

REPORT

# The Future of Networking

Embracing AI and Cloudification with Lumen<sup>®</sup> Private Connectivity Fabric<sup>SM</sup> (PCF)



LUMEN<sup>®</sup>

# Introduction

We're more than two years past the disruption of the pandemic that transformed traditional work into hybrid and remote work. SD-WAN is now a feature of secure access service edge (SASE), and dedicated internet access (DIA) is approaching displacement of the majority of the MPLS VPN private network infrastructure in enterprises.<sup>1</sup>

Cloudification and the pandemic made seismic impacts on enterprise network architecture and business relationships. In our post-pandemic digital era, if your IT isn't running out of the cloud already, you're lagging and falling further behind by the hour. It's also no longer feasible to meet employee and customer demands by trying to DIY enterprise IT or own and operate everything yourself.

Our current position as an industry in this new era highlights how we reactively stumbled into solutions barely adequate for today's cloud. We're stuck with a connectivity bottleneck and business services architectures with an unclear path forward for the private networking requirements of multi-cloud, multi-data center, Software-as-a-Service (SaaS) IT and the adoption of artificial intelligence (AI). In one sense, we managed to solve operating networks via swivel chairs by rebuilding new, shinier swivel chairs that were based in portals.

The industrial transition of how to operate an enterprise, the rise of AI, all content being delivered over-the-top (OTT), and cloud disaster recovery, storage and archive are all commonplace now.<sup>2</sup> We're facing a generational inflection point in network demand, and we need new architectural design patterns.

---

# AI data centers and on-ramps catalyze new internet architecture designs

AI has arrived, changing the players and geography of what we thought was “peak” cloud, stomping on the accelerator of enterprise transformation, and exposing the flaws of the hastily constructed scaffolding between the end of MPLS VPNs and the layers of internet overlays currently acting as a substitute. AI is all about unlocking data. Even after the first few waves of enterprise digitization, a sizable portion of data in enterprises is still trapped in multi-modal docs, shared silos/repos and custom databases. Formatted and unformatted data that large language models (LLM), at the very least, will give access to will spawn new enterprise-specific, SaaS-like apps that mine that data and continue multi-faceted returns to an enterprise.

Unlocking this data will create a major set of new problems to solve to meet the needs of the multi-cloud and AI cloud business world that cannot be met by our existing solutions.

The networking protocols, their behavioral assumptions and subsequent network architectures we have today, that are being taught in universities and written in our sacred textbooks are no longer relevant in a multi-cloud and AI world.

We need a cloud-first, cloud-triggered, cloud as a first-class citizen, cloud-networked internet.



## The relationship of WAN to AI

There are currently three main production steps to AI industrialization:

- Training and retraining
- Inference
- Reinforcement

And there are two general model or application architectures for AI that mirror those of the cloud evolution that preceded AI:

- **Contained**—where the models are not split between clusters at different sites
- **Distributed**—where models can reference other models that may reside in other locations or be sharded across different clusters or data centers

Most current AI literature and experience are focused on the contained architecture.

Training requires massive amounts of GPU resources that are only affordable and available in cloud infrastructure—either from well-known Internet-as-a-Service (IaaS) providers or new GPU-as-a-Service (GPUaaS) startups.

Available WAN bandwidth has implications for both consumers and providers—specifically around efficiency and performance.

The LLM evolution for general AI has resulted in data sets for recent large models that reach easily into petabytes. These have continued to grow over time.<sup>3</sup> Models for enterprise applications typically train across a wide range of data set sizes, from tens of terabytes of text and structured data today, with the expectation of an order of magnitude shift with multi-modal models. These models add voice and video data to some early use cases that have reached a petabyte.

Retraining of models—whether periodic, continuous or event-driven—can be a significant part of the model lifecycle with dependencies on how large the data sets are, how fast they grow and how much data variance is being introduced with growth.

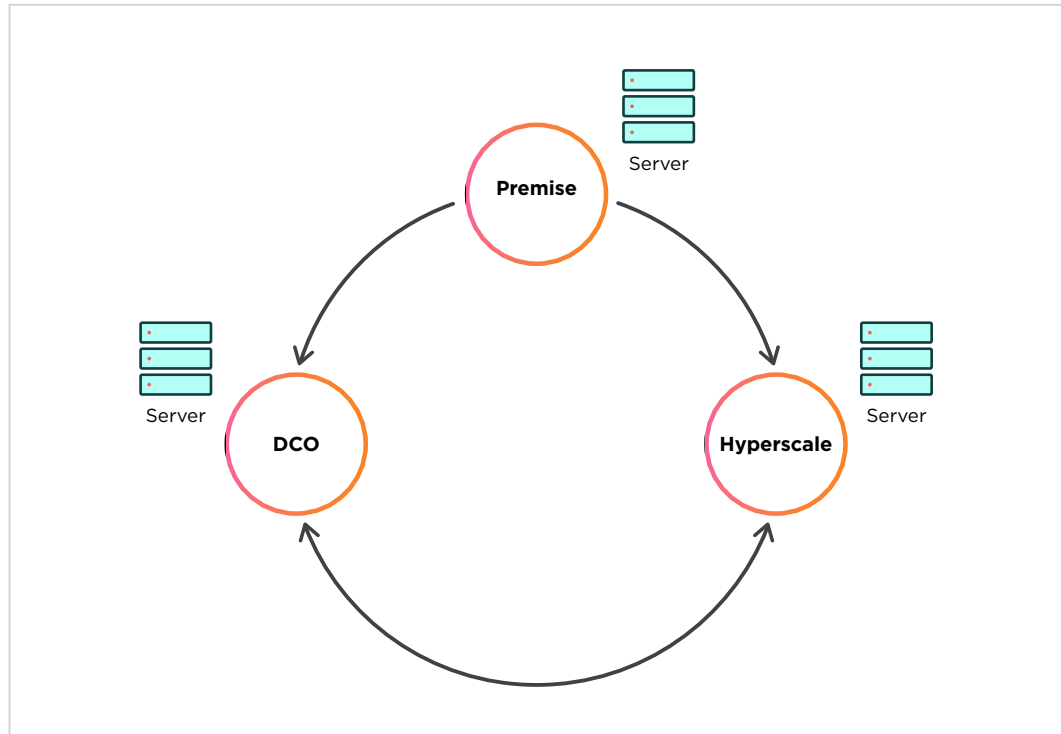
In regard to this training and retraining step, available WAN bandwidth has implications for both consumers and providers—specifically around efficiency and performance.

