

The AI Network Readiness Challenge

Practical Steps for Ensuring the
Network Will Support Strategic AI
Transformations

Author: Jim Frey, Principal Analyst

JULY 2026

in partnership with **verizon**

This Omdia White Paper was commissioned by Verizon, Inc. and is distributed under license from TechTarget, Inc.

Contents

Executive summary	3
The drive for AI network readiness	3
What AI network readiness means	4
Identifying network readiness for AI	5
Explorer	6
Pilot	6
Deployer	7
Optimizer	8
The agentic wildcard	9
The Verizon approach to AI network readiness: Verizon AI Connect.....	10
Conclusion	13



Executive summary

Failure is not an option when it comes to AI initiatives, and organizations must plan accordingly, including assessing readiness of networking as an essential infrastructure pillar. This comes down to understanding and aligning networking with AI adoption across deployment stages so that needs can be met as they evolve. Verizon offers Verizon AI Connect,¹ a suite of solutions that can help organizations meet their networking and infrastructure needs through the AI journey.

The drive for AI network readiness

AI has rapidly become a business imperative for profitability and competitive survival. Organizations large and small find themselves in a position where they must embrace and leverage AI just to keep pace. But deploying AI is not just about reaching an internet-based chat bot. It involves data movement, model tuning, and increasingly, agents that can orchestrate and execute semi- or fully autonomously.

Successfully embracing AI to transform business processes and workflows in production, at scale, requires a complete reassessment of IT infrastructure. Omdia research has found that 80% of organizations are making or planning significant infrastructure investments to support new AI initiatives,² encompassing all technology domains and, typically, leveraging

¹ Note: Verizon AI Connect is not available outside of the US, UK, and France.

² Source: Omdia Complete Survey Results, [IT Modernization Report: Cloud, Applications, and Infrastructure Amid the AI Revolution](#), June 2026.

This Omdia White Paper was commissioned by Verizon, Inc. and is distributed under license from TechTarget, Inc.

a hybrid cloud architecture. This includes networking within the data center, within the cloud, across the WAN, and at remote campus and branch sites.

But challenges with AI deployment and adoption are real. For instance,³

- Currently, less than half (49%) of organizations reported seeing a clearly positive ROI on production AI deployments. The significant investments undertaken when upgrading infrastructure for AI are a factor in those assessments.
- Managing that new infrastructure is creating its own difficulties. Nearly two-thirds (62%) of organizations agreed that managing the infrastructure for AI is currently more difficult than building the AI models themselves.

The network plays a key role in AI success. There is nearly universal agreement (88%) among enterprises that AI is making networking more critical,⁴ and this is driving changes across all network layers. Omdia has found that bandwidth capacity and latency within and between the data center and cloud are all commonly recognized constraints that require additional focus and resources.

The real question that IT leaders and practitioners alike must ask is whether or not their networks are ready for the journey. Are they already in an adequate state, or will changes be necessary, and when?

What AI network readiness means

Building an understanding of whether networks are ready for supporting AI starts with considering a range of relevant factors, from network architecture and AI sourcing strategy to scope and stage of AI adoption and deployment. Each of these will influence the choice of steps to be taken and the timing of changes.

It helps to first consider the general AI adoption and deployment lifecycle and journey. Much of the early stage of AI usage simply leverages external, cloud-based models and services, and this works well for AI teams that are familiarizing themselves with core capabilities. But as true enterprise AI projects get underway, the long-term aim often involves bringing AI technologies on-premises for cost, performance, and security reasons. Along the way toward production deployment, there will be development, test, pilot, initial release, and long-term scaling to consider. In many ways, this is typical of any new

³ Source: Omdia Research Report, [The State of AI Inference Strategies: Optimizing Deployment, Performance, and Impact](#), April 2026.

⁴ Source: Omdia Research Report, [Network Observability in the Agentic AI Era](#), April 2026.

This Omdia White Paper was commissioned by Verizon, Inc. and is distributed under license from TechTarget, Inc.

application lifecycle, although AI does bring along its own special requirements of the infrastructure, including the network, that must be accommodated.

