

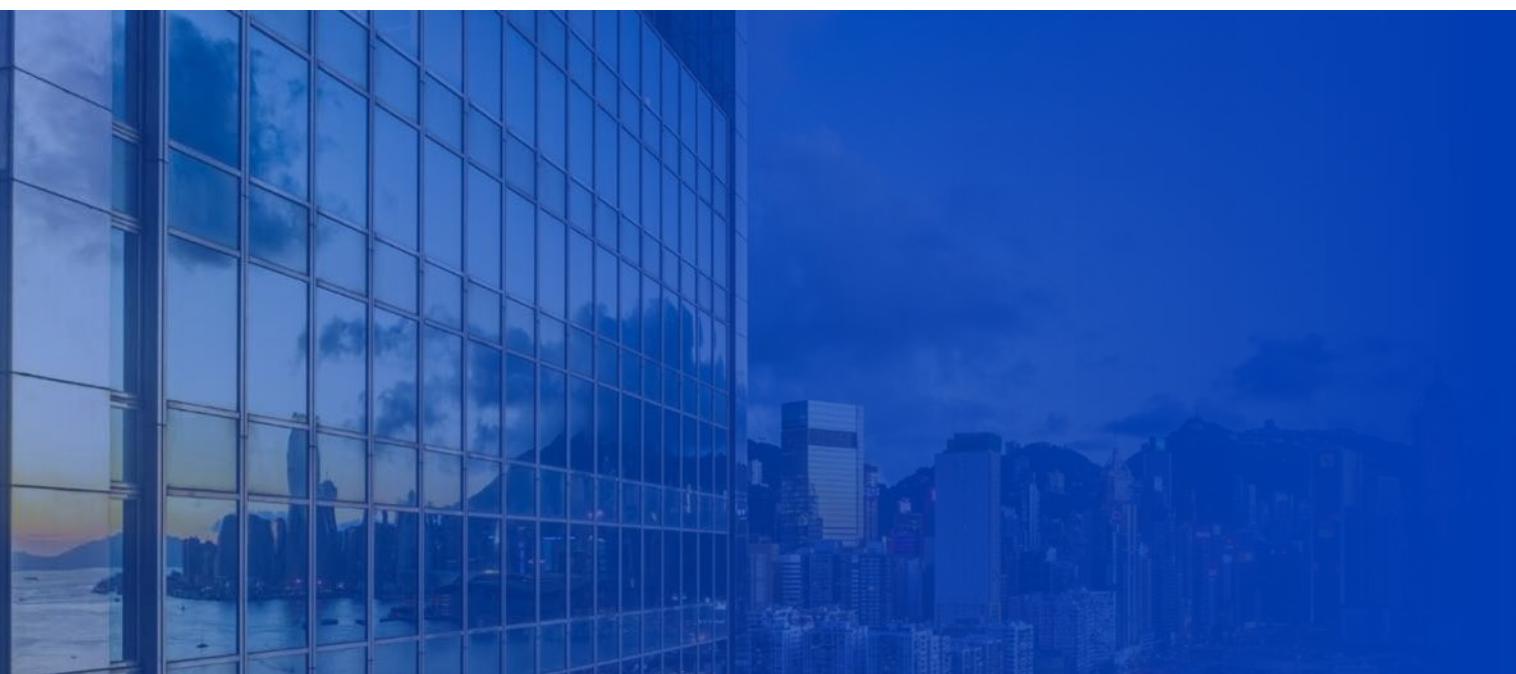
AI's Impact on Wide Area Networking

Demands on enterprise networks
and how service providers can
drive success with innovative
connectivity services

Author: Brian
Washburn
December 2025

In partnership with

HPE JUNIPER
networking



Enterprise AI functionality gains—and growing network needs—are real

Sweep the AI hyperbole aside. What is real is that AI is blowing up. Existing top public AI models receive billions of queries daily. Virtually all large enterprises use AI, and 72% have their own AI deployments. Enterprises see real savings, earn new revenue, and create new value from their AI ventures.

This report covers the real pressure that enterprise AI adoption has begun to put on wide area networks, large enterprise actions and plans, and how service providers are preparing for change. It includes the impact of AI on application requirements, data center needs, and wide area networking (WAN) traffic as a whole (see **Figure 1**).