For consumers, bandwidth shortages are already the most commonly cited performance issue with AI applications or workloads.<sup>4</sup> For enterprises that choose not to house their confidential data permanently at their training sites, the movement of data (loading) for training and retraining models can cause significantly larger spikes in utilization due to being a new application with a much larger peak-to-average ratio (PAR).

## An AI strategy is a cloud/data center partner strategy, a data strategy and a network strategy.

**Figure 1**

Contained AI training and retraining provokes large spikes of bandwidth utilization to load data into models—not unlike high-performance computing workloads.



AI providers have AI business operations that range from simply hosting GPUaaS (GPU cluster providers for enterprise model development, training and inference) to running their own massive business-critical AI factories and applications.

There are loose correlations in network resource needs between the data set size, cluster size and bandwidth for the smaller-scale hosted model of AI industrialization.

For GPUaaS cluster providers, operating out of a third third-party center operator (DCO) colocation space, using cluster sizes up to 4,000 GPUs and data sets around a petabyte, a provider can see demand bursts that require a 100 Gbps connection (multiple for redundancy) or even several 100 Gbps.

At the higher end of AI industrialization, providers factor in power efficiency within power contracts in a variable mix that determines profitability. The result is often a desired state of running the entire infrastructure complex at consistent and high-power levels that are gated by the bandwidth needed to get workloads and their data sets to them.

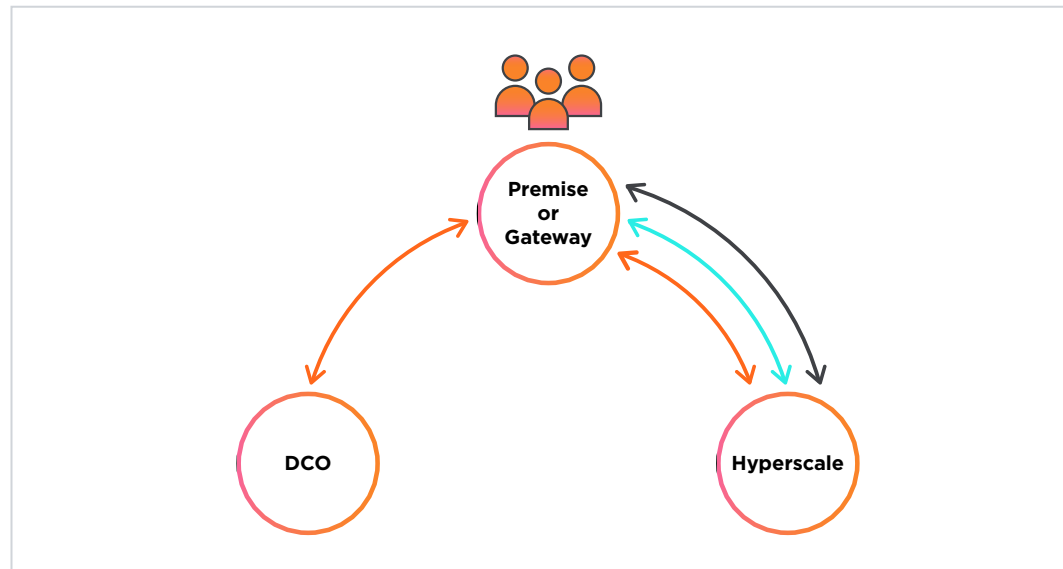
Unsurprisingly, there are larger clusters in some high-powered organizations (by an order of magnitude over the simpler hosting service configurations), and our discussions with them are already around high-count 400 Gbps wave connectivity and high-count managed fiber.

Early surveys of enterprises regarding AI inference plans, which define how AI utilizes models to make conclusions or predictions based on data, show a large amount of DCO and public cloud involvement.<sup>4</sup>

If you look closely, you'll see that enterprise AI inference behaviors are less demanding than training and more like business-critical SaaS—though at much higher bandwidth with lower latency requirements.

**Figure 2**

Visualizing contained inference as a cloud bandwidth consumption multiplier—existing SaaS, AI augmentation of existing SaaS, plus new business-critical AI-as-SaaS. Flows are asymmetric.



Combined with the pervasive inclusion of AI assistant applications in existing SaaS, one could argue that the enterprise apps emerging from the present stage of AI evolution will drive the average cloud consumption bandwidth of enterprises up by a conservative multiple of double to triple current levels—with dependencies on both the prompt type (less for text/query prompt contexts of today and potentially larger if the prompt is image, video or voice sample) and frequency.

Across the enterprise, link speeds are bound to transition. The combination of training, retraining and inference could move core 10 Gbps connections to 100 Gbps and 100 Gbps to 400 Gbps.

Access links are also bound to transition, bringing into question the cost and efficacy of public network overlay solutions and potentially making “fiber avoidance” moot.

