Omdia
by informa techtarget

Let's Get Physical: Why the Network Is Not One-Size-Fits-All

Five reasons to rethink connectivity in an AI-first world

January 2026

In partnership with:

Telstra

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

Contents

Introduction 3

What's driving network traffic? 4

Why physical infrastructure matters 8

Five reasons to rethink connectivity 9

How industry needs vary 12

Questions to ask 15

Conclusions 16

Appendix 17

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

Introduction

It's time to look at AI the right way up.

Beyond graphics processing units (GPUs) and data centres, AI's ability to reshape, empower, and disrupt industries depends on a digital asset that few consider: the network.

Eight out of 10 enterprises already run at least one type of AI application, and most are accessing a public AI model. This requires connectivity—and not just any connectivity will do. In an AI-first world of real-time decisions and on-demand experiences, securing network performance has never been more important.

Provenance matters too. Data sovereignty concerns are escalating globally. To ensure their resilience and accountability, enterprises must know how an AI workload routes across land and sea, not just where a dataset is processed.

The physical network underlay has become a mission-critical issue.

Do you know what yours can do for you?

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

What's driving network traffic?

Growth is the baseline

It's no surprise that enterprise bandwidth keeps growing. From low-earth orbit satellite broadband to 400G backbones, networking options are multiplying. They make it possible to connect more people, places, and things across more diverse environments.

The day-to-day reality of managing these digital interactions keeps enterprise networking teams busy:

- **Driving agility:** Enterprises are using network virtualisation to manage public and private transport. They must fine-tune the cost and performance trade-offs of the public Internet versus private connectivity choices.
- **Orchestrating clouds:** Enterprises are managing multi-cloud interconnects, direct connections, and VPNs. They must streamline cloud sprawl and optimise multi-cloud performance.
- **Securing assets:** Enterprises are using cloud-based security services, such as Secure Service Edge (SSE), to keep safe. They must calibrate their network performance to maintain critical application traffic flows.

AI traffic: It's more than you think

AI is impacting digital consumption and traffic patterns. AI brings more complexity to today's networking reality, and it's often hidden in plain sight.

- **Workhorse applications are becoming AI-enhanced:** Consider Microsoft 365, Salesforce, Oracle Fusion, HubSpot, and social media apps. Usually hosted in public clouds that enterprises must access, these apps are also now enhanced with AI features. AI is summarising meetings, suggesting next best actions, detecting trends, and creating content in real-time or near real-time.
- **A hidden AI traffic shift is underway:** Today, mostly without enterprises realising it, a third of global network traffic has become AI-enhanced. Over the next few years, B2B interactions will dominate backbone traffic. But an important consequence of retrofitting AI features into common business applications is the need for greater network reliability, lower latency, and higher speed.

