IDC PERSPECTIVE

AI and Cloud Connectivity — Impact on Networking and Connectivity Investments

Paul Hughes

EXECUTIVE SNAPSHOT

Across all verticals and all business sizes, AI is rapidly reshaping enterprise infrastructure requirements and investment timelines. As AI workloads as a share of total traffic across the network, the expectations for explosive data growth are placing unprecedented demands on cloud connectivity, network performance, and data movement. As enterprises scale AI workloads, particularly generative AI (GenAI) and real-time inference, the need for high-performance, low-latency, and resilient connectivity has become mission-critical.

This document analyzes how AI adoption is influencing cloud connectivity requirements across enterprises. It also identifies the most strategically important verticals, evaluates where network investments are being concentrated, and provides recommendations to IT buyers on how to align their network strategy to cater to AI workloads.

Key takeaways

- Enterprise AI adoption will be pragmatic and incremental, shaped by quantifiable ROI, cost, energy, regulatory, and environmental constraints, with success favoring enterprises that scale connectivity based on proven demand rather than hype.
- AI workloads and cloud connectivity are not just increasing traffic; they are redefining traffic patterns, creating a structural gap between legacy network design and AI requirements.
- Bandwidth demand is already increasing and will continue to grow, primarily driven by agentic AI, inferencing, and analytical/ML workloads. The number of enterprises reporting an increase in bandwidth demand was much smaller than the previous year.
- Enterprises are already responding to rapidly escalating bandwidth requirements, highlighting a direct correlation between AI adoption and network readiness. At the same time, investment patterns indicate a shift beyond isolated, reactive upgrades toward more coordinated connectivity architectures that integrate cloud, edge, and software-defined networking.

June 2026, IDC #US54567826

- As enterprises adopt distributed architectures, the challenge is no longer connectivity availability, but managing performance, visibility, and control across fragmented environments. The network is becoming harder to operate, even as it becomes more critical.
- Managed connectivity and service provider partnerships will become increasingly essential as enterprises rely on external expertise to align their network architecture with growing AI demands.

Recommended actions

- Enterprises should move away from legacy location-based network design toward modern architectures built for managing north-south and east-west traffic, edge-to-cloud data flows, and support AI traffic patterns.
- Enterprises should focus on delivering high-performance connectivity by prioritizing fiber backbone, cloud interconnect, and last-mile connectivity over generic bandwidth expansion. AI workloads require predictable latency and reliability, not just higher throughput.
- North-south traffic (between cloud and enterprise datacenter/edge) is a critical flow to address; enterprises need the right technology areas (cloud connectivity, cloud interconnect) to support this.
- East-west traffic (within and across datacenters and cloud environments) remains equally important, as AI training and distributed processing generate large volumes of internal data movement between applications, datasets, and compute clusters; enterprises must ensure high-speed, low-latency interconnect within their environments to support these workloads efficiently.
- Build an integrated, multilayer connectivity stack to adopt a coordinated approach combining cloud interconnect, edge infrastructure, and SD-WAN/network-as-a-service (NaaS) orchestration. Competitive advantage will come from how well these layers work together, not standalone deployments.
- Strengthen observability and automation capabilities by implementing a single pane of glass for visibility, monitoring, and control across hybrid environments. AI-scale operations require real-time insights and automated remediation.
- Enterprises should adjust their network architecture based on industry-specific needs and use cases. Tailor network investments based on industry dynamics such as low latency for financial services, compliance for healthcare, scale for retail, and real-time control for manufacturing, rather than applying a one-size-fits-all model.

SITUATION OVERVIEW

AI and the state of the enterprise

The enterprise timelines for digital business evolution and road map have accelerated, driving enterprises to evaluate investment and management priorities for future success. Enterprises were already seeing year-over-year (YoY) increases in bandwidth demand even before AI became ubiquitous. IDC's 2025 *Future Enterprise Connectivity Infrastructure and Services Survey* shows that 37% of enterprises globally saw their bandwidth requirements increase by more than 50% YoY, including 9% with requirements more than doubling (see Figure 1). Financial services enterprises saw the steepest demand increases, with 47.5% reporting bandwidth growth exceeding 50% YoY, followed by retail/hospitality, highlighting the impact of digital business initiatives, automation, and accelerated AI adoption.