**Your cloud core is going to become an AI fabric, and a large networking demand uplift will follow.**

The final AI production step is just beginning. Reinforcement can be continuous and more than likely distributed, driving high-speed peering between the reinforcement engine owners and LLM trainers. This is the ultimate core state of the AI infrastructure industrial complex.

Keep in mind, these discussions and estimates are for simple production relationships in training and inference of the contained deployment model.

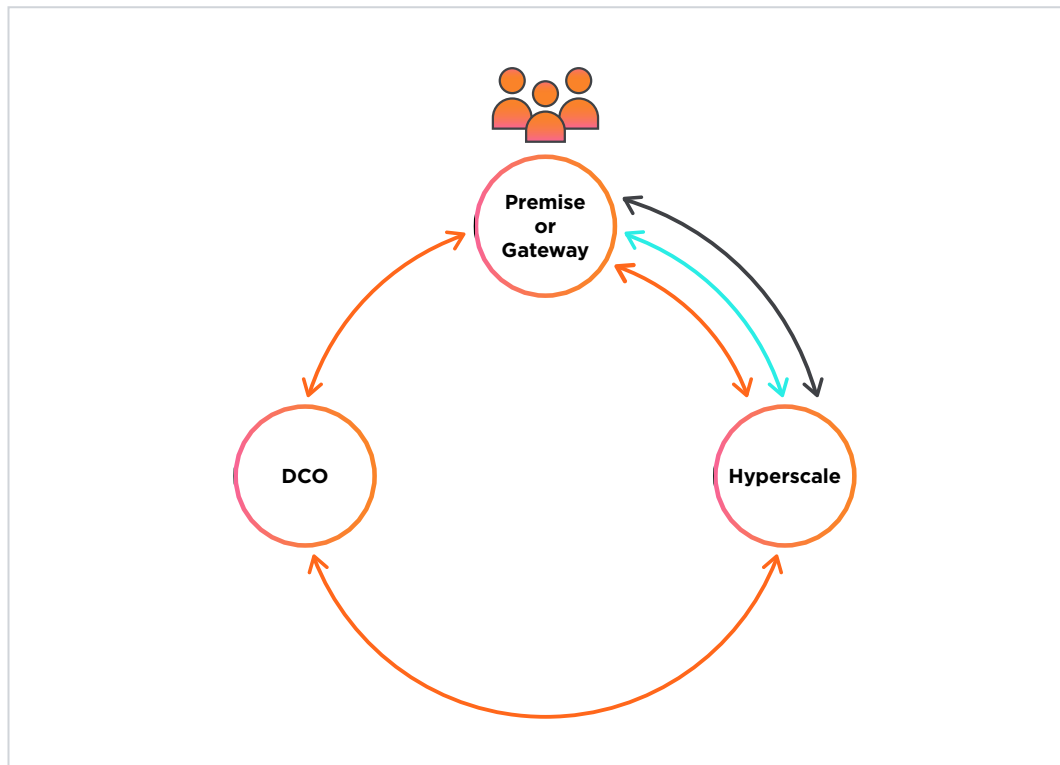
A more modular architecture is evolving that may make distributed feasible and potentially desirable. Distributed is already a reality in the sharding of large training models across different data centers.

If we look at the recent development patterns of cloud computing and view models today as application building blocks for tomorrow, a distributed model for inference comes into view: more complex applications leveraging models that call other models. In fact, this is already happening to a degree as we move into multi-modal AI.

Certainly, inference has to be imagined as becoming a distributed application environment. Running core models closer to the user for latency reasons that call “tool” models (hosted on cloud providers) for parts of their labors.

**Figure 3**

Distributed models can close the loop in the AI core with model-to-model or machine-to-machine flows between on-prem and cloud or cloud-to-cloud.



There is an even stronger and less fuzzy signal that the AI fabric will be different—a major step up in bandwidth requirements that can be gleaned from our discussions of fiber requirements with the tenants of greenfield builds for the AI future and for the linkage to brownfield expansions that will serve AI growth. The fiber counts are stunning—144 and 288 strands depending on size and general functionality of the build (DCO park or hyperscale). Plus, grooming across our existing infrastructure to connect their existing infrastructure to the new, massive bandwidth for future industrial AI production with high-bandwidth customer connection.

## The past can't serve the AI future and failed the existing multi-cloud IT environment

Basically, the protocols and limiting architectures passed down by our forefathers and held so sacrosanct are all wrong.

### We need more bandwidth

We must wean ourselves off the tunnel and VPN spaghetti. We are in need of more bandwidth—gobs and gobs of it—and have to get fiber into our access diet today. 10, 100, 400 and soon 800 Gbps will be necessary for the volumes of traffic in an AI-embedded IT world. This data movement change is the catalyst for the next generation of enterprise access to bandwidth.

We are in need of more bandwidth—gobs and gobs of it—and have to get fiber into our access diet today.

### Speed of access is king

The greatest attribute of SD-WAN and SASE is the management interface and automation—but it's fully tied to a best-effort, consumer transport layer that is its Achilles' heel in the era of AI. Transport/underlay and speed of access are now the most important considerations, and the key is how we can keep the user experience (UX) simple, efficient and API-driven with a graphical user interface (GUI) on top like a cherry on a sundae.

### Early cloud architecture isn't suited for AI

The architectural constructs that supported early cloud adoption (meet-me rooms, on-ramps, hub and spoke, gateways to oblivion) were designed with an objective of avoiding pulling fiber, but instead pull traffic into an overstuffed closet of networking artifacts of a past era. They have evolved into an anti-hourglass design ill-suited for the demands of AI.

### Signaling must be cloud-driven

The idea of signaling intent head to tail, where the head of the signal is the enterprise premise, is fundamentally flawed in a cloud world. In the hybrid, multi-AI world, the head of the service is in the cloud. It has been instantiated through cloud UX, marketplaces and DevOps tool chains. Any signaling now and in the future would be cloud-driven.