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

Omdia commissioned research, sponsored by Telstra International

Figure 1: AI's impact on global backbone traffic

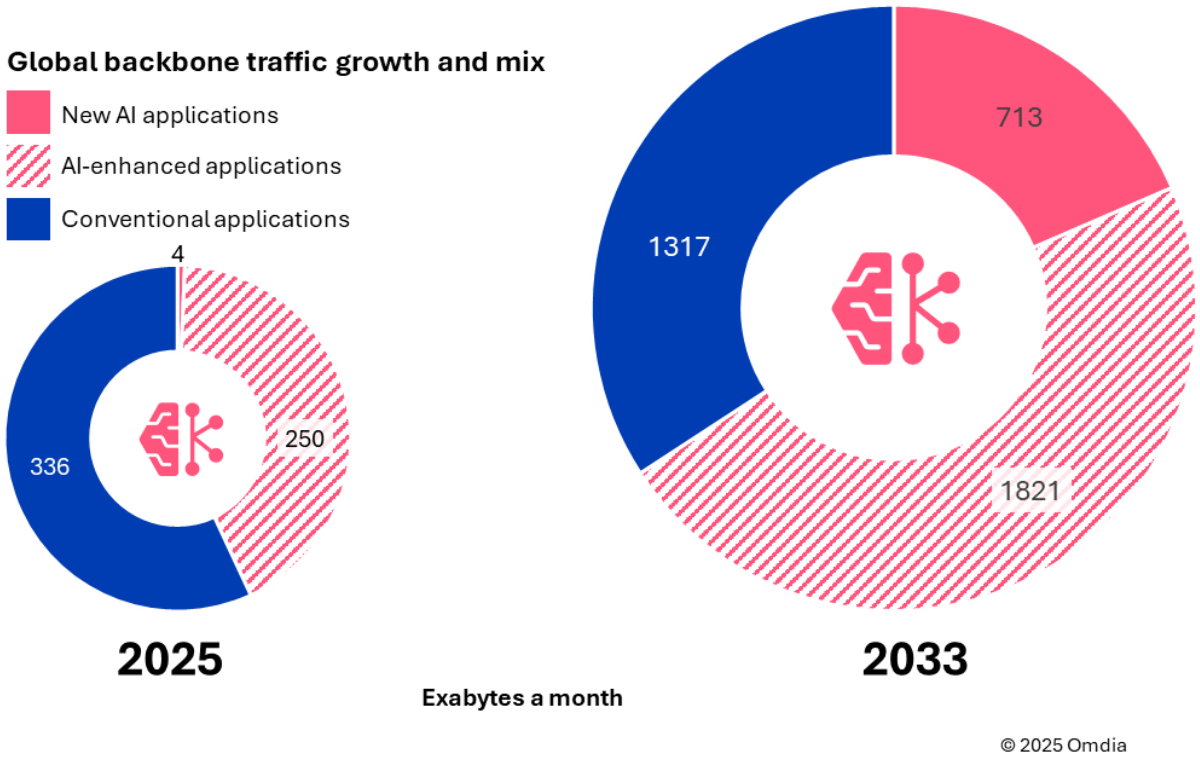

Source: Omdia

- **Each AI model creates a different network traffic profile:** For example, Generative AI (GenAI) causes sporadic traffic spikes when someone makes a query. But Agentic AI encompasses autonomous AI systems capable of multiple independent actions. This exerts even more network performance demands. Agentic AI creates sustained traffic flows beyond human interaction speed. Soon, machines will talk on networks more than humans. Forget about planning your enterprise network architecture around a predictable human-centric traffic busy hours: some AI applications will always be busy and run micro batches 24/7.

