

Sponsored content | White paper

A guide to upgrading your network for AI workloads

In the new AI ecosystem, traffic is growing exponentially. That's why it's critically important for organizations to ensure that their networks are ready to support AI-driven digital transformation initiatives.



CIO

Sponsored by

LUMEN®

Organizations eager to gain competitive advantage from artificial intelligence (AI) are racing to deploy the technology for a variety of potential benefits: improved customer engagement, optimized business processes, data-driven insights, and real-time actionable intelligence.

According to an IDC InfoBrief¹ on digital transformation, 81% of larger enterprises and 67% of midsize organizations are aggressively pursuing AI initiatives or plan to in the next 12 months. And AI has supplanted security and cloud deployments as the main investment priority for enterprises in 2024, according to the report.

But before organizations get too far down the road, it's critically important to make sure the network is ready to support AI-driven digital transformation initiatives. AI requires that organizations rethink, rearchitect, and reinvest in their networks, which likely have been cobbled together over the years and typically rely on outdated modes of connectivity such as Multiprotocol Label Switching (MPLS) or legacy virtual private networks (VPNs).

"The networking protocols, their behavioral assumptions, and subsequent network architectures we have today are no longer relevant in a multicloud and AI world," says Lumen Chief Technology and Product Officer Dave Ward.

An AI-ready network needs to deliver the high bandwidth and low latency that enable the movement of massive data sets across multicloud and edge environments. It also needs to be customizable, modular, resilient, secure, and on-demand and run on a fiber optic underlay that can dynamically scale to meet rapidly increasing demands as well as support additional services such as security and policy enforcement.

Unfortunately, a staggering 86% of CIOs don't think their enterprise networks are prepared for the AI ecosystem, according to IDC's Enterprise Horizons report.² And when CIOs were asked to name their biggest challenge when it came to preparing their networks for AI, the No. 1 answer was the lack of an external IT partner.

For CIOs to get their networks up to speed, they need to team up with an experienced network service provider that has the technical expertise and the high-speed fiber optic capabilities necessary to deliver secure, dedicated bandwidth for AI workloads.

"The networking protocols, their behavioral assumptions, and subsequent network architectures we have today are no longer relevant in a multicloud and AI world,"

— **Dave Ward , Lumen Chief Technology and Product Officer**



The challenges associated with preparing the network for AI

In many ways, the AI revolution is already here. In other words, even as the organization makes decisions on how to proceed in building its own AI-driven applications or to what extent it plans to develop its own large language models (LLMs), AI is coming in: through the back door, through ChatGPT, or embedded in applications that end users rely on to do their job.

Lumen Senior Vice President James Feger points out that “even if you don’t believe you are interacting with AI, there’s a great likelihood that the applications you’ve been using for years are becoming more chatty on the network, because they’re sending more and more traffic out to a cloud resource to run AI.”

If the WAN network isn’t adequate to support the AI ecosystem, organizations run into unacceptable latency that impacts application performance, particularly in light of the real-time requirements of AI-driven applications.

Ward explains, “There’s currently a lot of hype around software-based VPNs, SASE [secure access service edge], and SD-WAN [software-defined wide area network]. While that’s all well and good, they weren’t meant to be leveraged for high-capacity, very noisy applications. They were meant to be used for basic remote office connectivity. And in the ecosystem of AI and the proliferation of distributed traffic associated with AI, SD-WAN and SASE can become a technological bottleneck.”

Feger adds, “If you have bottlenecks or you have unacceptable latency or you have performance degradation, you can ultimately

work through those. But the inability to launch new applications is probably one of the biggest concerns a lot of industries have, and that would include healthcare, finance, retail, education, and manufacturing. You could go through practically every industry, and the inability to launch new applications is a direct hindrance to the business.”

The key takeaway for CIOs is that AI applications are fundamentally different from traditional apps. “In the past, when a new application would launch, it would be a blip on the network,” says Feger. In the new AI ecosystem, applications are constantly pulling data from end users or other applications to acquire intelligence and to share intelligence, so traffic is growing exponentially. And farther down the road, AI will be talking to AI, putting even more strain on the network.

“The large language model evolution for generative AI has resulted in data sets that reach easily into petabytes,” says Ward. Models for enterprise applications typically train across a wide range of data sets, from tens of terabytes of text and structured data today, with the expectation of an order-of-magnitude shift with multimodal models.”