Decisions that are made along the AI adoption journey will be specifically relevant to determining network readiness. For example:

- The use of cloud services for AI application development, in early stages or as part of ongoing, sustained operations, brings with it network requirements for cloud connectivity.
- Moving AI workloads on-premises has significant impacts on the network, both within and around AI training and inference clusters in data center or colocation facilities, between those sites for data movement and cluster synchronization, and to campus and branch sites for access and data transfer.
- Choosing to distribute AI inferencing workloads out into remote “edge” sites brings other considerations and requirements, depending on whether edge workloads will be hosted locally on-premises or will leverage edge services from a provider.
- Data hosting and transfer needs can have a major impact on networking, particularly regarding network capacity and performance targets.

Security is another essential angle of consideration, and network security must be an intrinsic element with any network planning and architecture to support AI. Specific security approaches will vary based on deployment stage and chosen location of AI workloads but generally should include privacy, encryption, and zero-trust principles.

Identifying network readiness for AI

All of these factors come together to translate into a range of specific network demands that will arise based on where an organization is in its AI adoption and deployment journey. Since networking needs change and evolve with the progress of AI projects, it is useful to consider them by lifecycle stage so that networking capabilities and capacities can be aligned and ready for testing, development, deployment, and scaling as needed.


Each stage brings with it a specific set of networking challenges and requirements that networking teams need to be ready to meet.

This Omdia White Paper was commissioned by Verizon, Inc. and is distributed under license from TechTarget, Inc.

Explorer

This is the earliest stage, and specific network requirements are still relatively simple. Essentially, the most important goal is to establish an enabling foundation of connectivity for AI exploration and development. Specific networking considerations for this stage include:


- Highly performing, highly reliable connectivity to cloud and internet-based AI services so that development teams can readily access resources.
- Data transfer volumes may begin to grow during this stage, so it's important to ensure sufficient bandwidth is available to facilitate timely and successful delivery. This issue will be of greatest concern when gathering and moving data from remote sites to a data center or cloud or from data centers to the cloud (commonly referred to as “north-south” traffic).
- Given the ubiquitous need for ensuring control over proprietary data, this is also a good stage to start working on network security to provide a protective shield during experimentation and testing.
- Finally, the Explorer stage provides programmability opportunities to tie network connectivity capabilities directly into orchestrated infrastructure management. This ensures that network resources can be readily accessed when needed, ideally in near-real time, and not held onto after they are no longer needed.



Pilot

At this stage, networking attention turns toward understanding and managing growth and infrastructure optionality and setting the stage for moving toward full production deployment of AI applications and workloads. Important network readiness considerations include:

- Accommodating hybrid environments will be a must. It's likely that an organization will be faced with a mix of cloud, on-premises, and colocation sites that need reliable, high performing, secure connectivity.
- This may be the first time that traffic, both in the form of data transfers and application interactions, really starts to increase at a baseline level. Organizations should make sure that sufficient bandwidth capacity is available or can be added as needed as soon as growth dictates.




This Omdia White Paper was commissioned by Verizon, Inc. and is distributed under license from TechTarget, Inc.

- Network latency will also become a greater concern at this stage, as true operational testing begins in earnest. Understanding and optimizing performance is also a top objective.
- Characterizing network costs and planning methods for controlling them can also be important. Traditional internet connectivity may not only be too slow; it may also end up being too costly at scale, so this is the time to look for more cost-effective (and secure) private connection options.

Deployer

Once major AI proofs of concept and pilots are finished, it's time to start scaling up the production environment for AI. While an organization might have been able to tweak its network to support the prior phases, in the Deployer stage, organizations should switch their mindset toward true enterprise-grade choices. Getting the network ready for AI deployment includes:



- Reassessing early planning assumptions around network performance, resilience, and cost. Are they working as expected, or are cracks starting to form? As usage and data volumes ramp up, new bottlenecks may emerge, so designs may need to be adjusted.
- Inferencing is where AI value is delivered, and part of the production deployment of AI will likely include moving inferencing workloads out into the organization, closer to data sources and output consumers. This means reassessing the readiness of campus and branch networks and understanding the relative impact of the two most common edge inference options—locally on-premises or at the service provider edge.
- Network readiness from a security perspective now shifts from simple network-based protection and into a fully embedded objective. At this point, enterprise-grade security should be brought to bear, adding layers that were not as critical during initial phases. Private, dedicated connectivity is a big plus here versus using the internet for WAN and cloud connections. Further, plans should include, at a minimum, consistent access and policy management and zero-trust principles.
- Finally, this is the stage at which production operational monitoring needs to be fully up and running to continuously assess health, status, and usage of the AI applications and the infrastructure that supports and delivers them. This effort should include and leverage the network viewpoint to recognize traffic loads and patterns.