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

Figure 2: Comparing AI application latency

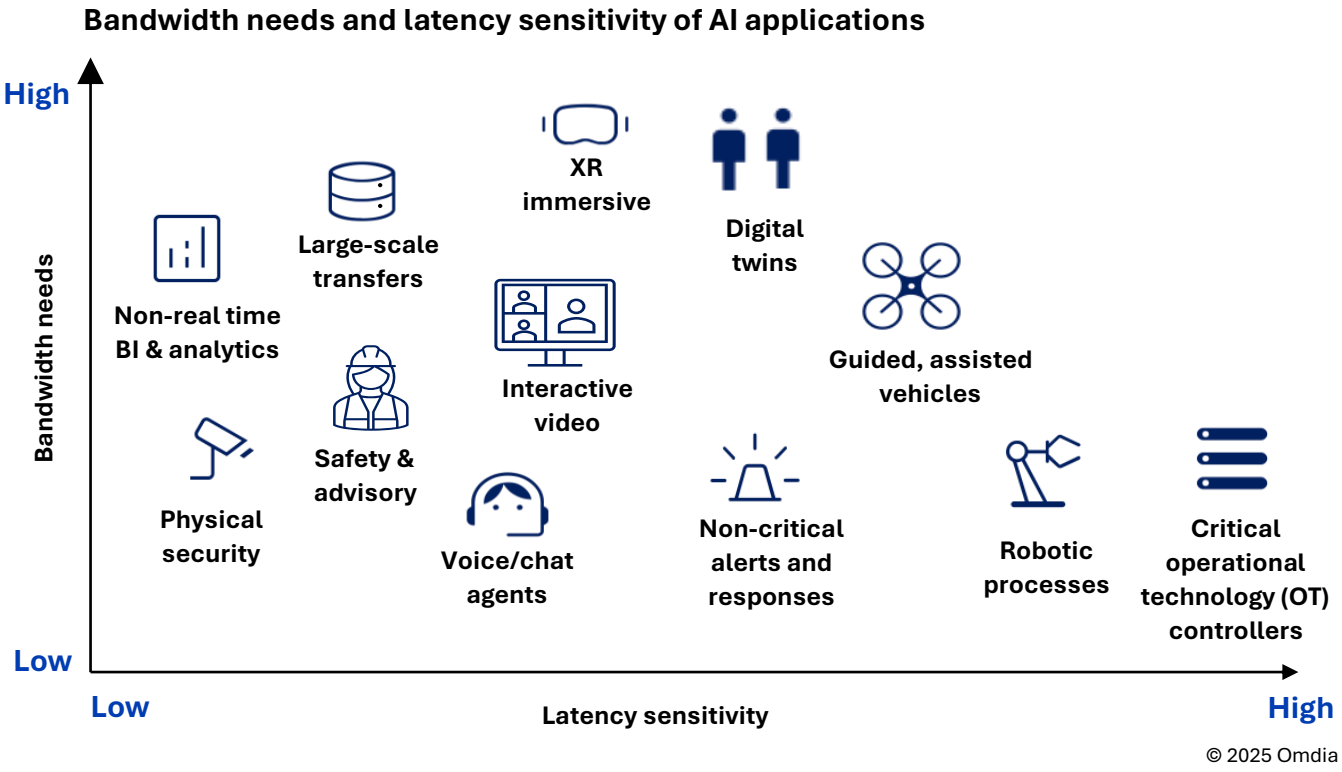

Source: Omdia

Superb network performance is not optional

Many enterprises are already experiencing quality issues with existing digital infrastructure. Greater reliability is now the number one priority driving enterprise networking decisions, according to a 2025 Omdia survey of 419 large global enterprises. Next on the wish list: higher performance and increased bandwidth. Growing AI use across enterprises will only escalate these issues, so the time to act is now.

- **Appetite for cloud consumption keeps rising:** Eight out of 10 enterprises expect to double their cloud connectivity in the next 18 months, with dramatic rises in 100G connectivity planned.
- **Cheap public Internet is getting riskier:** 60% of enterprises plan to rebalance from the public Internet to business-grade options. Internet performance can be uncertain in some regions of the world because local networks are optimised to favour residential, not business customers. Rock solid performance is required for some digital workloads, drawing enterprises to seek dedicated connectivity and business-grade internet services.

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

- **AI traffic is the coming challenge:** While much AI use is experimental and immature, many enterprises are starting to consider network impact. Over two-thirds of current AI users expect to quadruple their AI bandwidth within 18 months (see fig 3).

Enterprises already manipulate multiple AI models. Some AI capabilities are well known in some departments, like AIOps, widely used in IT and networking environments. Others, like Agentic AI, have capabilities that multiple departments could eventually benefit from. Most AI models, whether experimental or fully deployed, are hosted publicly (see fig. 3). This is driving enterprise expectation that their AI traffic will grow fourfold.