Figure 1

Enterprises experiencing bandwidth surge globally — demand driven by AI workloads, cloud migration, and digital transformation

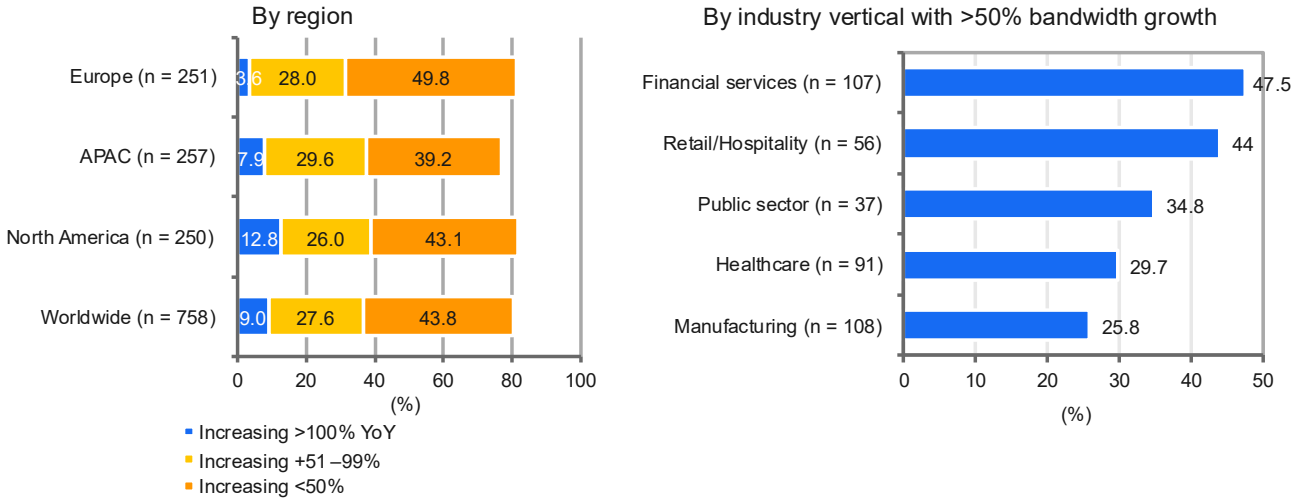

Source: IDC's *Future Enterprise Connectivity Infrastructure and Services Survey*, August 2025

Across the enterprise landscape, adoption of new AI and cloud technologies is having a dramatic impact on the timelines and strategies moving forward. IDC's 2025 *Future Enterprise Connectivity Infrastructure and Services Survey* showed that over 50% of enterprises stated that AI was a top area of investment and is having a significant

transformational impact on their connectivity strategy and road map, with slight differences across different verticals (see Figure 2).

Figure 2

Impact of AI on connectivity and the network

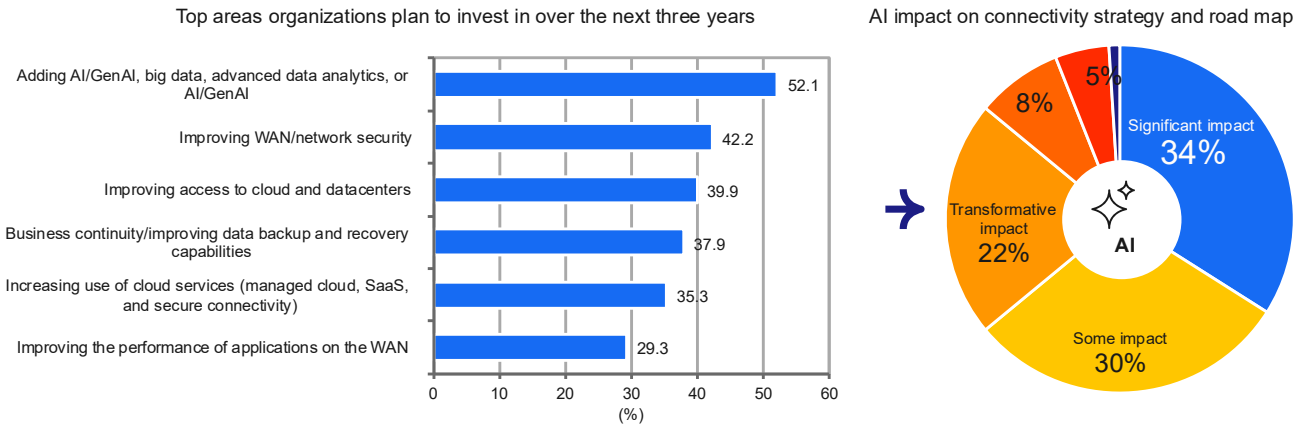

Source: IDC's *Future Enterprise Connectivity Infrastructure and Services Survey*, August 2025 (n = 758)

Enterprise AI adoption has accelerated due to both predictive and productive factors. IDC has seen enterprises increase usage of AI applications that support everything from software code generation to marketing documentation over the past two years. This, along with automation and predictive analytics that leverage the increased availability of network, system, and process data, has put significant pressure on networks to scale to meet the demands of data traffic and workloads. Training workloads require massive data ingestion and compute clusters, whereas inference workloads demand low latency and distributed deployment. As such, the survey data should not come as much of a surprise, as AI workloads are expected to drive exponential growth in east-west and edge-to-cloud traffic. The nature of those workloads put increased pressure on network and IT staff to address:

- **Latency sensitivity**. Ultra-low latency is critical for real-time AI applications.
- **Bandwidth demand**. The movement of massive datasets and model outputs from the corporate network to the cloud and back requires high-bandwidth connectivity.
- **Regulatory constraints**. For certain industries and regions, compliance with data localization requirements must be part of the strategy.
- **Cost pressures**. With AI ROI a top concern, network designs and implementations must focus on efficient data movement and compute placement.