### SDN is lagging

The software-defined networking (SDN) protocols of the 2010-2019 era held so religiously as the saviors of IT are now rendered useless by their limitations. All interfaces into vendors' hardware are now considered APIs, as vendors were too slow to open them up, and NETCONF/Yang are antiquated protocols and schemas far too rigid to compete with gRPC and Protocol Buffers. We are long past the point where controlling a box was the hard part.



## Controllers met a new reality

Debates of logically centralized or decentralized controllers and so many slides at so many conferences by so many vendors were neutered by the operational realities of running a network. Operators moved directly to a cloud-native design pattern. Controller vendors are mostly out of business, and those that survive do so not only because they're a critical piece of the ecosystem infrastructure—but also because they're cloud native.

## Legacy routing protocols have hit their limit

The legacy routing protocols are now so fragile that extending them further to completely new needs like wave connectivity would be challenging and too time consuming. We tried once before to enable a true multi-layer, multi-service interconnected mesh architecture in the “standards” way. Generalized Multi-Protocol Label Switching (GMPLA) fell flat, overloading protocols and layers—a Gordian Knot by the International Telecommunication Union (ITU) and Internet Engineering Task Force (IETF). The days of arguing this out with box and arrow diagrams at a conference are long over.

## We need to get out of our own way

Hub-and-spoke architectures, gateways, fish diagrams, tunnels, multi-encapsulation or deep extension header forwarding architectures are all obsolete or simply in the way. In an AI world, it's about massive speed, riding over photons and moving gobs and gobs of data. Extra headers and encapsulations eat bandwidth (remember ATM?) capacity, and since waves are predominant in the AI world, they aren't even viable. Vendors, standards bodies and conference speeches unfortunately took a 10-year tour of inventing 42 flavors of encapsulations and headers.

To complete the networking challenge and opportunity AI presents, we need fundamentally different solutions using new API-driven platforms to control rapidly changing mesh topologies that connect at all layers where needed. The previously conceived control and management planes are one unified plane. If there isn't an API for it, build one, or don't do it.

If there isn't an API for it,  
build one, or don't do it.

The interplay of API-driven provisioning and needs-based signaling directionality—in conjunction with the dynamic routing protocols deployed to detect topology changes, failures and connectivity of the internet—are what we need as an operational architecture. Full citizenship of API-driven, cloud-native platforms for provisioning and telemetry along with the OAM and routing protocols.

Solutions going forward are going to evolve from a diverse set of SMEs from ecosystem partners that get together and hammer out APIs and get things working to solve a problem. Standards bodies that define standardized APIs and nothing underneath are a path forward for the operational systems to interplay. Let's agree on how to communicate via APIs, but let's not dictate what API users do with them.

The solution will require a shift in architecture that finally uncouples the layers of networking required for the future from a single protocol and combines them with a platform approach that can work across all the layers of the OSI model.

Private Connectivity Fabric (PCF) describes the architecture within those solutions.

## What is the Private Connectivity Fabric (PCF) and what does it solve?

PCF is an architecture that enables network mesh building with full control of the endpoints, participants and the three and a half layers that can potentially connect them: waves, Ethernet, internet (IP) and private networking. It allows for bandwidth control, latency, load and redundant paths, driven by APIs and controlled by routing protocols within the PCF domain.

We know that the answer to “What’s next?” starts with fiber and waves. Fiber infrastructure offers security, availability, low latency and capacity that provide business continuity for present and future demands. For many customers today, when bandwidth requirements approach or exceed 10 Gbps, a dedicated wave service is the best solution. AI will tilt many customers into a 400 Gbps form-factor connection—and 800 Gbps and 1.6 Tbps can’t come fast enough.



We know that easy consumption and dynamic connectivity are now an entrenched requirement from the enterprise move to cloud.

We know a private fiber network by itself is hard to consume and operate and is relatively static—at least, for solutions on offer today.

We know the connectivity is between premises, multiple DCOs, multiple communications service providers (CSPs) and multiple internet transit providers.

We know that fiber and waves attack our bandwidth constraints, but that ubiquity and reach still require Ethernet and IP as you move a few hundred kilometers from the core.

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, SaaS, storage and CSP/DCO access. The service layer (e.g., VLANs, PWEs and VRFs) can now be delivered in an as-a-service model with complete IT and workflow control. It’s what MPLS promised but couldn’t deliver because of the protocol and critical architectural assumptions.

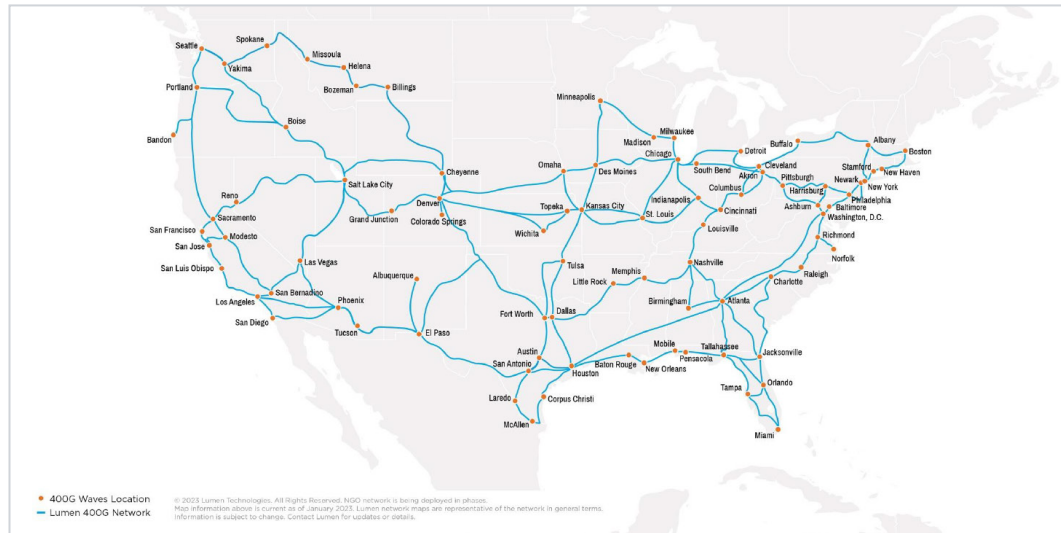
At Lumen, we are placing extra emphasis on a currently missing component/tool for building PCFs—optical connectivity that allows for the aforementioned full control of participants and endpoints of a fabric. Plus, a way to interconnect and integrate those fabrics with the rest of the enterprise network estate via Lumen Digital, a portfolio of on-demand, composable, digital solutions for networking, edge and security.