Figure 3: Enterprise AI use—diverse and often publicly hosted

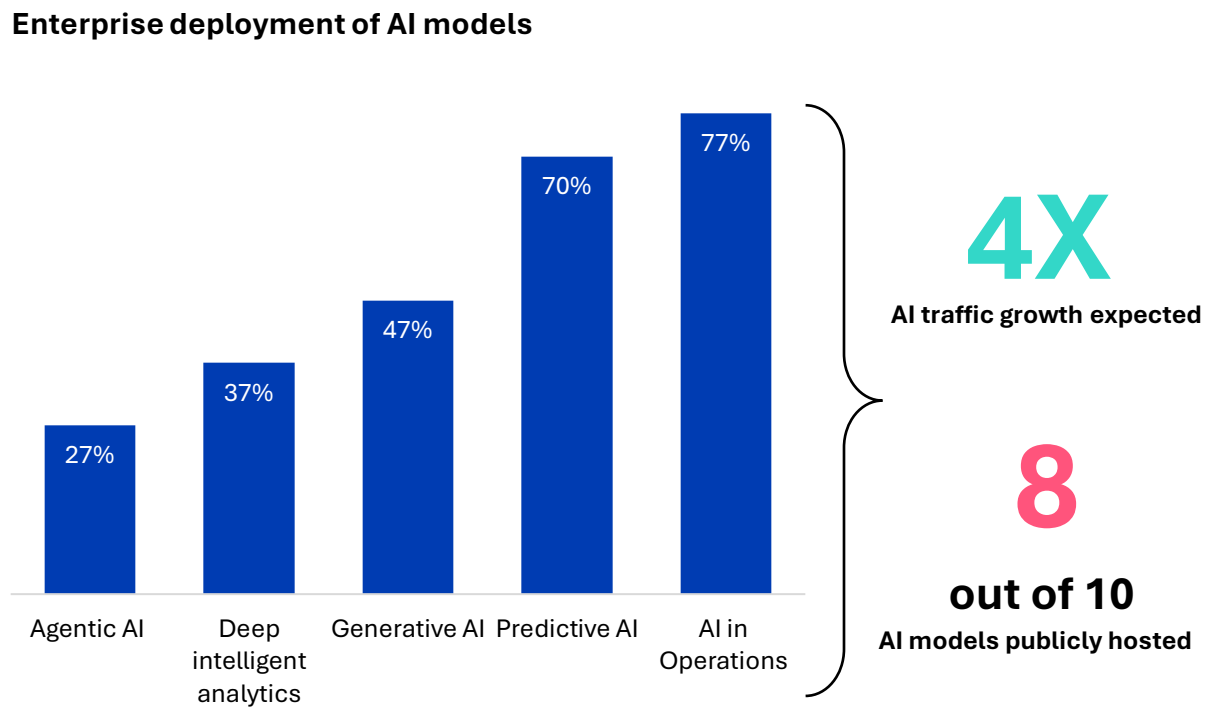

Source: Omdia

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

Why physical infrastructure matters

To thrive and stay safe in an increasingly AI-first world, enterprises must be aware of various networking realities:

- **The public Internet is a best-efforts environment:** Routes are picked based on policy, which might be what's cheapest, not what's best. Congestion control is reactive; routers forward packets until they're full. The best way to mitigate performance issues is to work with a partner that is a Tier 1 provider in key regional geographies: this status means shorter paths, richer peering options, and single-hop direct peering with other service providers, hyperscalers, CDNs, and traffic exchanges.

- **Resilience must consider physical infrastructure:** Underlying physical infrastructure from the fibre right of way, operation of a cable landing station, or the availability of power redundancy. Severing a submarine fibre cable could be due to an accidentally dropped anchor or a bad actor. Enterprises need to be informed about such risks and be prepared to deal with them.

- **AI is taxing on networks:** Training an AI model needs major computing power to ingest huge data volumes. If that data exists in multiple physical locations, then it also involves serious bandwidth for bulk data transfers to compute. This activity has little tolerance for jitter, meaning training stalls if packets take a sub-optimal path. That's expensive if a GPU cluster idles due to network congestion. Running an AI model can impact network traffic as it depends on live data flow. Latency is critical, and this is driving demand for more local edge inference nodes. More traffic is driven into hyperscale AI endpoints, increasing the need for high-capacity links.

- **Data sovereignty applies to routing, not just processing location:** Data sovereignty is about exerting control over digital assets without undue external influence or dependency. But the need for sovereignty applies not only when data is stored or processed. Sovereignty also matters the moment data moves. Routing policy determines where data moves, and on the public Internet, this can be difficult to control. Nevertheless, emerging regulatory activity is focusing on routing provenance and origin validation.

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

Five reasons to rethink connectivity

Reason 1: Complete security demands physical and virtual control

What	Logical precautions such as encryption and authentication are highly valuable within a managed SSE platform. But they can only go so far to protect enterprises. The physical reality of networks can't be ignored: all data passes through tangible assets like a subsea cable or a data centre.
Implication	Choice of network 'underlay' demands deeper scrutiny. This refers to the underlying physical network infrastructure, such as routers, switches, and transmission kit. It is also about land and subsea fibre routes. Not least, network underlay encompasses the global Internet backbone infrastructure where traffic is exchanged. Resilience depends on underlay redundancy on multiple levels: equipment, connectivity, and traffic routes.
Action	Ensure secure, high-performance connectivity to your SSE platform. Choose higher-performance, lower-latency network paths that strengthen your security posture.