- **Vertical-specific needs**. For traditional enterprises in areas, such as retail, financial services, healthcare, or manufacturing, every AI investment still needs to justify itself against a concrete return. Enterprises that benefit most are those that build their AI strategies on pragmatic foundations: clear use cases, measurable ROI, and infrastructure that scales with proven demand rather than projected hype.

The AI era will bring generational change to how enterprises operate, but that change will arrive incrementally and with friction, not overnight. Much of the capital flowing into AI infrastructure, including datacenters, semiconductors, and compute capacity, is currently concentrated within a relatively small ecosystem of hyperscalers, chipmakers, and platform providers that are investing heavily in one another. To a significant extent, these large technology companies are also each other's biggest AI customers and strategic partners. This circular investment dynamic is helping sustain the current pace of infrastructure expansion, but it has yet to reflect a similar speed of adoption across the wider enterprise economy so far.

Recent IDC enterprise interviews highlighted a current lack of trust around data security, cost, ROI concerns, and businesses simply not ready to make the jump yet. That said, strategic alignment between AI deployment and network architecture is slowly becoming a competitive business differentiator, driving future network investments toward supporting edge infrastructure, private connectivity, and high-capacity fiber. It also has created an opportunity for connectivity providers to take the lead in helping enterprises plan, deploy and manage their AI enabled networks.

Impact of AI workloads on enterprise cloud connectivity

Cloud architectures continue to undergo structural transformations within the enterprise as data volumes increase and performance needs evolve, both locally and offsite. Enterprises are now migrating from centralized IT models to a distributed architecture environment that can better respond to the growing demands of AI, data-intensive applications, and digital services. Traditional models built around a single or limited number of core datacenters are no longer sufficient to support modern workloads that require low latency, high availability, and real-time responsiveness. As a result, enterprises are adopting architectures that distribute compute, storage, and networking resources across multiple environments, including public cloud regions, on-premises systems, and edge locations. This shift enables enterprises to place workloads closer to users and data sources, improving performance while reducing bottlenecks associated with centralized processing.

Where many enterprises still struggle is in the internal alignment of network architecture to the needs of AI and the evolving role of cloud connectivity. IDC's *Enterprise Connectivity Infrastructure and Services Survey* shows that nearly 30% of enterprises rank the alignment of network infrastructure with AI applications as a top challenge, and aligning the network with cloud connectivity requirements as one among the Top 5 challenges (see Figure 3).

Figure 3

Top-ranked connectivity challenges by enterprises

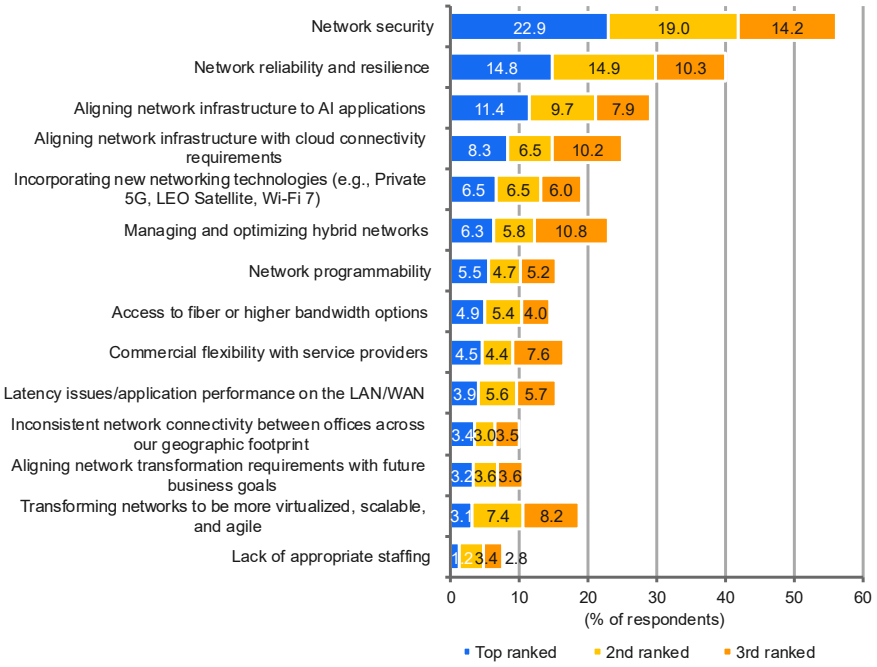

Network security remains the top challenge as diverse threats and a larger attack plane become targets for bad actors.