AI is changing how enterprises connect and communicate. This gives service providers new opportunities to capture business and grow revenue, if they can anticipate and meet these changes.

For this report we conducted international surveys and interviews with global WAN and IT thought leaders. We would like to thank the following executives and organizations for their time and insights:

- Tony Nguyen, executive vice president/CIO, CENIC
- Mirko Voltolini, vice president, Colt
- Jeff Ary, vice president, Lumen
- Jean-Louis Le Roux, executive vice president, Orange
- Paul Ruelas, associate vice president, Verizon

We additionally thank the executives from all 18 interviewed organizations whose valuable contributions made this report possible.

Figure 1: Key wide area networking insights from enterprises implementing AI



Source: Omdia

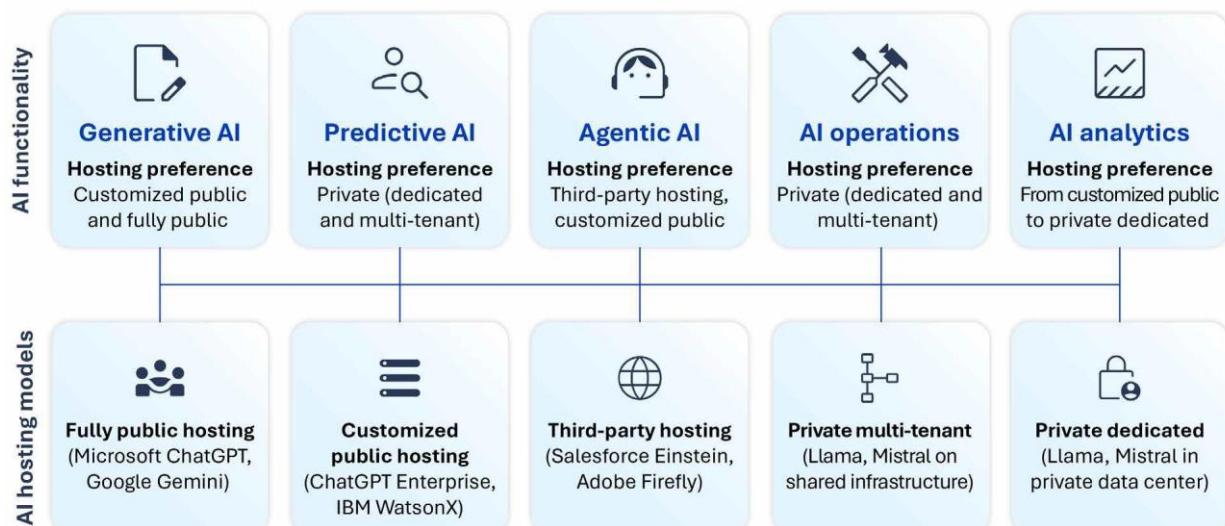
© 2025 Omdia

Enterprises treat AI as a versatile tool kit

AI is versatile. The average large enterprise AI adopter has more than seven different active AI functions. Models can be tuned to emulate human-like interactions and content (generative AI), monitor and forecast trends (predictive AI), manage specific IT systems (AI in operations—AIOps), find and interpret patterns in data (AI analytics), or use tools under supervision (agentic AI). Predictive AI, AIOps, and generative AI represent about 85% of deployed enterprise AI functions. Agentic AI and AI analytics have small shares.

Each AI can be hosted in a range of ways from fully private to fully public. Enterprises mix and match. They choose AI functionality for the task, host with an acceptable level of privacy and control, and set a minimum bar for performance (see **Figure 2**). Enterprises favor private AI models for workforce management, data collection and response, business intelligence, and analytics. They lean into public AI models for a range of enterprise applications (many now hosted as software as a service—SaaS), collaboration, and internet-related activities.

Figure 2: A matrix of AI choices: Functionality and hosting models



Source: Omdia

© 2025 Omdia

Enterprises predict booming demand for network capacity

Enterprises expect the share of AI traffic on their networks to grow by about 5× in the next 36 months. Today, they estimate AI functionality averages about 3% of total network traffic, but they expect that share to balloon to 15% over the next three years.

But that is not all. Enterprises consume a great deal of AI traffic without direct intent. Search engines, upgraded conventional applications and SaaS, media, collaboration, and social media all incorporate AI traffic. Enterprise applications are loading up on AI-enhanced features such as chatbots and virtual support, intelligent selection and filtering, intelligent alert triggers and notifications, and automated management responses.

