

Infrastructure built for specialized cloud providers

Deterministic connectivity that helps you scale GPU platforms faster, help protect SLAs, and monetize capacity with confidence.

The Challenge

As AI platforms scale, the network becomes a critical part of how the business operates. Connectivity starts shaping performance consistency, customer trust, and how confidently you can grow.

Often, the network becomes the limiting factor as GPU density increases. Latency variability, slow inter-metro turn-ups, and fragmented network paths introduce unpredictability into training runtimes, put SLAs at risk, and leave valuable GPU capacity sitting idle.

That's when the network stops being just infrastructure and starts acting as a control point - determining how quickly you can expand into new regions, how reliably performance scales, and how efficiently capacity turns into revenue.

The Opportunity

As GPU platforms grow, taking control of the network becomes essential. Predictable performance, faster expansion, and the ability to bring capacity online quickly all depend on it.

How Lumen can help

Deliver deterministic performance at scale

Lumen provides high-capacity, low-latency connectivity between GPU clusters, data centers, and metros. Traffic is performance-isolated and engineered for AI workloads, helping keep training and inference consistent as demand grows.

Protect platform SLAs and customer trust

Built-in resiliency, redundancy, and security reduce variability across network paths, to enable performance to hold up under load and SLAs remain enforceable as environments expand.

Maximize GPU utilization and unit economics

By reducing activation delays and network bottlenecks, Lumen helps minimize idle GPU capacity and align network investment with real usage, so deployed infrastructure can start to generate returns sooner.

Common use cases

- **Distributed AI training and inference:** Deterministic, low-latency connectivity between GPU clusters helps deliver consistent training runtimes and predictable inference performance, even as workloads scale across sites.
- **Rapid multi-metro expansion:** High-capacity inter-data-center and inter-metro connectivity support faster region launches and customer onboarding, to reduce time-to-activation as new capacity comes online.
- **High-volume AI data movement:** 400G transport enables efficient movement of massive datasets between locations, helping keep GPUs utilized and unit economics in check as demand increases.



"Prometheus Hyperscale's partnership with Lumen furthers our aligned missions of building the backbone of the AI economy. Lumen's future-ready fiber is the connective fabric we need to provide best-in-class connectivity and drive AI innovation."

– **Trenton Thornock**

Founder and CEO, Prometheus Hyperscale

Network built for AI scale

Lumen offers flexible connectivity options designed to support GPU platform growth, helping you scale capacity, manage performance, and expand with greater confidence.

High-capacity connectivity where GPUs run

Bring bandwidth to where compute lives, without redesigning the network as demand grows.

- **Optical Capacity** (100G/400G) to support high-bandwidth connectivity between GPU clusters, data centers, and metros
- **Private Inter-DC Connectivity** for low-latency, deterministic traffic across distributed training and inference environments
- **Elastic Capacity Scaling** to better align bandwidth growth with GPU utilization and customer demand

Performance-isolated network services

Reduce variability by isolating AI traffic and engineering the network for consistent performance under load.

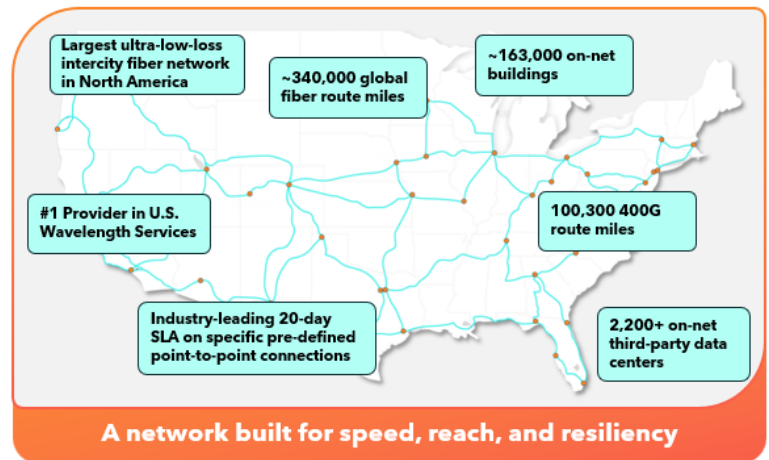
- **Dedicated, SLA-backed connectivity** designed to support GPU-intensive workloads
- **High-throughput transport** (up to 400 Gbps) optimized for large-scale AI data movement
- **Multi-path diversity** to support resilient, multi-data-center architectures
- **Built-in resiliency and security controls** to help reduce variability, protect data in motion, and support consistent SLA delivery

Architectural flexibility as platforms mature

Choose the level of control that fits today, with a clear path as the platform evolves.

- **Managed Wavelength Services** for rapid scale and simplified operations. Provides high-capacity, deterministic transport between sites without the operational overhead of owning the network.

- **RapidRoutes**, for pre-engineered, validated capacity on select point-to-point routes, with a **20-day SLA** on qualifying connections to help accelerate site and region turn-up.
- **Network-as-a-Service (NaaS)** to support dynamic capacity needs and more flexible consumption models. Help enable on demand bandwidth changes and cloud-like consumption as traffic patterns shift.
- **Dark Fiber** when utilization density and scale justify full ownership and operational control of unlit fiber strands. Supports long-term platform infrastructure for organizations that want to engineer and operate their own network.
- **Managed Optical Fiber Networks (MOFN)** for increased control and performance isolation. Extends private optical connectivity across data centers and cloud regions while reducing operational burden.



The Lumen network offers up to **400 Gbps** transport, **125+ cloud on-ramps** and **2,200+ third-party data centers**.

Why Lumen?

Lumen delivers secure, high-performance connectivity built on one of the world's most expansive and deeply peered networks. With 30+ years of experience, we simplify complex networking challenges to help connect people, data, and applications across clouds, data centers and locations. Our continued investment in AI-ready infrastructure and metro expansion helps ensure your business is future-ready.

866-352-0291 | lumen.com | info@lumen.com