Enterprises continue to seek multinetwork access to address network resiliency and reliability challenges.

Source: IDC's *Future Enterprise Connectivity Infrastructure and Services Survey,* August 2025 (n = 758)

Compounding these challenges is the transition to latency-sensitive applications, particularly those powered by AI inference. Use cases such as real-time analytics, autonomous systems, and interactive digital experiences require rapid data processing that cannot tolerate the delays introduced by distant datacenters. Distributed architectures address this challenge by leveraging edge computing, where data is processed locally or regionally before being sent to centralized systems for aggregation or training. This approach not only enhances user experience but also optimizes bandwidth utilization by reducing the need to transmit large volumes of raw data across long distances.

Another important factor is the increasing adoption of hybrid and multicloud strategies. Enterprises are no longer relying on a single cloud provider or deployment model; instead, they are distributing workloads across multiple platforms to improve resilience, avoid vendor lock-in, and meet regulatory requirements. Hybrid architectures allow enterprises to maintain sensitive workloads on-premises while leveraging the scalability of public cloud environments. Multicloud strategies further extend this flexibility by enabling workload

portability and redundancy across different providers. However, these approaches also introduce complexity, requiring advanced networking solutions to ensure seamless connectivity and consistent performance across environments.

Finally, the migration to distributed architectures is being shaped by economic and regulatory considerations. Data sovereignty laws often require enterprises to store and process data within specific geographic regions, necessitating localized infrastructure. At the same time, enterprises are seeking to optimize costs by strategically placing workloads where compute and storage are most efficient. Distributed architectures support this by enabling dynamic workload placement and more efficient use of network resources. As a result, enterprises are investing in technologies such as software-defined networking, edge platforms, and interconnection services to manage this complexity and fully realize the benefits of a distributed operating model. Cloud architectures are undergoing a structural transformation within the enterprise, as data volumes increase and performance needs evolve both locally and offsite. Enterprises are moving away from centralized models toward distributed environments that include:

- Edge computing, where processing occurs closer to users or devices
- Hybrid cloud, combining on-premises and public cloud resources
- Multicloud, distributing workloads across multiple providers

These architectures increase flexibility and resilience but also introduce new challenges in data movement, orchestration, and network performance.

Enterprise connectivity challenges and cloud connectivity requirements for AI

As AI workloads scale, enterprises are confronting a fundamental mismatch between legacy network architectures and the new traffic patterns generated by AI. Traditional networks were designed for east-west enterprise traffic and for predictable and elastic, north-south traffic, but not for the intense, continuous north-south data flows that characterize AI workloads between enterprise environments and hyperscale clouds. This is creating a structural gap between network capability and AI demand that is already becoming a constraint on adoption.

- **Infrastructure bottlenecks**. These can no longer be treated only as latency risk; they are already constraining AI scalability (as shown in Figure 3). IDC's 2025 *Future Enterprise Connectivity Infrastructure and Services Survey* data shows that 29% of enterprises struggle to align networks with AI workloads, and 24% face challenges with cloud connectivity, with even higher pressure in financial services and manufacturing. The core issue enterprises face is not just capacity but an architecture mismatch. Existing networks were built for internal traffic, not for the high-volume, cloud and AI-centric data flows. As a result, incremental upgrades are

insufficient, and enterprises need to step up their investments in fiber, cloud interconnect, and last-mile performance. Without this, technological advancements in AI will outpace network readiness, limit performance, and reduce business impact.

- **Security and data privacy concerns**. As enterprises transmit increasing volumes of sensitive data between on-premises systems and cloud environments to support AI training, inferencing, and analytics, the attack surface expands. Data localization and regulatory compliance requirements, particularly in financial services, healthcare, and the public sector, further complicated cloud connectivity strategies, requiring enterprises to build compliant pathways for data movement without sacrificing performance.
- **Environmental and regulatory pressures on AI datacenters**. Globally, environmental impacts, cost to local residents, and "NIMBY" are becoming major constraints on AI infrastructure growth. High energy usage, water consumption, and rising costs of running large-scale datacenters are making regulators, communities, and enterprises more cautious. As a result, decisions around infrastructure placement, cloud partners, and connectivity strategy must now balance performance, cost, sustainability, and regulatory risk.
- **Operational complexity**. The shift toward multicloud architecture is introducing significant operational complexity. IDC's 2025 *Future Enterprise Connectivity Infrastructure and Services Survey* shows 24% of enterprises rank aligning the network to cloud connectivity requirements as one of the Top 3 connectivity challenges. Managing and optimizing hybrid networks is also a major pain point (22.8% overall, rising to 29% for financial services), reflecting the growing complexity of multicloud and hybrid deployments. At the same time, AI is emerging as a key tool to address this complexity. AI-driven network management automation is the most prominent area of impact, with 26% of respondents identifying it as the primary benefit and an additional 24% as the second most significant. This highlights a strong focus on improving operational efficiency through automation of routine tasks, reduction of manual intervention, and the adoption of self-healing network capabilities. However, enterprises must still manage multiple network environments simultaneously while ensuring consistent performance, security, and visibility.