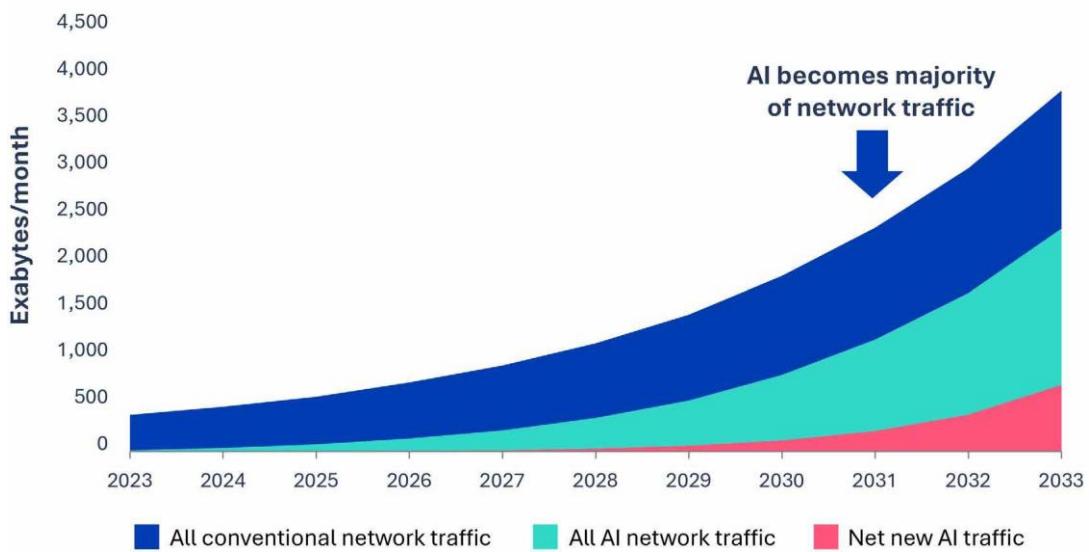
As **Figure 3** illustrates, Omdia estimates that AI traffic from direct investment plus passive consumption is about 14% of all network traffic today and will grow to 31% of all network traffic in the next three years. AI's use of video, image, and future media is a traffic growth engine powered by cognitive analytics, communications and collaboration, and new forms of multimedia content. The big flip is coming: AI is currently projected to eclipse conventional network traffic around 2031.



“We have saved almost \$400,000 worth of time. That’s effort we can dedicate to other projects: We’ve been able to take on 35% more work. AI also benefits our incident response. It used to take seven hours to remediate a network problem in the field. We’re now tracking under five hours. With automation and AI, we have cut two hours from remediation.”

– Electrical utility, North America

Figure 3: Projected global total network traffic, AI and non-AI, 2023–33



New AI functionality drives premium service quality

All large enterprise AI adopters have key functions where speed and reliability are important. Eighty-seven percent of enterprises are willing to pay a premium for their critical AI functions. The need for end-to-end AI performance is pulling enterprises to dedicated optical services and, in some cases, back into MPLS. Organizations need to ensure key AI transaction traffic is delivered in sequence, processed, and responded to quickly (see **Figure 4**).

As enterprises entrust AI with increasingly business-critical tasks, they cannot afford platform downtime, response lags, or failed transactions from delayed or dropped traffic. Enterprises expect the lowest possible latency with AIOps; they need business performance for predictive AI, generative AI, and agentic AI for good user experiences; and they want guaranteed traffic delivery and completed transactions for AI in analytics.

Figure 4: Share of enterprise AI tools that need premium network quality support



Source: Omdia

© 2025 Omdia



“I am thinking about future technical requirements, and the need for AI edge, cameras and VR, and network, storage and compute requirements...For VR environments, we could have AI tracking clinician performance and adapting training models in real time. For vision-driven AI, a single HD camera stream generates about five Megabits per second. This is a future that is coming soon.”

– Healthcare provider, North America

Action items to keep pace with AI needs

Plans to increase network capacity

New AI builds are driving huge bandwidth contracts for network operators. Network providers are fielding requests from major AI players for network capacity with potential to scale to 10 terabits per second (Tbps) and beyond.