Reason 2: AI network traffic patterns are less predictable

What	For AI to thrive, best-effort Internet doesn't cut it. AI workloads accumulate various inputs for batch aggregation. This results in synchronised surges to GPUs, creating 10–100x throughput spikes compared to web traffic. Physical fibre cables, peering points, and routing logic matter when conveying AI workloads safely.
Implication	Deeper and more diverse Internet interconnection offers more choices to maintain the required AI application performance latency. Major Internet backbone owners offer better routing control (and redundancy) with single-hop infrastructure.
Action	Favour service providers with strong hyperscaler adjacency in all the geographic regions that matter to the business. Expect multiple direct cloud interconnections and edge PoPs in addition to extensive local exchange participation. Check access arrangements to SSE providers. Are routing policies multi-path and low latency? Activities like payment authentication or autonomous navigation need a strong network to maintain a good user experience.

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

Reason 2: The public Internet was not designed for AI

What	Real-time inference, distributed training with edge to cloud orchestration are key AI processes. They depend on a network fabric that is predictable, secure, and sovereign. The public Internet, designed for best-effort delivery, is at odds with the compliance and determinism that AI processes demand.
Implication	Overcoming Internet performance challenges requires a two-pronged approach: the quality of physical network underlay coupled with an identity-aware, policy-driven routing layer to steer traffic.
Action	Look for service providers that can help you apply application classification rules to traffic routing to secure the best performance. Can they send high-sensitivity inference traffic to direct cloud on ramps while other applications use the public Internet? Deterministic, controlled flows to the cloud avoid common Internet hiccups such as hairpinning traffic via latency-inducing junk routes.

Reason 3: Data sovereignty requires a network audit

What	Sovereignty brings complex new compliance responsibilities to enterprises, and they vary by country. Data localisation and sovereignty rules may also introduce cross-region networking traffic asymmetries. App and data may stay in country, but identity and access control may sit in another region. Sovereignty may restrict or prevent cloud data replication, archiving, or disaster recovery to foreign regions.
Implication	Greater scrutiny of network underlay can determine if digital services running over a supplier's Internet backbone can adapt to these new needs. Sovereignty requirements may make routing transparency mandatory in several geographies. Sovereignty will increasingly drive choices in underlay topology, potentially with local first routing and edge peering.

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

Action	AS path transparency is important. This refers to entities running a specific network. Check what a service provider can tell you about transit provenance—where packets traverse different providers and jurisdictions. Can they tell you whether a route stayed within a permitted jurisdiction, and alert if not?

Reason 5: Resilience has a wider definition in a multi-cloud world

What	Resilience used to mean redundancy and diversity, but that's becoming a dangerously narrow definition. Enterprises are distributing compute across hyperscalers, edge zones, and sovereign regions. The attack surface is wider due to remote work and multi-cloud SaaS use across unmanaged endpoints.
Implication	Securing how data moves between clouds needs greater focus. Resilience is becoming multi-faceted, factoring type of workload, required data paths to ensure performance, and not just about securing uptime.
Action	Treat cloud on-ramps as critical infrastructure. Evaluate which cloud regions a service provider connects to, available capacity, and whether diverse physical paths to cloud on ramps exist. Check on their historical Border Gateway Protocol (BGP) stability and Mean Time To Resolve (MTTR).

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

How industry needs vary

Drivers behind digital consumption

Enterprises may be consuming more digitally, but the reasons why vary by industry. One business priority is front and centre for all: the need to secure physical and virtual assets (see fig. 4). Too much is at stake: processes, devices, and equipment are all now network dependent, in addition to customer and partner interactions and financial transactions.

Figure 4: Key reasons to invest in better connectivity

Top business priorities driving network investments

	Financial services	Energy & mining	Manufacturing	Transport & logistics
Securing physical & virtual assets	√	√	√	√
Improving customer experience		√	√	√
Automating processes / boosting productivity			√	√
Improving product / service quality	√	√		
Making data-driven decisions	√			

N=419

© 2025 Omdia

Source: Omdia

Other common trends include:

- **Moderating reliance on the public Internet:** Eight out of 10 enterprises plan to rebalance their public Internet and private network investments. Some moved too sharply away from MPLS and are re-examining other private connectivity choices. Some were never attracted to Internet options and are now finding them useful in the mix. Calibrating the right balance for the business, application, security, and performance needs in the applicable region is the name of the game today.
- **A big appetite for cloud connectivity:** The norm among enterprises is to double cloud bandwidth within the next 18 months. However, many enterprises expect to do much more. For some, this is due to historical under-investment. However, the rise of AI models running in public clouds is also causing a spike (see fig. 5).