## Fiber first—the map is changing.

Interconnection (telecommunications services) are currently provided over a well-known infrastructure grid. Look at any national fiber network provider and the maps all have a high degree of overlap.

**Figure 4**

Lumen 400G wavelength network.



There are junctions between providers in select major metro areas and data centers that have saturated the surrounding areas.

Cloud consumption and expansion have pushed this infrastructure into a barely manageable situation with big tradeoffs that include interconnection through those 20-year-old “meet-me” points. Those carrier hotels were built between 1990 and 2010 (when average expected lifecycle was 15 years) and now have the capacity and power density limitations that plague the industry.

The commercials around the meet-me rooms known as cross connects (XCs) are a massive expense driven by the fiber hotel owners that make no architectural sense. They need to be completely revamped with new architecture.

This already strained design—where interconnection is hair pinned back through a choke point—completely shatters data patterns and flows with the emergence of AI. Power is now a huge issue, and serious AI requires a lot of power. In search of power, hyperscalers, startups and major data center operators are having to go farther from those well-known cores and into places that are not currently on anyone’s fiber map.

Structure Research notes, “Hyperscale platforms are experiencing the ripple effects of this and are now actively pushing for alternative carrier hotel ecosystems to be developed as an extension to the existing ones.”<sup>5</sup> And, in our own experience at Lumen, we are seeing interest in direct connection between hyperscalers as well as interest in directly connected regional AI ecosystems.

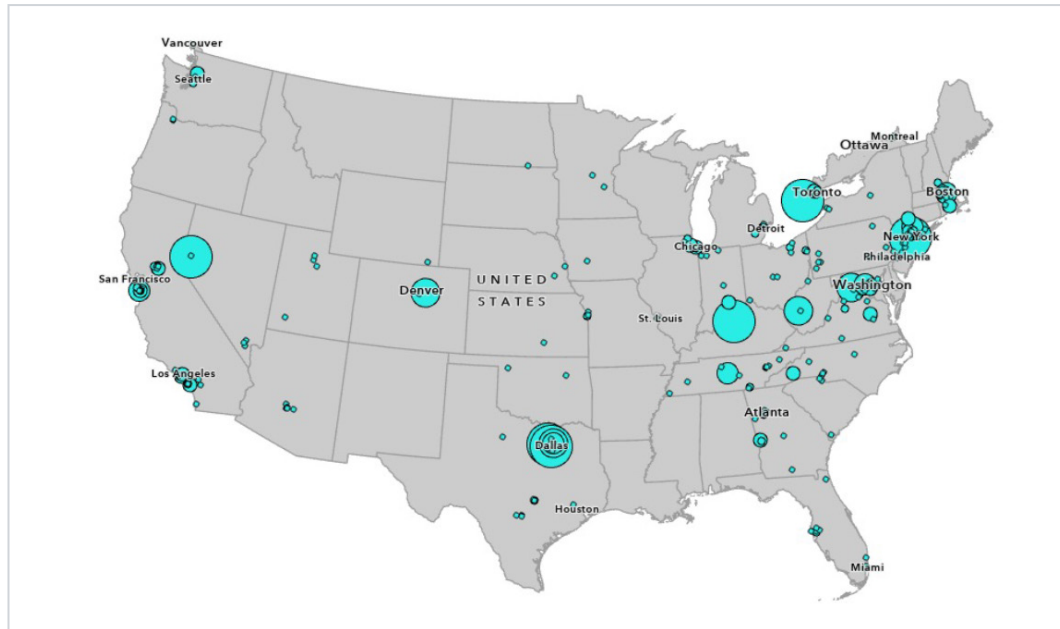
New power corridors are emerging, creating secondary and tertiary data center markets where gigawatt projects can be built, and there is tangential overlap with the existing national fiber and interconnection maps.<sup>6</sup>

- Building is shifting closer to power sources for efficiency in the Midwest and central Pennsylvania, rural Virginia and into Indiana, Wyoming, Nevada, Texas and Utah, where government incentives are attractive.
- The expansion of metro areas already known for massive data center concentration into what might be dubbed a “metro-plex” (expansion into nearby counties for space, water and power) is also stretching the national fiber map.
- Regional AI communities of interest are beginning to emerge.

Providers such as Lumen are expanding their footprints to make new fiber infrastructure accessible and imagining a new interconnection ecosystem to serve this expansion. Our recent deals with greenfield site customers to connect their new locations are a first step in this direction.<sup>7</sup>

**Figure 5**

Lumen tracked data center new bandwidth demand for greenfield and brownfield builds 2024.



### From the ground up—dig, build and connect.

Providing dark fiber and managed, lit services—enabling wave meshes—is a key architectural focus of PCF at Lumen. But while the functionality of PCF has roots at the wave level, there is a level below the “ground floor.” The very bottom of PCF functionality actually starts with where and how fiber infrastructure is built.