These network providers cannot use the old practice of waiting until they receive a service order to ramp up capacity. The AI industry is forcing them to predeploy optical gear and activate new wavelengths in 15–30 days, not three months.

Enterprises also expect an ongoing rise in traffic over the next 12–18 months, on average tripling their bandwidth to data centers and clouds. AI is playing a part in this boost. An AI race that started with North America and China is joined by markets in Europe, the rest of Asia & Oceania, the Middle East, pockets in Latin America, and eventually Africa.

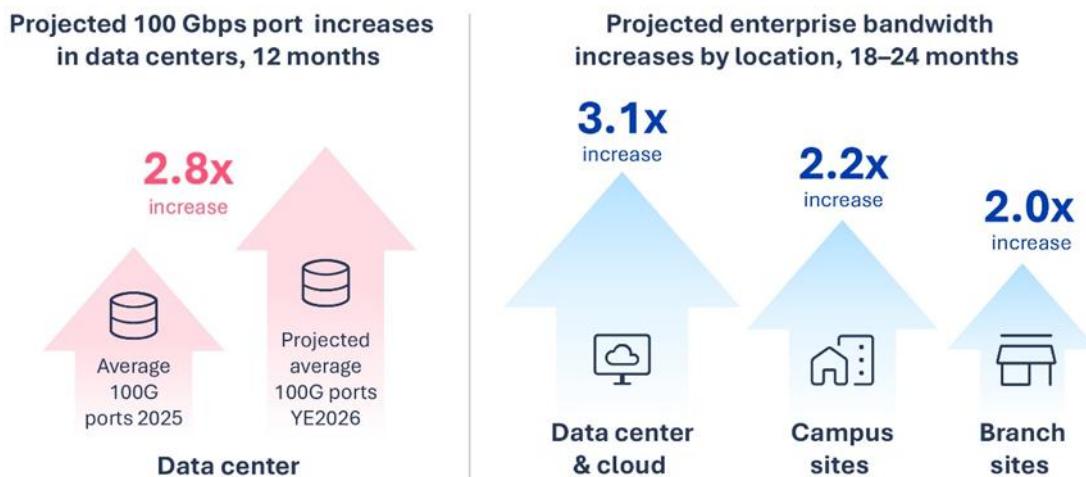
Large enterprises are also bracing for growth. They expect upgrades to 100Gbps ports almost to triple in the next three months, and they are eyeing speeds of 400Gbps and higher. Enterprise campus and branch sites feeding data centers are growing moderately (averaging about 2x bandwidth increase) over the 18–24 months (see **Figure 5**).



“We’ve had a 15% reduction in documentation. We’ve reduced HR clicks for manual data entry tasks. We’re looking to reduce our patient appointment no-shows by about 5%. AI-driven nudges through our patient portal could lead to 25% more loggings. We are looking at ROI in terms of cost avoidance, which translates into dollars saved per year per patient.”

– Healthcare provider, North America

Figure 5: Projected port and traffic growth across data center, campus, and branch