On top of that, traditional edge-to-data-center or data-center-to-cloud traffic patterns are obsolete in an AI ecosystem. “AI needs data, data needs data centers, and data centers need to be connected,” says Feger. The new ecosystem needs to seamlessly connect access networks, backbone networks, cloud networks, and data center networks. “And ultimately, for the AI ecosystem to thrive, you need to be able to connect any of these networks from wherever you are.”

“Even if you don’t believe you are interacting with AI, there’s a great likelihood that the applications you’ve been using for years are becoming more chatty on the network, because they’re sending more and more traffic out to a cloud resource to run AI.”

— **James Feger,**
Lumen Senior Vice President

Steps to take on the road to AI readiness

Here are the key steps an organization should take to make sure that the network can support the requirements of an AI ecosystem.

Mindset shift. The first step is a mindset shift away from the idea of fixed connections between point A and point B and toward the concept of a flexible, resilient network topology that provides connectivity and services in a mesh- or fabric-style architecture.

The organizations that succeed in their digital-first strategies by leveraging their high-speed underlay network to drive positive business outcomes will be the ones that concentrate on the linkage between endpoints and also make services a priority.

Feger adds, “You need to be able to have an architecture mindset as to where the ecosystem is going and what’s driving the ecosystem. And today what’s driving the ecosystem is AI.”

Capacity planning. Organizations are becoming more distributed, with edge data centers, multiple cloud service providers, and employees working from everywhere. CIOs must provide the raw bandwidth that can handle current day-to-day traffic flows,

but they must also account for future AI-driven growth as well as other bandwidth-intensive events such as disaster recovery and data backups.

It makes sense to work with a service provider that can deliver on-demand capacity that can scale as requirements increase. Ideally, the service provider should have the ability to monitor traffic patterns and automatically scale up to stay one step ahead of potential bottlenecks.

Architecture. Organizations need to not just think about volumes of data but also traffic flows in terms of how workloads function, which applications are driving the most traffic, and what’s the best way to route that traffic. The network architecture needs to be flexible, modular, redundant, resilient, and scalable, which implies a mesh or a fabric rather than fixed point-to-point connections, as well as dedicated private links rather than best-effort public broadband.

Policy enforcement. Another key consideration is how to enforce policies related to access control, security, and performance in a multicloud environment. Organizations need to have visibility into network traffic, they need to be able to aggregate and analyze telemetry data across the network, and they need enforcement points. Multicloud gateways, deployed at the data center or the edge, address this issue by aggregating traffic data and providing a holistic picture of what’s happening in the network from an application perspective.

The Lumen Advantage: Private Connectivity Fabric

Lumen® Private Connectivity FabricSM (PCF) is an architecture that ena[®]bles organizations to

deploy a network mesh encompassing the full spectrum of connectivity options, including waves, Ethernet, Internet Protocol (IP), and private networking. This approach enables bandwidth control, latency, and redundant paths, driven by APIs and controlled by routing protocols.

PCF leverages the company's vast network,

one of the world's most connected.

This provides the capacity to handle massive data sets, combined with a near-zero-latency edge fabric and the ability to control every layer in the network architecture.

At its core, PCF is composed of ports and services within that port. The port is the head/tail of the pipe, and the services enable bandwidth, redundancy, and connectivity control for voice, security, unified communications, software as a service (SaaS), and storage.

Ultimately, CIOs recognize that AI readiness is critical for the success of digital transformation initiatives. And enterprises can't do it alone; they need to partner with a service provider that has the underlying physical infrastructure as well as the complementary services that make up a full-stack solution.

Although hyperscalers target future AI revenue streams, most enterprises focus on improving efficiency and making data-driven decisions. This is where Lumen Private Connectivity Fabric excels, addressing today's AI data needs without sacrificing performance or security.

Learn more about Lumen PCF at

[Lumen.com/PCF](https://lumen.com/PCF).

This content is provided for informational purposes only and may require additional research and substantiation by the end user. In addition, the information is provided "as is" without any warranty or condition of any kind, either express or implied. Use of this information is at the end user's own risk. Lumen does not warrant that the information will meet the end user's requirements or that the implementation or usage of this information will result in the desired outcome of the end user. This document represents Lumen products and offerings as of the date of issue.

¹ IDC InfoBrief: "Empowering Digital Transformation."

https://assets.lumen.com/is/content/Lumen/idc-info-brief-empowering-digital-transformation-lumen-digital-idc-us52378224-ib?Creativeid=54ca170a-4a56-4fcb-ad96-Do_d8f29bef60f2

² IDC, Enterprise Horizons 2024, June 2024.