation and deeper into agentic AI capable of autonomous action.

“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:

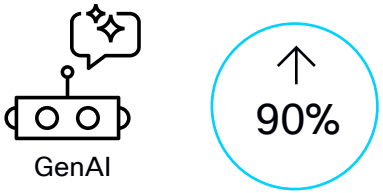

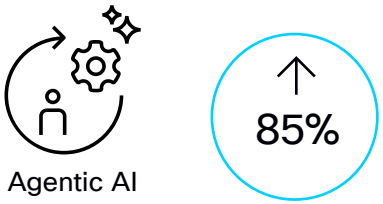

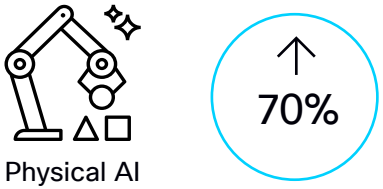

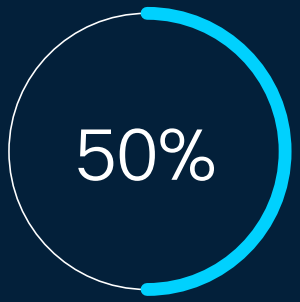

of AI-driven demand is concentrated in wireless networks within campus locations (Wi-Fi).

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

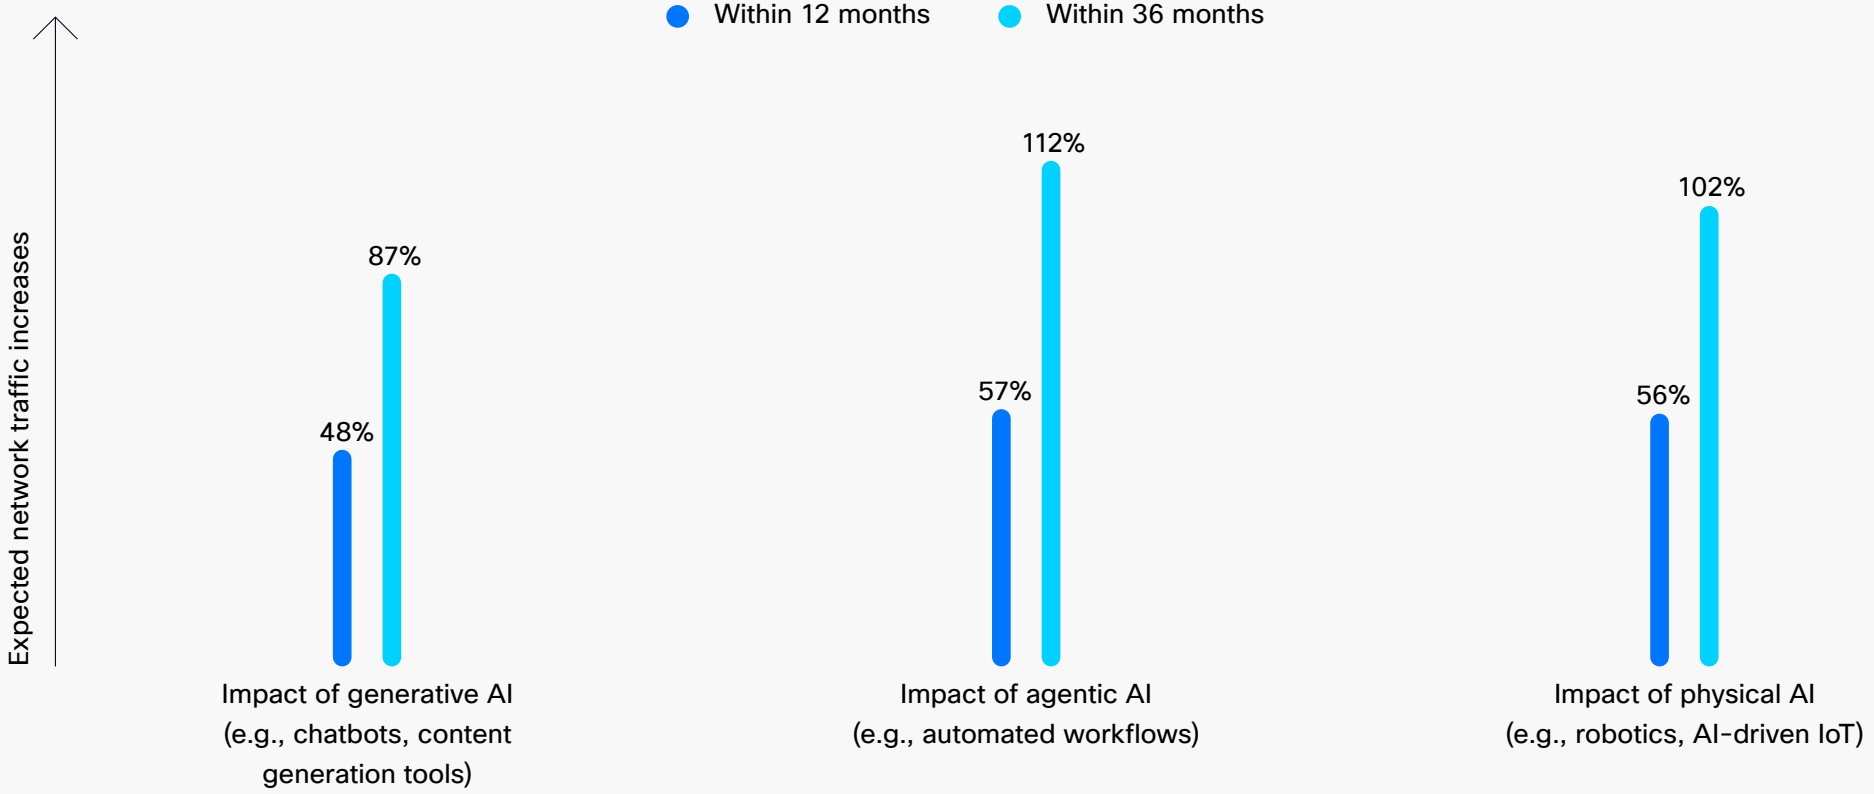

Combined figures reflect the compounded effect of the three projected AI traffic increase estimates rather than a simple average; layered growth can produce values above 100%, where 100% represents a doubling of current traffic.

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:

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