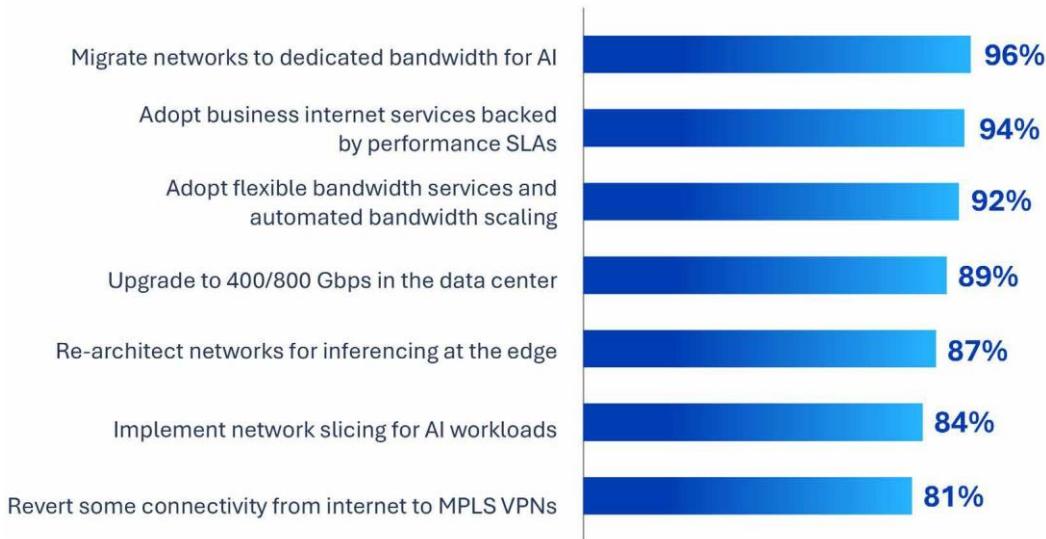


This Omdia White Paper was commissioned by HPE Juniper Networking

Enterprises are executing, or at least evaluating, a range of network capacity and performance options (see **Figure 6**). Dedicated network ports, performance service-level agreements (SLAs), and the flipping of links back to MPLS VPN will address network availability and performance. High-speed 400/800G data center ports and bandwidth for edge inferencing add capacity where it is needed. Optical capacity on demand and flexible bandwidth hedge for unexpected changes.

As enterprises project their near-term (12 months) to long-term (36 months) planning windows, most of these options are in play.

Figure 6: Enterprise AI adopter plans over the next 36 months



Source: Omdia

© 2025 Omdia



“We are using AI with computer vision to analyze manual work tasks for training materials. Tablets and mobile phones capture video, and AI writes up standard operating procedures. We are rolling out this training across all sites. We additionally analyze captured video for production chain optimization.”

– Advanced components manufacturer, Europe

Distributed edge inferencing delivers demands and surprises

Almost all (96%) of large enterprise AI adopters have at least one private or customized model they need to train and inference. Many AI adopters (63%) also have at least one trained model that is mirrored to more than one location. To keep AI models aligned with each other and guarantee they are making good decisions, the great majority of AI adopters (84%) rely on distributed inferencing.

The world’s biggest AI models have an insatiable appetite for training data: books, video and images, audio and music, web content, news, computer code, and constant user interaction. Some models are approaching 100TB. Feeding these models, and keeping them synchronized, is one reason why AI platform operators need massive bandwidth.

Enterprises' trained private models and customized public models are much smaller, but the differences are critical to the business. The average trained or customized enterprise model takes hundreds of gigabytes to set up. The average enterprise inferencing data stream is measured in gigabytes per week. But there is also a lot of activity. More than 60% of private and customized AI models operate multiple instances that need to be kept synchronized, and 88% of enterprises must retrain their models regularly—typically every 6–12 months—to keep models from drifting and to guarantee reliable results.

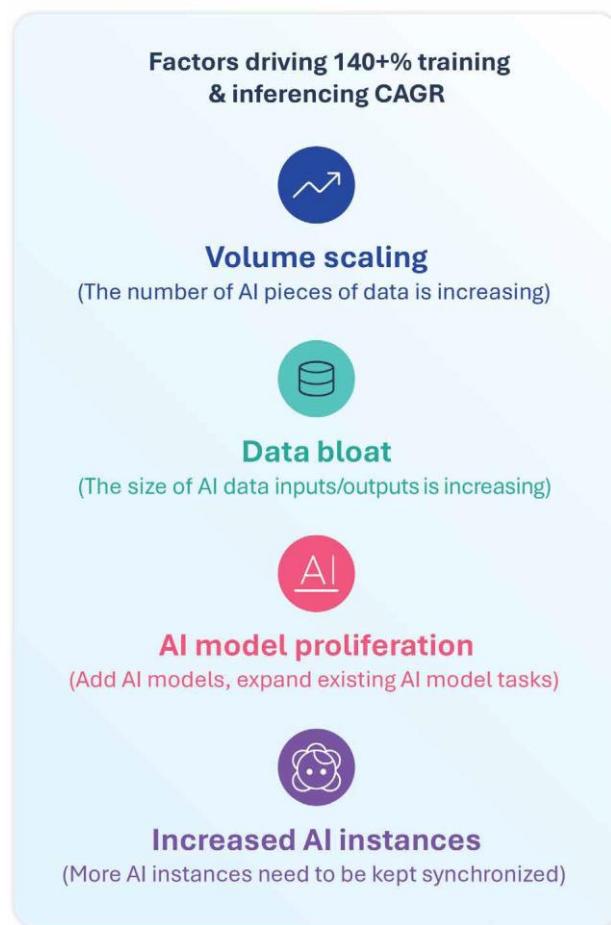
The net effect is that administrative data for AI traffic is blowing up at a CAGR estimated at more than 140% (see **Figure 7**). AI training and inferencing is still just a sliver of total AI traffic, which is a slice of all network traffic. But it is a critical slice, and most enterprises that use inferencing have ordered bigger ports (1.5–2×) for inferencing. Nearly half of enterprises also demand low latency (<100ms) for inferencing data transfers to keep models synchronized.