This Omdia White Paper was commissioned by Verizon, Inc. and is distributed under license from TechTarget, Inc.

Optimizer

The AI building and deployment process has been completed at this stage, and productive maturity has been achieved. But the work is far from complete. In fact, there are important network readiness objectives at this stage as well, including:

- Workload placements may begin to shift. Many organizations take this opportunity to move cloud-based AI workloads back to on-premises or colocation sites, most often for cost, performance, and control reasons. Even with that, it is highly likely that some AI workloads will remain in the cloud. Connectivity needs to support this potential diversity and fluidity, ensuring adequate proficiency and security regardless of where AI workloads land.
- Programmable networks become a necessity. The operational monitoring systems set up during the deployment phase can ideally be used not only for reactive troubleshooting but also to trigger proactive automated adjustments, tuning network configurations to optimize cost and performance.
- This is also the right time, if it was not put in place by now, to ensure that network functions are included in workload management automation and orchestration, coupled with data pipelining and dynamic resource balancing.



Summary definitions of these four stages are shown in **Table 1**.

This Omdia White Paper was commissioned by Verizon, Inc. and is distributed under license from TechTarget, Inc.

Table 1: Lifecycle stages

Lifecycle stage	Definition
Explorer	Organizations at this stage are on the initial steps of their AI journey, having engaged AI technologies ad hoc. Experimentation and assessments could well be underway, and the focus is on development, testing, and planning around initial AI use cases.
Pilot	Once AI projects have been defined and committed and initial development is underway, teams enter the Pilot phase. This will include experimentation around model selection, data types and sources, and process adaptation, culminating in proofs of concept, pilot testing, and even initial production launch.
Deployer	With basic viability established, the next stage involves scaling up capacity and use cases in full production. During full deployment, initial architectures need to be revalidated, operational monitoring commences in earnest, and there may be interest in repatriating some AI applications and infrastructure from cloud to on-premises. Focus can also turn toward distributing AI inferencing to edge locations.
Optimizer	Once maturity is reached in full and AI applications are up and running at scale, the next set of objectives is focused on sustained optimization. Important here is identifying ways to maintain and improve performance, manage costs, and automate operational tasks and workflows.

Source: Omdia

The agentic wildcard

Throughout all of this, it's important to recognize one specific type of AI workload that can have significant networking impacts. Agentic AI can, depending on the approach taken to implementation, be very “chatty” on the network, driving far higher and continuous levels of traffic than human operators and often carrying very low latency expectations. Agentic AI agents can be deployed centrally, such as within ITOps, NetOps, or SecOps management platforms, or can be used by individual workers as digital assistants, installed wherever their normal workplace might be.

If AI agents are part of the landscape in an organization, extra bandwidth capacity should be added into any LAN segments where they will be active and added as a category of usage for cloud and SaaS connectivity if those agents are authorized to access external

This Omdia White Paper was commissioned by Verizon, Inc. and is distributed under license from TechTarget, Inc.

services. Further, network-level guardrails, such as AI-aware firewalls, should be part of the architecture to provide protections against accidental—or intentional—misuse or abuse of resources.

The Verizon approach to AI network readiness: Verizon AI Connect

When it comes to navigating network readiness for supporting strategic AI projects, it is helpful to have access to a service provider that has already built solutions and has proven answers to the many challenges organizations face. Such is the case with Verizon AI Connect offered by Verizon. Verizon AI Connect is a foundational suite of integrated networking and infrastructure services that can help enterprises deploy, manage, and scale AI workloads.