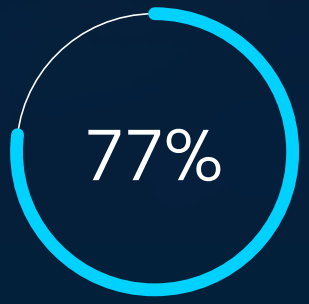

said the technology has already expanded their attack surface in the last 12 months

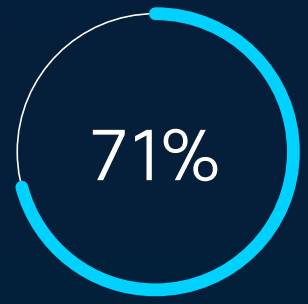

believe that threats are evolving faster than existing controls can adapt

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.

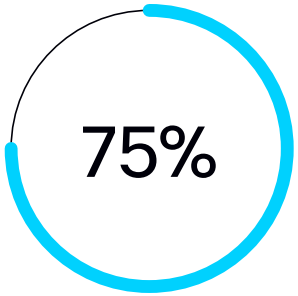

of IT leaders agree there are higher long-term costs due to reactive upgrades or remediation.

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

75%

Higher long-term costs due to reactive upgrades or remediation

73%

Inability to meet customer expectations

72%

Missed business opportunities

72%

Increased security risks or an expanded attack surface

71%

Decreased operational efficiency

70%

Business reputation risks

67%

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.