What happens next: Distribution and alignment

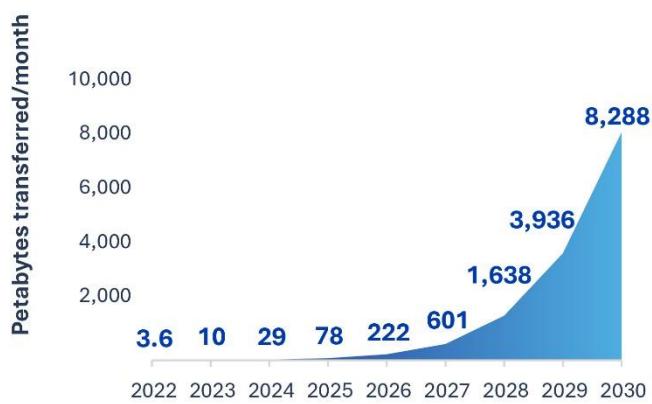
About half the time, enterprises train the AI model and shift to commercial operation in the same place. But it is just as common for enterprises to train AI models in one place, then transfer or distribute the model somewhere else to do its job. They may move trained models from centralized cloud to a private data center or from on-premises to a hosted edge. Today, most customized and private AI workloads (60%) are handled in centralized clouds and hosted edge. About 20% of workloads reside onsite, and about the same percentage are hosted in a private data center.

The concept of “edge” is splintering. Market realities push AI to decentralize and specialize (see **Figure 8**). Edge conventionally implies proximity to the customer for the best performance, but there is also a resources edge: proximity to plentiful power, efficient cooling, and inexpensive real estate. There is also a geopolitical edge: reaping the benefits of sovereign/commercial compliance and government financial incentives.

Figure 7: Enterprise AI inferencing traffic and its growth drivers



Explosive growth is forecast for training & inferencing of net new AI traffic



Source: Omdia

© 2025 Omdia

Figure 8: Many drivers are forcing AI data centers to become distributed



Source: Omdia

© 2025 Omdia

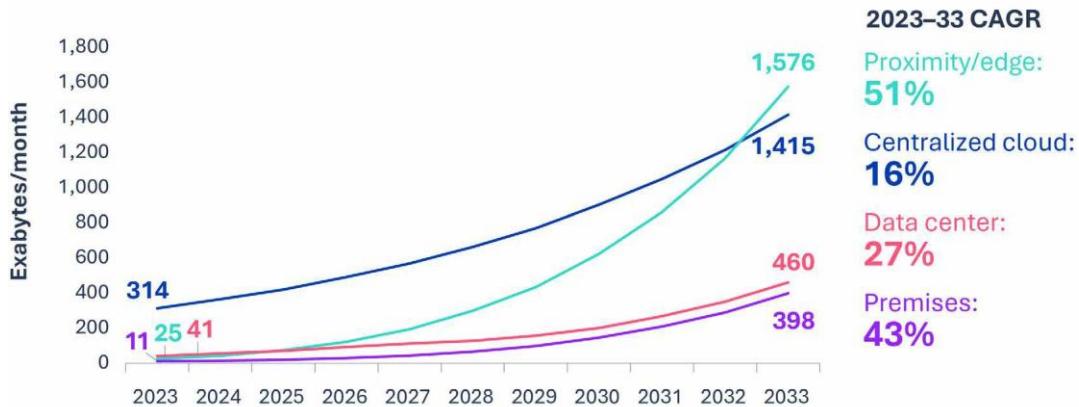
Over time, Omdia expects AI traffic growth momentum to sit squarely with a broad definition of edge. New sites will be distributed, exploiting some combination of customer proximity, resource proximity, and geopolitical proximity. That leaves centralized cloud to refer to the world's massive hubs for compute processing (these also continue to grow traffic). Premises-based AI is also plentiful but, as **Figure 9** shows, it generates less WAN traffic than other hosting targets.