Connectivity requirements for AI at scale

To support AI workloads at scale, enterprises must address four interconnected categories of data movement, each placing distinct demands on connectivity infrastructure.

- **Network performance, low latency, and bandwidth**. Real-time AI applications, from inferencing to agentic AI, are a cause of network delays impacting performance. Enterprises must invest in the bandwidth capacity and low-latency infrastructure to ensure AI applications can perform as designed. Bandwidth demands are increasing dramatically across all AI workload types.

- **East-west traffic and data in motion**. AI training and inferencing workloads generate substantial east-west traffic across datacenters and cloud regions. As enterprises distribute compute across hybrid environments, data in motion must be managed with both performance and security in mind. Data is now routinely protected in transit using encryption protocols such as Transport Layer Security (TLS), which has become a baseline requirement. This shifts the focus toward network architectures that can handle encrypted traffic at scale without introducing latency or throughput bottlenecks. As a result, enterprises must prioritize high-performance datacenter interconnects, optimized east-west traffic flows, and resilient network design to support secure, reliable, and low-latency data exchange across distributed AI environments
- **Edge-to-cloud traffic**. IoT devices, sensors, and edge computing nodes are generating real-time data that feeds AI models running in the cloud. As more intelligence moves to the edge and as edge infrastructure investment grows, the connectivity pathways between edge and cloud become mission-critical arteries. Enterprises that fail to invest in this layer will find their AI insights delayed, degraded, or incomplete.
- **User-to-inference traffic and last-mile performance**. When end users interact with AI applications, whether via enterprise tools, customer-facing services, or automated workflows, the response must be near-instantaneous. Last-mile connectivity quality directly determines the user experience of AI-powered services. Datacenter interconnect quality and the reliability of the final connection to cloud-hosted inference engines will increasingly determine whether AI investments deliver their promised returns**.**

Impact on vertical industries

Enterprises across industries are expecting AI to have a significant impact on their connectivity requirements. In IDC's 2025 *Future Enterprise Connectivity Infrastructure and Services Survey,* more than half of respondents across most verticals report that AI will be a major or transformative driver of network change, highlighting a strong level of awareness and forward planning.

Figure 4

AI Impact on enterprise connectivity by industry

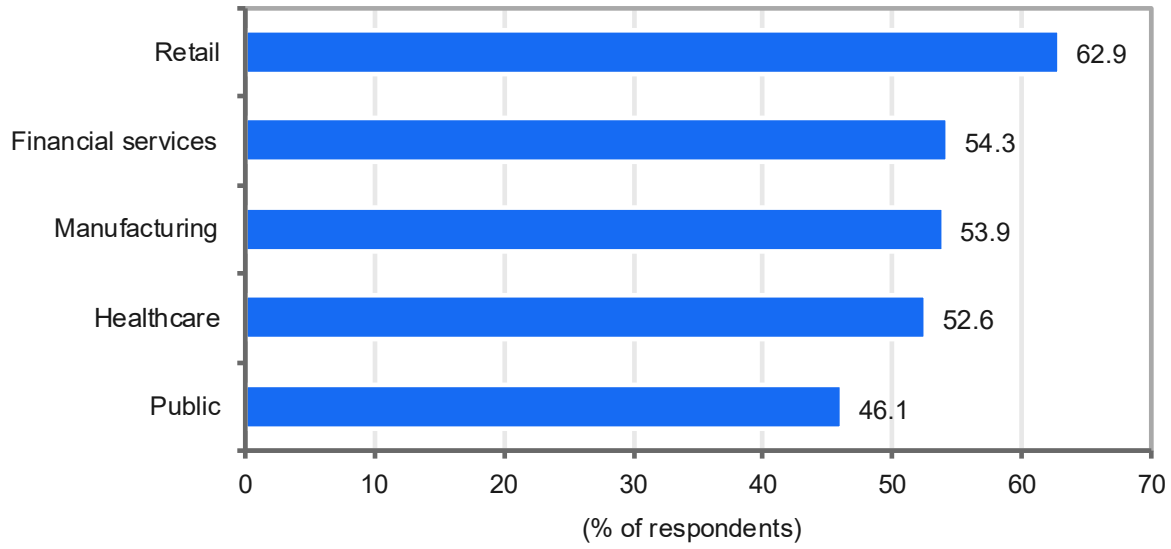

Source: IDC's 2025 *Enterprise Connectivity Infrastructure and Services Survey*, August 2025, (n = 758)

As shown in Figure 4, Retail leads with 62.9% of respondents indicating a significant impact, followed by financial services (54.3%), manufacturing (53.9%), and healthcare (52.6%), while the public sector reports a relatively low but still notable 46.1%. These differences reflect how AI adoption varies by industry in terms of data intensity, distribution of workloads, and the importance of real-time processing.