At Lumen, we collaborate with individual customers to design, build, light and operate fiber infrastructure to requirements including conduit, fiber and huts (as required).

While this may seem like an exclusive function engaged most often by hyperscalers, Lumen has recently seen an increase in interest from regional groups of related businesses for infrastructure builds along with the evolution of master-planned data center parks.

Beyond the drivers of individual or group fiber projects, we continuously research and explore opportunities to expand our own national fiber footprint, which is how we engage with the majority of our customers today for wave services.

## Connect—a switch for fiber.

If your conversations about business connectivity include the words “multi-cloud” and “ecosystem,” you have logically jumped the shark between existing point-to-point optical connections and the need for reorganizable point-to-multipoint or mesh networks.

With PCF, Lumen follows market signals from a developing market for waves on-demand and pushes even further by introducing managed programmable switched optical connectivity at a metro and regional level. This enables the on-demand interconnection of PCFs at the wave layer—a core building block of cloud ecosystems with access to bandwidths in the 100-400 Gbps ranges today with 800+ Gbps in the future and a functional step beyond the current generation of point-to-point fiber networking offers for mesh building. This solution is built on an evolution of the Lumen ExaSwitch® platform (launched in 2023) that now includes Ethernet, IP and integrated coherent optics.<sup>8</sup>



Our focus today is strongly metro and regional for programmable optical switching—metro fiber rings with tributaries and a switch function providing logical interconnections by design. National scale is achieved by interconnecting these regional switches using a switch “degree” connected to our Lumen national fiber network to provide Wavelengths-as-a-Service (WaaS) across our footprint.

For programmability, the operational switching technology is being absorbed in a platform API scaffold, as is the management of the connection—a controllable, on-premises component that connects to a Lumen optical connection/switching point.

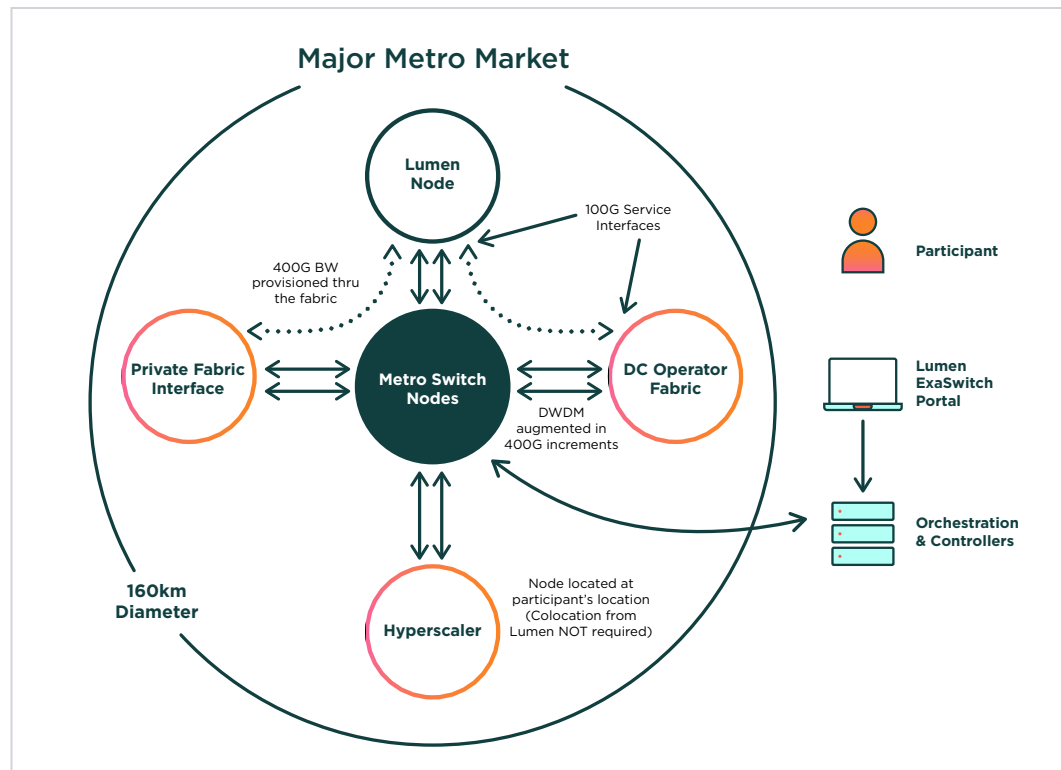
Fundamentally, the infrastructure components here are not new. They are routers with Dense Wavelength Division Multiplexing (DWDM) pluggables and transponder shelves around a Reconfigurable Optical Add-drop Multiplexer (ROADM) delivering waves, wave grooming into fibers, ENET and IP-as-a-Service. The key is using the right routers, pluggables, transport gear and software versus building it all at one silicon company that specializes in only one of those pieces.

For the economics to work, demand and capacity have to be big, and we’re finally reaching the demand point where that is possible. It has to start at 100 Gbps (400 Gbps is even better) and have different deployment sizes from L to XXXL.

With these building blocks, an operator can then serve a 160-kilometer diameter, creating a metro and data center/mega data center park DCI architecture.

**Figure 6**

Optical switching can create a metro-level “meet-me room” replacement.



As we have done with Ethernet and IP, we are applying programmability—a whole bunch of Network-as-a-Service (NaaS)/SDN automation on the top—and demand focus to known technology to make it a more useful tool in the present and future connectivity contexts. And, like those other services, delivered through a digital service platform.

## Interconnect and scale

The introduction of PCF appears fiber-first but stresses interconnectivity across all layers.