“We are building regional data centers to distribute AI loads. We are targeting a flat architecture with low latency, where traffic is routed to the closest data center. We are also distributing our landing points for subsea cables.”

– Leading national service provider

Figure 9: Projected total network traffic by processing destination, 2023–33



Source: Omdia

© 2025 Omdia

This leaves service providers with two fundamental future challenges. The first is a change in balance for upstream versus downstream traffic. New AI traffic growth is more balanced than historical downstream video growth. AI traffic needs more upstream traffic for processing and recognition from video and images. Augmented reality also needs real-time object recognition and analysis for its interactive experiences.

The other challenge is intelligent orchestration to unite network and compute resources. This orchestration needs to be smart enough to arbitrate competing interests for resources end to end. It must make real-time decisions to assess, capture, and release combinations of resources that best meet each competing interest's policies and service levels.



“Operator networks are becoming hybrid to reduce latency and improve caching. There are all sorts of technical reasons why to distribute core networks. In Europe, it’s now normal for an operator to have 20 data centers in one country.”

— Leading regional service provider

Conclusions and recommendations

Enterprises: Choose your mission and build

Every enterprise takes its own journey in the B2B AI revolution, motivated by its investment priorities. Some industries and business cultures lean first into customer engagement; others on automating their operations, or on securing their assets. Still other organizations just want to save money. Your organization’s goals will guide initial AI focus: the IT department, sales and support teams, or finance and executive administration.

- **Start with easy wins:** Address issues that drive your C-suite decision makers; choose areas where industry peers use AI; let your existing partners pitch their AI game. Intelligent transcription, workforce scheduling, site surveillance, and in-house automated IT support are examples of horizontal and vertical easy wins.
- **Gems in, gems out:** If you plan to train or customize AI, your peers warn that it is critical to align and clean up internal data sources first.
- **Process not product:** AI grows, adapts, changes. Check AI output for accuracy, and recalibrate models regularly to be sure. AI that only infers its own conclusions can eventually drift into mistakes.
- **Security minded:** Your peers caution to be mindful of AI hosting models for privacy and security. Curate and analyze your AI data sources for good measure. Even fully private AI models can be compromised by a malicious injection of bad information.
- **Network minded:** Key AI functions have critical performance needs. Enterprises are adopting 100Gbps network ports and have 400G/800G in sight. They turn back to MPLS for its performance and add capacity in data centers for distributed inferencing. A robustly built WAN should never be the cause of AI quality issues.
- **Success feeds success:** Winning AI adoptions start with the business case; match the need to the best-suited hosting model, and closely track results. Over time they accelerate off proven AI use cases and built-up expertise. If your business lags in AI, you will need to force a concerted catch-up effort or else fall ever further behind.

Network providers: Move to a proactive mindset

Network providers are split on the way they see the market. A frenzy of AI investment started in North America and parts of Asia; it is rising in Europe and elsewhere. Long-term cloud and neocloud network contracts are potentially worth billions of dollars. But enterprise and consumer AI use rides with existing network traffic. Some network providers are adding fiber and optical gear now, seeing the changes. Future large-capacity orders will want fast delivery with delivery timeframes of 15–30 days, not 60-, 90-, or 120-day service orders.

- **Performance matters:** AI functions need availability, latency, and delivery guarantees that match the mission. Some enterprises said they buy extra capacity just to improve latency and transactions. Providers need to present capacity and performance guarantees as options for enterprises to make sure the network is not at fault for AI function failures.
- **Brace for broadband rebalance:** Wireless, cable, and DSL broadband were tuned for downstream, passive, mass-market media consumption. Much AI functionality is interactive. Providers will need to support balanced upstream/downstream infrastructure and service plans.
- **Enterprises need AI help:** Enterprises implementing AI need functional models, hosting and connectivity options, training content, sensors / controllers / field devices, connectivity between platforms, AI implementation expertise, and expertise in their industry vertical. Providers may play a role here, by pulling together partners and by supplying direct professional services expertise.
- **Turn that help inside out:** Telecom service providers are early and complex AI adopters. They can flip internal ramped-up expertise—and even host instances of trained/customized platforms—to enterprise customers that need help deploying AI.