The variation across industries highlights that AI-driven connectivity requirements are not uniform, as each vertical has distinct use cases, data flows, and latency requirements that shape both the scale and urgency of investment. These differences are important for enterprises planning network upgrades and for service providers tailoring solutions to vertical-specific needs.

- **Financial services**. Financial services show a high level of urgency in connectivity transformation, driven by the growing use of AI in latency-sensitive environments, such as algorithmic trading, real-time fraud detection, risk analytics, and transaction processing, where real-time data directly increases traffic. These use cases require continuous data exchange and near-instant processing, placing sustained pressure on network performance. As AI becomes more embedded in decision-making workflows, firms have no choice but to prioritize low-latency connectivity, cloud interconnects, and secure data movement to support distributed AI models while meeting strict regulatory and compliance requirements.

- **Healthcare**. Healthcare connectivity requirements are being reshaped by the adoption of AI in diagnostics, medical imaging, and clinical decision support to improve patient outcomes. Each of the above-mentioned use cases is highly data-intensive and increasingly real-time, but network evolution is moderated by strict regulatory controls and risk sensitivity around patient data. As a result, enterprises are investing in private and hybrid connectivity models that allow controlled data movement between on-premises systems and cloud-based AI platforms. The shift toward real-time processing of imaging and patient data is driving demand for highly reliable, secure, and low-latency network architectures.
- **Retail and hospitality**. Retail and hospitality are experiencing strong connectivity demand driven by distributed AI use cases such as personalized customer experiences, demand forecasting, dynamic pricing, and in-store analytics. These workloads operate across thousands of locations and endpoints, requiring networks that can support real-time data synchronization and consistent performance at scale. The challenge is less about core bandwidth capacity and more about managing highly distributed environments, driving investments in SD-WAN, edge compute, and centralized orchestration to ensure seamless connectivity across stores, channels, and platforms.
- **Manufacturing**. Manufacturing connectivity is increasingly shaped by the integration of AI into production environments, particularly through automation, predictive maintenance, computer vision for quality inspection, and process optimization. These use cases generate continuous, high-frequency data streams from machines and sensors, requiring a combination of edge processing for real-time control and cloud connectivity for analytics. As a result, manufacturers are prioritizing low-latency, deterministic networks that can support operational reliability. Network performance becomes directly tied to production outcomes, making connectivity a critical component of industrial operations rather than a supporting IT layer.
- **Public sector**. Public sector connectivity requirements are influenced by large-scale, distributed operations and growing adoption of AI in areas such as smart cities, first response and public safety analytics, citizen services, and infrastructure monitoring. While demand is increasing, investment is often constrained by tighter budgets, longer procurement cycles, legacy systems, and data sovereignty requirements. This creates a need for highly secure, resilient, and compliant network architectures that can support AI workloads while maintaining strict control over data. As AI adoption expands, connectivity strategies are evolving to balance performance with governance across complex and widely distributed environments.

Enterprise investment expectations by technology

Enterprise connectivity investment priorities are increasingly shaped by the need to support AI-driven, distributed, and latency-sensitive workloads, pushing enterprises to rethink

network architectures beyond traditional connectivity models. The focus is shifting toward integrated, high-performance ecosystems that combine cloud interconnection, edge processing, and flexible consumption-based networking. This reflects a broader transition from static infrastructure to programmable, scalable, and multienvironment connectivity platforms that can dynamically support evolving AI use cases.

As shown in Figure 5, companies are prioritizing investments across different technologies. Cloud interconnect emerges as the highest priority, with 77.0% of enterprises increasing investment, highlighting the central role of hyperscale cloud connectivity in AI deployments. This is followed by strong investment in private networks (61.6%) and mobile/wireless WAN (61.3%), indicating growing demand for controlled, flexible connectivity across distributed environments.