Network sites and ecosystem participants aren't all one size or one function and don't all have a single requirement for connectivity—nor are they typically all in the same regional geography.

- DIA and SD-WAN already provide a hardened solution architecture driven by the economics and practicality of using public network overlays for remote and branch, which has already been proven for bandwidths of 1-2 Gbps.
- The economics and technology for using Ethernet for building private network meshes as MPLS VPN substitutes have been an emergent solution for a bandwidth tier of up to 10 Gbps, proven for the use case of dedicated connection to public cloud.
- Wave services on shared fiber infrastructure are emerging as the dominant technology for the tier above 10 Gbps.
- Private or specialized fabrics (for example, Lumen AI fabric) of fiber and waves are directly connected to the Lumen topology, allowing for the Lumen fiber, wave, Ethernet and IP network to be fully utilized to build PCFs.

In addition to platform programmability and network operations for optical connections and mesh building, Lumen provides a service exchange capability above waves with Ethernet and IP.

Lumen enables the interconnection of private and shared fiber network wave services and, through the Lumen optical nodes at our wave switch and access points, access to the Lumen Digital portfolio for interconnection and integration of services above the private wave level. This includes national access to managed cloud connectivity, security, edge compute, SD-WAN/SASE services, Lumen® Ethernet On-Demand and Lumen® Internet On-Demand.\*

**Figure 7**  
Lumen Digital



PCF provides both an architecture and a platform to seamlessly integrate connectivity across waves, Ethernet and IP, delivering connectivity across bandwidth tiers to expand and scale a PCF using site-appropriate or function-appropriate technology.

### Private networking future: NaaS is just a consumption model

Analysts rightly point out that NaaS has been slow to take off and is confusingly applied as a concept to everything, not MPLS VPN. That's unfortunate—because we need automated, composable networks.

That confusion is because NaaS is just a consumption model.

The NaaS concept has admirably internalized the idea that networks are API-driven platform-controlled-and-delivered domains with the protocol control complexity for the domain layers it serves hidden from a rapidly changing IT consumer.

The current diffusion of the NaaS definition shows that consumption models and fulfillment requirements across the multiple layers required to build complete enterprise networks are different—which is entirely logical. Not every connection can be on-demand but can be consumed as a digitized network service—fully managed, automated and integrated via API into a composable network solution.

The NaaS experience as applied to Ethernet has also shown the weaknesses of old telco provider-to-provider interfaces, and potentially new requirements for platform-driven architectures. With no formal methodology for OSS/BSS communication, the traditional network-to-network interface (NNI) is opaque.

Most tellingly, as implemented for either IP, IP overlays or Ethernet overlays, NaaS instances to date have no ability to access the higher bandwidth, direct connectivity of optical wave services.

PCF is an architecture delivered via a technology stack and platform vision of Lumen Digital, with multiple consumption models (NaaS, MOFN, service portability) and a roadmap beyond NaaS-as-IoD (Internet On-Demand), NaaS-as-EoD (Ethernet On-Demand), NaaS-as-VRFoD (Virtual Routing and Forwarding), NaaS-as-SD-WAN and even NaaS-as-WoD (Waves On-Demand).

## Private networking future: The metro is the meet-me room

The movement into the programmatic future of optical is just beginning, with point-to-point in high-density metro areas being exposed as Waves On-Demand. However, there is little leadership in moving into mesh-building constructs and interconnection with wave switching—building and operating networks, not just connections—at that level.

PCF also allows us to move past some telco history that has been institutionalized and exposed within the current state of NaaS—the tyranny of the meet-me room and its cross-connect.

If that sounds hyperbolic, you haven't really experienced trying to build cloud-connected on-ramp alternatives to MPLS VPN at scale. What started out as a reasonable idea—that carriers would cluster at a fiber junction to enable interconnection—has morphed into a geographic sinkhole of connectivity where a possibly unwanted third party (perhaps even a competitor) must manually interconnect you to the cloud with all the opaqueness in operation and extortive cost you might expect of such a process.

What started out as a reasonable idea—that carriers would cluster at a fiber junction to enable interconnection—has morphed into a geographic sinkhole of connectivity.

We've seen this movie before. In the annals of internet history, you'll find MAE-East and MAE-West: the first internet exchanges where providers would come together and peer. Now we are to the point where meet-me rooms and fiber hotels have met the same fate—trying the same thing over and over, hoping for a different result. But the industry got to the same point with no shock—just a lot of frustrated network engineers and architects.

The existing on-ramps and meet-me room architecture have a “come here” model suited to 2020 versus meet where it makes business sense in a much more distributed digitized present. Again, an artifact of a period when the big DCOs were the islands folks wanted to visit. Now they're the islands that enterprises want an escape from to avoid being locked in, and the expansion of data center footprints in search of power for AI is really making this model look dated.

Building PCF around metro fiber and programmable optical switching is a winning argument for a new architecture with more automation, fewer hands and clearer responsibility for transport that resets the industry meet-me room concept. It's the path forward for the new AI data center exchanges and how the data center parks will be connected, and enterprises get an “AI on-ramp.”



We can enable so much more with PCF. The private peering of the traditional internet exchange can move out of the meet-me room to encompass a full metro area—using automated cross-connects at the wave level to replace VLANs on a switch. Imagine Dallas, Chicago or all of northern Virginia as a high-speed internet exchange, or a region the size of New York to Philadelphia as a single exchange.

Stay in your building because the metro fiber is already there. We can meet in a metro ring instead of in a room.

Ecosystems of businesses have always existed in large individual businesses and across entire industries, such as those you find in auto manufacturing, healthcare, media and entertainment or new affinity groupings around AI. But the tech to interconnect these ecosystems with scale and scope has always fallen short.