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

Figure 5: Some industries expect to quadruple cloud connectivity

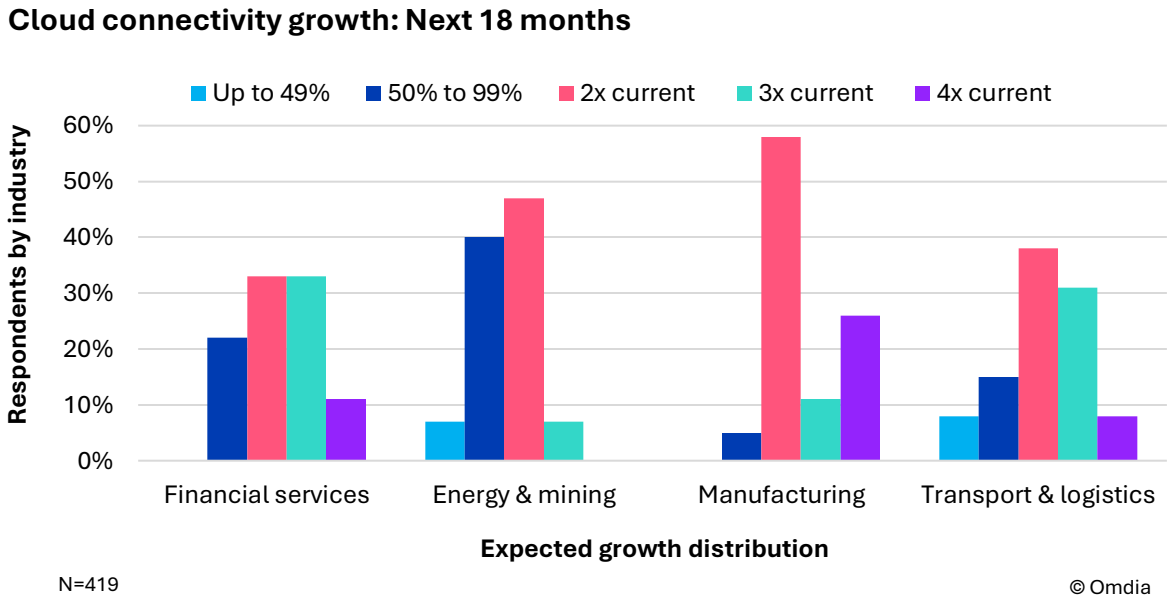

Source: Omdia

Planning for the next 18 months: An industry view

Financial services: Smooth consumers

- **AI use:** 77% have at least one model deployed
- **Key networking issues:** Greater reliability, but also more flexibility
- **Cloud/data centre bandwidth:** Double is the norm, 40% expect to triple or more
- **Branch bandwidth:** 70% expect capacity to grow at least 100%
- **Most critical bandwidth-hungry apps:** Payments, high-frequency trading, ERP

Takeaway: Already major bandwidth consumers, financial services firms have developed a nuanced understanding of ebbs and flows in their digital consumption. They spend proportionally more on networking than other industries. This is why financial services firms need to control consumption needs more accurately. Scrutiny of on-demand provisioning and management tools is increasingly important in supplier evaluation.

Energy & mining: Thirsting for knowledge

- **AI use:** 86% have at least one model deployed
- **Key networking issues:** Greater reliability and higher performance
- **Cloud/data centre bandwidth:** 60% expect to need at least 10x 100G, and to double their current bandwidth.
- **Branch bandwidth:** 50% growth is expected.

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

- **Most critical bandwidth-hungry apps:** Geological/seismic exploration, asset management, IoT data collection

Takeaway: The power of data collected from a wider range of distantly connected assets is exciting for energy & mining firms. With this, they can increase yields, anticipate demand, and improve safety. Running more remote operations, machinery, and autonomous processes amps the need for network reliability. Diverse connectivity is required across various terrains, adding complexity. They need partners that can make the supply and support of hybrid connectivity easy.