Figure 5

Enterprise connectivity investment outlook, with the percentage of enterprises expecting to increase investment over the next 12 months

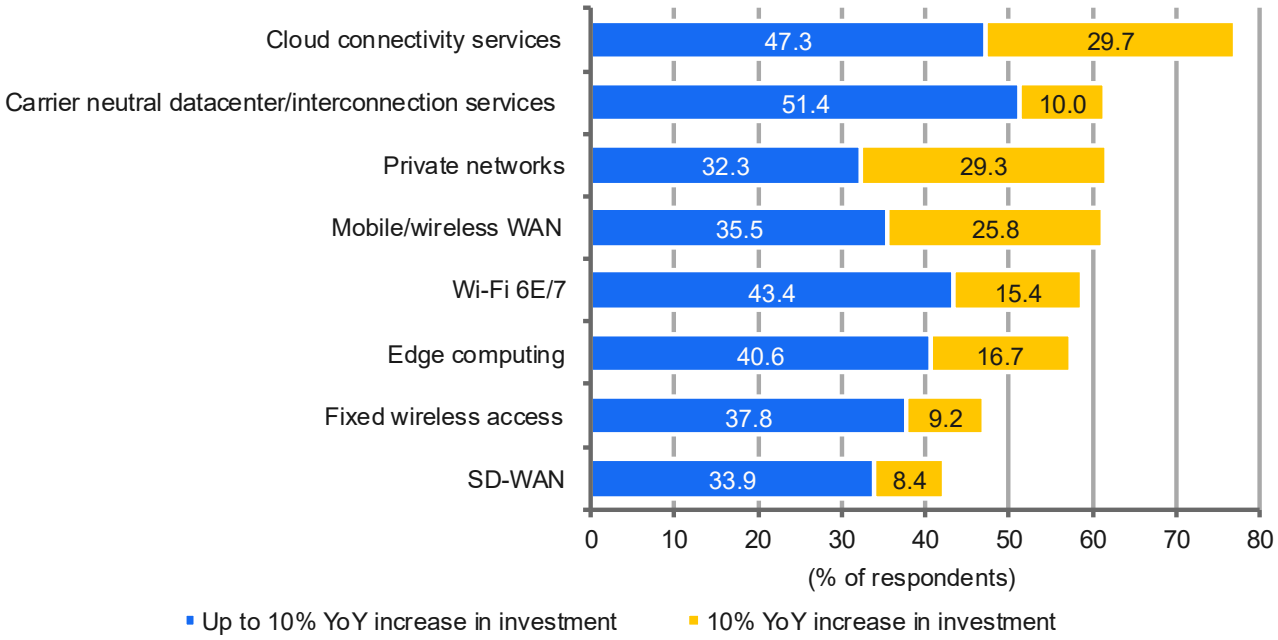

Source: IDC, 2026

At the same time, investment in Wi-Fi 6E/7 (58.8%) and edge computing (57.3%) reflects the need to modernize last-mile and edge infrastructure to support real-time, data-intensive applications closer to users and devices. Interconnection platforms (61.4%) also show solid

momentum, underscoring the importance of neutral ecosystems that enable seamless data exchange across multicloud and hybrid environments.

While SD-WAN and fixed wireless access continue to be important foundational components, investment is increasingly shifting toward higher-value layers such as cloud interconnect and edge computing that directly support AI-driven workloads. This reflects a transition from access and transport-focused investments toward architectures optimized for data movement, low latency, and distributed processing.

ADVICE FOR THE TECHNOLOGY BUYER

- **Align AI strategy with network architecture**. Enterprises must tightly align their AI strategies with underlying network architecture, as AI cannot scale without a robust and responsive connectivity foundation. Every AI initiative relies on reliable bandwidth, low latency, and seamless cloud connectivity. As workloads grow more distributed and data-intensive, networks must evolve to handle increased traffic, dynamic routing, and real-time data exchange. Conducting a comprehensive network assessment is essential to identify capacity constraints, latency bottlenecks, and integration gaps. Enterprises should also design their networks considering AI-specific traffic patterns and ensure that they can support both centralized training workloads and distributed inferencing at scale.
- **Invest in hybrid and edge infrastructure**. To effectively support AI at scale, enterprises must invest in hybrid and edge infrastructure that brings compute closer to where data is generated and consumed. Edge computing reduces latency, improves responsiveness, and minimizes bandwidth costs by enabling real-time processing at or near the source. As AI workload patterns evolve, a hybrid approach combining on-premises infrastructure, private networks, public cloud, and edge nodes ensures that flexibility, scalability, and resilience are required. Connectivity should be viewed as a portfolio of integrated capabilities, including fiber, SD-WAN, mobile networks, and cloud interconnect, rather than a single fixed pipeline. This approach allows enterprises to dynamically balance workloads across environments, optimize performance, and maintain operational continuity even in complex, distributed ecosystems
- **Establish a single pane of glass for network monitoring, observability, and resilience**. As AI and cloud connectivity become increasingly embedded in critical business processes, ensuring network reliability and performance is paramount. Enterprises should establish a unified view for monitoring and observability, providing end-to-end (E2E) visibility across datacenters, cloud environments, and edge infrastructure. Advanced observability tools powered by AI and analytics can help detect anomalies, predict performance degradation, and enable proactive remediation before issues impact operations. Given the growing complexity of hybrid

and multicloud environments, enterprises should also consider managed connectivity and observability services to reduce operational overhead and ensure consistent performance. Strong observability not only improves resilience but also enables enterprises to maintain service levels and trust.