At a high level, the services that comprise Verizon AI Connect span high-performance network connectivity, network programmability, data center, and edge compute infrastructure and services. While any or all related services may be of value to organizations at any stage of network readiness preparations, specific components align well with the various phases as follows:

- **Explorer:** The emphasis here is on providing an enabling foundation to help AI research and development get rolling fast. Verizon can help, with:
 - **High-speed, low-latency connectivity**, which helps organizations set up and access testing environments, using wavelength services for ultra-low latency connections to major cloud providers to facilitate moving very large or growing datasets.
 - **AI-ready security**, via private connectivity (rather than using the public internet) and multi-layer governance, which provides protection for project environments.
 - **Data center colocation**, which offers on-demand compute capacity and facilities, making it easier to overcome space, power, and cooling barriers for research and development of AI applications.
 - **Programmable network capabilities** that provide agility during testing and trials, including the flexibility to quickly and easily increase and decrease network capacity as needed, on demand.


This Omdia White Paper was commissioned by Verizon, Inc. and is distributed under license from TechTarget, Inc.

- **Pilot:** During this phase, it is essential to address early infrastructure constraints to ensure that initial implementations operate smoothly. Verizon can help, with:
 - **Cost-effective data transfer.** As data volumes grow, network capacity can quickly be stretched, and costs can spiral. Verizon Wavelength Services offer high-capacity bandwidth to move large datasets more efficiently, and potentially more cost-effectively, than traditional internet-based approaches.
 - **AI-ready security** keeps initial rollouts protected end-to-end, with multilayer governance and private connectivity, to ensure data is secure and in compliance with privacy objectives.
 - **Scalable facilities** to host AI workloads via data center and colocation services. This helps resolve the increased demand for space, power, and heating/cooling as volume increases and AI projects move from development into initial production deployments.
- **Deployer:** While working to achieve scale across the organization, a new set of capabilities is needed to protect performance, control costs, and harden the environment. Verizon can help, with:
 - **Ultra-low latency for edge computing,** via extensive fiber network and wavelength services, moving inferencing closer to where the data is, improving the performance of AI applications and reducing data transfer congestion upstream.
 - **Flexibility and cost effectiveness,** via solutions that scale with AI needs, ensuring network investments deliver maximum value as AI initiatives grow.
 - **Enterprise-grade network security** via private connectivity and multi-layer governance designed to protect data and minimize exposure to internet threats, establishing a strict zero-trust security posture for AI applications and systems.
- **Optimizer:** Once the shift from deployment to fully functional AI production is in place, the focus can turn toward achieving the best possible cost and performance outcomes and making networks an integral part of automated workflows. Verizon can help, with:
 - **Specialized AI compute power for AI workloads.** Verizon's strategic partnerships with leading-edge vendors enable access to the most recent and advanced GPU chips via GPUaaS. With platforms optimized for training, fine-tuning, and inferencing, AI teams can stay on top of the best potential configurations at all times.
 - **Programmable, software-defined networking** via self-service tools and open APIs. This reduces manual ticket-based management with portal-based self-service capabilities to dynamically scale, control, and connect network functions across AI infrastructure automation and orchestration.

This Omdia White Paper was commissioned by Verizon, Inc. and is distributed under license from TechTarget, Inc.

- **High-end, virtually unlimited network scale for north-south and east-west connectivity.** Verizon's dedicated private connection solutions, including wavelength services and dark fiber, minimize congestion and latencies for re-platforming projects as well as synchronizing distributed GPU clusters.

Importantly, these services are all backed by the deep technical and operational expertise of the Verizon team and by service-level agreements that ensure a commitment to delivering to expectations. Verizon stands as a capable partner, ready and able to help facilitate, accelerate, and ensure that networks will be ready, regardless of where organizations are on their AI transformation journey.



Conclusion

AI adoption has become a strategic business imperative. But getting there is not a simple process, as complex infrastructures must be designed and deployed. Network readiness plays a critical role in ensuring the delivery of the foundational, reliable, and high-performing connectivity essential to AI success.

Organizations should prepare by understanding the readiness of their networks to support AI project demands. These requirements change over time, depending on progress along the AI adoption and deployment lifecycle, from exploration through pilot and into production, and eventually on to optimization. Verizon AI Connect is a suite of solutions designed to meet networking and infrastructure needs across all phases of AI transformation, from enabling initial development to securely scaling and tuning production operations, all delivered by Verizon, a proven and trusted partner.

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
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.