Manufacturing: The quality tyrants

- **AI use:** 85% have at least one model deployed
- **Key networking needs:** Higher performance and more bandwidth
- **Cloud/data centre bandwidth:** 60% expect bandwidth to double, a third to triple
- **Branch bandwidth:** 60% expect growth of 100% or more
- **Most critical bandwidth-hungry apps:** OT controllers, quality control, SCM/ERP

Takeaway: Extraordinary growth in capacity needs cannot compromise core principles. Manufacturing firms need predictable latency, jitter, and packet loss because many processes depend on consistent network timing. Despite working across multiple production sites, partners, and jurisdictions, they must strive for global consistency, particularly around contractual service levels.

Transport & logistics: Real-time realists

- **AI use:** 68% have at least one model deployed
- **Key networking issues:** Greater reliability and higher performance
- **Cloud/data centre bandwidth:** Most will double their current bandwidth, and over a third expect to triple it.
- **Branch bandwidth:** At least 50% growth is expected.
- **Most critical bandwidth-hungry apps:** Route planning, SCM/ERP, transport asset inventory management

Takeaway: Time is money in transport and logistics. Companies must work with service providers that can provide an array of options—fixed, mobile, and non-terrestrial—to keep the business and customers informed about asset location, whether human, physical, or virtual. Real-time visibility and telemetry are core to a business, as is demonstrating a chain of custody for goods. Continuity of information flow is non-negotiable: they should also pay particular attention to backup networking policies to ensure this.

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

Questions to ask

Questions about you

An emerging AI-first world demands certainty and supplier transparency. This applies within an enterprise's operations too. Do you know what's running over your digital infrastructure?

Table 1: Questions to ask internally

Topic area	Questions
Applications	Do we understand the network performance requirements of our mission-critical applications?
Network	Can our network infrastructure flex to cope with less predictable resource consumption needs, particularly as AI workloads grow?
Routing	Does our current routing infrastructure support business-grade latency?
Security & sovereignty	What are our sovereign data handling requirements across our global operations? If we have deployed SSE, have we designed our network to optimise user traffic through the SSE POPs?

Questions about suppliers

Do you know what you are paying for and the risks you are exposed to?

Table 2: Questions to ask externally

Topic area	Questions
Applications	Can the supplier optimise the performance of priority applications by directly interconnecting with SaaS providers?
Network	Can the supplier provide diverse and resilient network options in target geographies for different working contexts, from fixed, mobile to ad hoc locations in urban, rural, or challenging terrains?
Routing	Does the supplier have strong peering relationships? Can they share any statistics about their internet backbone volumes and direct peering arrangements with other carriers and content providers?
Security & sovereignty	Where do packets go? Can they provide optimised routes to security cloud services such as Secure Service Edge (SSE)?

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.

Conclusions

An AI-first world still depends on physical infrastructure—on land, under sea, and, increasingly, in space too. To thrive in the AI super cycle, enterprises must secure digital performance: predictable latency, verifiable data routes, and trusted compute zones.

Importantly, they must understand data sovereignty is not just about where data resides, but also where it transits internationally. Trusting in logical controls to ensure security, accountability, and performance is only part of the solution. Enterprises need to set a higher bar and expect the same from connectivity suppliers.

- **Relearn the Internet:** The Internet is best efforts, but your business can't afford to be. Understand how your service provider's peering status impacts routing and latency. Seek single-hop routing for better performance. But consider the workload too, because a direct private connection may be a better choice.

- **Curate your underlay:** Security starts beneath the application. Zero-trust overlays alone can't fix an untrusted path. The security of the physical network is the first layer of defence. Your service provider must tell you how it can mitigate against issues such as cable cuts or power outages without compromising security or performance.

- **Treat routing as compliance:** Cloud compliance once meant proving data at rest was local. Expect routing attestation to become a line item in risk and compliance audits. Your service provider should clarify where, how, and with whom it interconnects, ideally offering choices in case of sovereignty issues.

Appendix

Methodology

Between September and October 2025 Omdia surveyed 419 enterprise technology decision makers in large multi-site enterprises about their current digital environment, networking investment plans, applications, and AI usage. The respondents were drawn from 15 countries, with representative industry samples in construction, healthcare, energy and utilities, financial services, professional services, retail, and transport organisations.

Camille Mendler, Research Director, Omdia
askananalyst@omdia.com

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 capitalise 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
askananalyst@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.

© 2025 TechTarget, Inc. All rights reserved. Unauthorised reproduction prohibited.