Advice for the technology Supplier

- **Think vertically and develop industry-specific connectivity solutions.** Technology suppliers should adopt a vertical-first approach, as traditional connectivity is no longer sufficient in an AI-driven enterprise environment. Different industries are now demanding highly tailored solutions; for instance, financial services require ultra-low latency and regulatory compliance, the healthcare sector demands secure hybrid connectivity, retail depends on distributed edge networks, and the manufacturing sector increasingly relies on OT/IT convergence. Suppliers should align their go-to-market (GTM) strategies around AI-led use cases within each vertical, positioning connectivity and infrastructure as enablers of business outcomes such as fraud detection, predictive maintenance, and real-time analytics. Moving forward, vendors should focus on delivering pre-integrated, industry-specific solutions that combine connectivity, security, and edge capabilities to address both operational and regulatory requirements.
- **Expand edge infrastructure offerings and managed connectivity services.** Technology suppliers should invest in expanding edge nodes, private network capabilities, and edge computing platforms that bring processing closer to users and devices, enabling real-time AI inference. At the same time, demand for managed connectivity services is rising, as enterprises seek to offload network complexity and focus on core business initiatives. Providers that offer integrated, service-level agreement (SLA)-backed solutions combining SD-WAN, edge compute, cloud interconnect, and observability into a single managed offering will have more advantage to be best positioned to capture growth. Building hybrid, multinetwork architectures that integrate fiber, wireless, satellite, and cloud will further enhance flexibility and resilience, creating a differentiated value proposition.
- **Strengthen partnerships with cloud providers.** As enterprise AI adoption accelerates, hyperscalers are becoming essential infrastructure partners rather than competitors. Technology suppliers must deepen collaboration with hyperscalers and develop strong co-sell, co-deploy, and integration models. This includes expanding capabilities in private cloud interconnect, cloud on-ramps, and multicloud networking to ensure seamless and secure workload mobility. Additionally, suppliers should invest in joint solutions that integrate connectivity with cloud-native AI services, data pipelines, and developer ecosystems. Those that position themselves as trusted partners capable of orchestrating connectivity across both network and cloud

domains will be better equipped to support enterprise AI ambitions and capture long-term growth opportunities.

LEARN MORE

Related research

- *MWC26: The AI Continuum* (IDC #EUR154461426, April 2026)
- *2026 Enterprise Connectivity Investment Analysis by Region — Fixed Connectivity* (IDC #US54365826, March 2026)
- *Enterprise Connectivity Technology Supplier Heat Map* (IDC #US54376326, March 2026)
- *2026 Enterprise Connectivity Investment Analysis by Region — Wireless/Mobile* (IDC #US54353626, March 2026)
- *The Year Ahead: Top Enterprise Connectivity Trends to Look for in 2026* (IDC #US54260526, February 2026)
- *NRF 2026 Highlights — Enterprise Connectivity (Indirectly) Becomes Cool Again as AI Takes Center Stage* (IDC #US54242526, February 2026)

Synopsis

This IDC Perspective examines how the rapid growth of AI workloads is reshaping enterprise network infrastructure requirements and connectivity investment priorities. Drawing on IDC's 2025 *Enterprise Connectivity Infrastructure and Services Survey*, the report shows that bandwidth demand is accelerating across all industries, AI traffic is on track to overtake traditional network traffic by 2027–2028, and enterprises are already increasing investment in SD-WAN, network-as-a-service (NaaS), edge computing, and cloud interconnect. This document also identifies vertical-specific connectivity requirements across financial services, healthcare, retail, manufacturing, and the public sector, and provides actionable guidance for technology buyers and suppliers navigating the AI connectivity imperative.

"AI does not exist without connectivity. As enterprises accelerate their AI road maps, the network is no longer an infrastructure but is considered a strategic asset. Enterprises that align their connectivity strategy with their AI ambitions today will be the ones that deliver on the promise of AI tomorrow," says Paul Hughes, research director, IDC.

ABOUT IDC

International Data Corporation (IDC) is the premier global market intelligence, data, and events provider for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries. IDC's analysis and insight help IT professionals, business executives, and the investment community make fact-based technology decisions and achieve their key business objectives.

Global headquarters

One Beacon Street
Suite 33100
Boston, MA 02108
USA
508.872.8200
X: @IDC
blogs.idc.com
www.idc.com