With PCF structure and automation, we can build business and industry ecosystems that scale. We can create new segments within the cloud economy—for example, enabling access to all cloud storage through a local connection to a switched optical exchange.

## Engaging with PCF

### How are customers currently engaging with PCF?

- Hyperscalers are already engaging with PCF below the “ground floor.” We’re building and lighting infrastructure for these partnerships. Several have also already begun to trial switched optical interconnection for high-bandwidth peering.
- Data center operators (DCOs) are already engaging by building wave-based fabrics between campuses and are in discussion with Lumen about AI ecosystems and metro optical switching. Some AI specialists are engaging around the construction of community fiber rings in regional concentrations. All DCOs can engage with Lumen to expand their customer connectivity options through Lumen Digital services or wave termination through participation in metro switched optical interconnects.
- Enterprises are interested in wave-based core fabrics and/or Ethernet fabrics with Lumen Digital extensions to build partner/provider ecosystems using composable components. Interconnecting via waves with hyperscalers, DCOs, tributary sites and partners is at the forefront of these conversations and driving our discussion with hyperscalers to engage cloud on-ramps behind switched optical interconnection to enable easy wave connectivity.

## The “next” in networking

As an industry, we have consolidated around a few lower-bandwidth, purpose-driven networking solutions for post-cloud, post-pandemic enterprises:

- A reasonable ubiquity/bandwidth tradeoff at the IP layer for remote access and branches in enterprises using public networks.

As we leap into an AI-driven network future, we need to reconsider the architectures and solutions that will move us forward.

- A reasonable 1 Gbps/10 Gbps solution in Ethernet On-Demand and private networking (VRFoD)—driven largely by the public cloud on-ramp and DCI market and presented as NaaS by some startups—is spreading. Ethernet On-Demand adoption shows that there is an acceptable bandwidth/ubiquity tradeoff between private mesh building with access to higher bandwidth tiers and the alternative of MPLS VPNs—again, if made consumable and operable behind platform APIs.

But there is a real void at the “next” level and a real opportunity to move forward. We can already see a top-down demand from AI giants for higher bandwidth services enabled by working at an even lower layer—waves and dark fiber. Demand for wave alternatives to Ethernet Private Line (EPL) at 10 Gbps and above was already growing for cloud, and AI deployments will only serve to amplify demand.

Building campus-, core- and cloud-connecting fabrics from waves with Ethernet framing will be the next step required to unlock higher bandwidth private network connectivity. It, too, will need to be consumable and operable behind platform APIs. This is the new AI on-ramp.

In the process, we will have to redraw the national network maps we’ve been using for decades and re-invent interconnection architecture concepts that go beyond the meet-me room.

While we’re doing this, the designs of our networks and the platforms that make them accessible as products—which enable the next generation of networking—have to account for the fact that organizations have needs across a solution space that incorporates all layers without overburdening protocols or over-stretching existing solutions like tunneled overlays.

- Need managed fiber? Check.
- Need waves on-demand, portability and ability to groom into fiber? Check. Check. Check.
- Need Ethernet, PWEs, IP and VRF for services and connectivity across a national backbone to get to all DCOs, CSPs, AIDC and SaaS? Check, check ... you get the point.



PCF is a heresy against old religions forced upon us by the protocols, architectures and SDOs of a long gone era. This is the era of unlimited network design patterns via seeing the world as a mesh of connectivity instead of a bicycle wheel.

PCF proposes a path to unlocking the latest and most urgently needed networking tier while integrating with the solutions available across the higher layers. It moves us past the current MPLS VPN stalemate and horrific static, hub-and-spoke networking for private

networking by using platform automation and composability to provide selective connectivity across all layers. And it does this without a standards box-diagram protocol enhancement and solution debate—the “next” for enterprise networking.

---

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. All third-party company and product or service names referenced in this article are for identification purposes only and do not imply endorsement or affiliation with Lumen. This document represents Lumen products and offerings as of the date of issue. Services not available everywhere. Lumen may change or cancel products and services or substitute similar products and services at its sole discretion without notice. ©2024 Lumen Technologies. All Rights Reserved.

---

#### Footnotes

1. TeleGeography, *2023 WAN Manager Survey*, 2023.
2. Gartner, *Top Trends in Enterprise Data Storage*, 2023.
3. Epoch AI, *Trends in Training Dataset Sizes*, 2022.
4. Flexential, *2024 State of AI Infrastructure Report*, 2024.
5. Structure Research, *Global Data Center Colocation and Interconnection*, 2023.
6. ASG, *Research Note and Investment*, July 2024.
7. Fierce Network, *Lumen’s got its eye on the AI gold rush*, August 8, 2024.
8. PRNewswire and Lumen, *Service providers can now route network traffic between data centers*, June 12, 2023.

\*Services available over the Lumen optical backbone.

## Why Lumen?

Lumen stands as a trustworthy partner for your [Private Connectivity Fabric](#) needs, offering a future-ready infrastructure that scales with the pace of digital innovation. With robust security measures and comprehensive operational support, including space, power and network management by national teams, Lumen helps ensure your business is well-prepared for tomorrow’s challenges.

Opt for Lumen nationwide deployment to mitigate risks and avoid the uncertainties of self-build projects, while benefiting from rapid deployment that helps accelerate your time to market. Our holistic approach extends beyond Private Connectivity Fabric, delivering a tailored, all-encompassing suite of network services.

866-352-0291 | [lumen.com](https://lumen.com) | [info@lumen.com](mailto:info@lumen.com)

Services not available everywhere. Business customers only. Lumen may change, cancel or substitute products and services, or vary them by service area at its sole discretion without notice. ©2024 Lumen Technologies. All Rights Reserved.

LUMEN®