“Industries like manufacturing spend a lot on data that comes from cameras. Industries like retail have cameras, and inputs from tablets and phones. Image classification has become very common in healthcare and in manufacturing.”

— Leading global hyperscaler

A future that is already happening

Drone swarm light shows, personalized gene editing, algorithm-managed investing, virtual workers monitoring people and virtual workers managed by people, and a virtual actress shopping talent agencies: these are already happening.

Where AI at times turns the world upside-down, enterprise AI often starts with simple and straightforward tasks. Examples include internal information that is searchable with natural-language queries, automatic transcription and summarizing of meeting minutes, and alerts triggered by anomalies with incoming data.

AI is already accelerating the next phase of simple and straightforward solutions: monitoring through cameras, sensors, and systems data. These success stories also start with simple tasks. If inventory is low, issue a reorder. If there is a manufacturing defect, set the unit aside. When traffic conditions change, reoptimize routes.

Not all of what is coming will be simple. AI virtual workers exist today. Futurists expect an explosion in agentic AI that organizes into clusters to collect data, make decisions, and execute them. These agentic workloads will need extremely strong performance to do their job at a human (or faster) pace. There is potential for these agents to scale into the millions.

Currently, the single largest limit to AI is the ability within enterprises for people to change and to adapt to change. For both enterprises and service providers, leaders will reap benefits of efficiency and creating new value; laggards will struggle to keep pace with change.

Appendix

Survey/executive interview demographics

North America

Enterprise interviews:

-  Electrical utility
-  Financial services
-  Healthcare provider
-  Research network

Service provider interviews:

-  Tier 1 global networks
-  Regional networks
-  Hyperscalers
-  Interconnect specialists

18

Enterprise & provider executive interviews

314

Large enterprise survey respondents

UK	18%
Canada	6%
USA	43%
Germany	16%
France	17%

Europe & Asia

Enterprise interviews:

-  Government agency
-  Manufacturing
-  Research university



Service provider interviews:

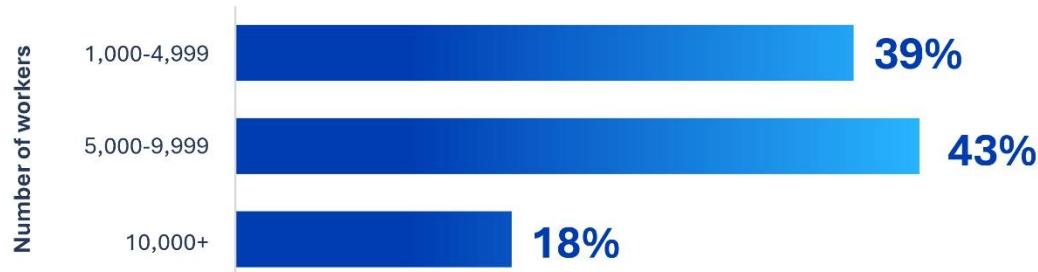
-  Tier 1 global networks
-  Regional networks
-  Wireless specialists

Service provider interviews:

-  Tier 1 national networks



Enterprises surveyed by size





Omdia consulting

Omdia is a market-leading data, research, and consulting business focused on helping digital service providers, technology companies, and enterprise decision makers thrive in the connected digital economy. Through our global base of analysts, we offer expert analysis and strategic insight across the IT, telecoms, and media industries.

We create business advantage for our customers by providing actionable insight to support business planning, product development, and go-to-market initiatives.

Our unique combination of authoritative data, market analysis, and vertical industry expertise is designed to empower decision-making, helping our clients profit from new technologies and capitalize on evolving business models.

Omdia is part of Informa TechTarget, a B2B information services business serving the technology, media, and telecoms sector. The Informa group is listed on the London Stock Exchange.

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Omdia's consulting team may be able to help your company identify future trends and opportunities.

Get in touch

www.omdia.com
askanalyst@omdia.com



Copyright notice and disclaimer

The Omdia research, data, and information referenced herein (the "Omdia Materials") are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together "Informa TechTarget") or its third-party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document.

The information and opinions expressed in the Omdia Materials are subject to change without notice, and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an "as-is" and "as-available" basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third-party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.