09:25–09:45 | Cloud 2.0 – The Network for the AI Economy

ETHAN BLODGETT

Vice President, Product Management

LUMEN®

What is Cloud 2.0?

A transformation of cloud and enterprise core requiring 5 new capabilities

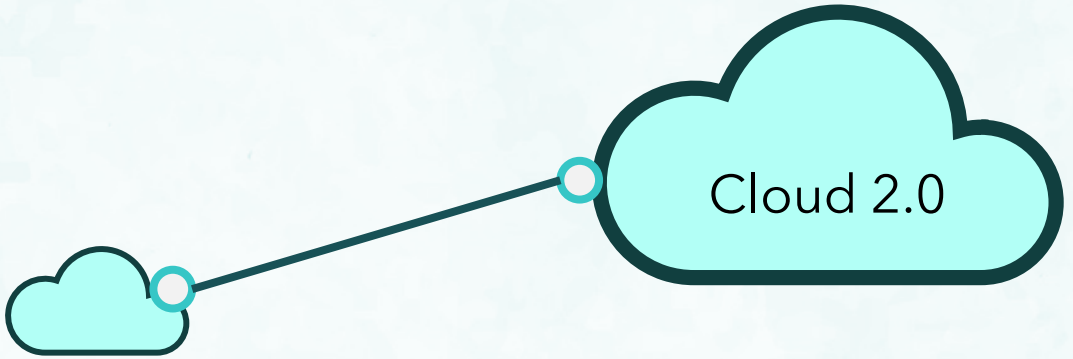

1. **Extreme bandwidth and low latency** – Scale from 400G toward 1.6T to use GPUs most cost-effectively
2. **Data Center Interconnect (DCI) foundational element** – Power the multi-cloud fabric
3. **Expansion into AI corridors** – Extend fiber and optical into areas where power exists and DCs planned
4. **Distributed on-ramps** – Programmable, high-bandwidth cloud, AI on-ramps, landing stations pre-lit to Lumen
5. **Programmable, API-first networks** – Deliver on-demand fabrics, integrated into marketplaces

AI is Driving the Shift of Cloud 1.0 to 2.0 – Networks Must Keep Up

Cloud 2.0 Demands a Network Reset

AI is redefining the data center footprint, which is growing 10x from 2025 to handle exabyte-scale demand

Legacy networks were designed for voice, internet and VPNs - **not bandwidth heavy AI**

Cloud 1.0 connectivity with internet overlays and carrier-neutral facilities **has hit its limits**

Without purpose-built connectivity, billions spent on GPUs, SaaS, and hyperscale **underdeliver**

Cloud 2.0 demands a reset, AI will overwhelm legacy network architecture

Business Imperatives:

1. **Plan for architectural shift**, more than incremental upgrades
2. **Evaluate how to evolve** the enterprise WAN for AI
3. **Prepare for Cloud 2.0,** enabling on-demand fabric and new NaaS services
4. **Move beyond** "flat" WAN, explore stratified layers and DCI at the core

Rarely a Mention of the Network Yet it's the Cornerstone

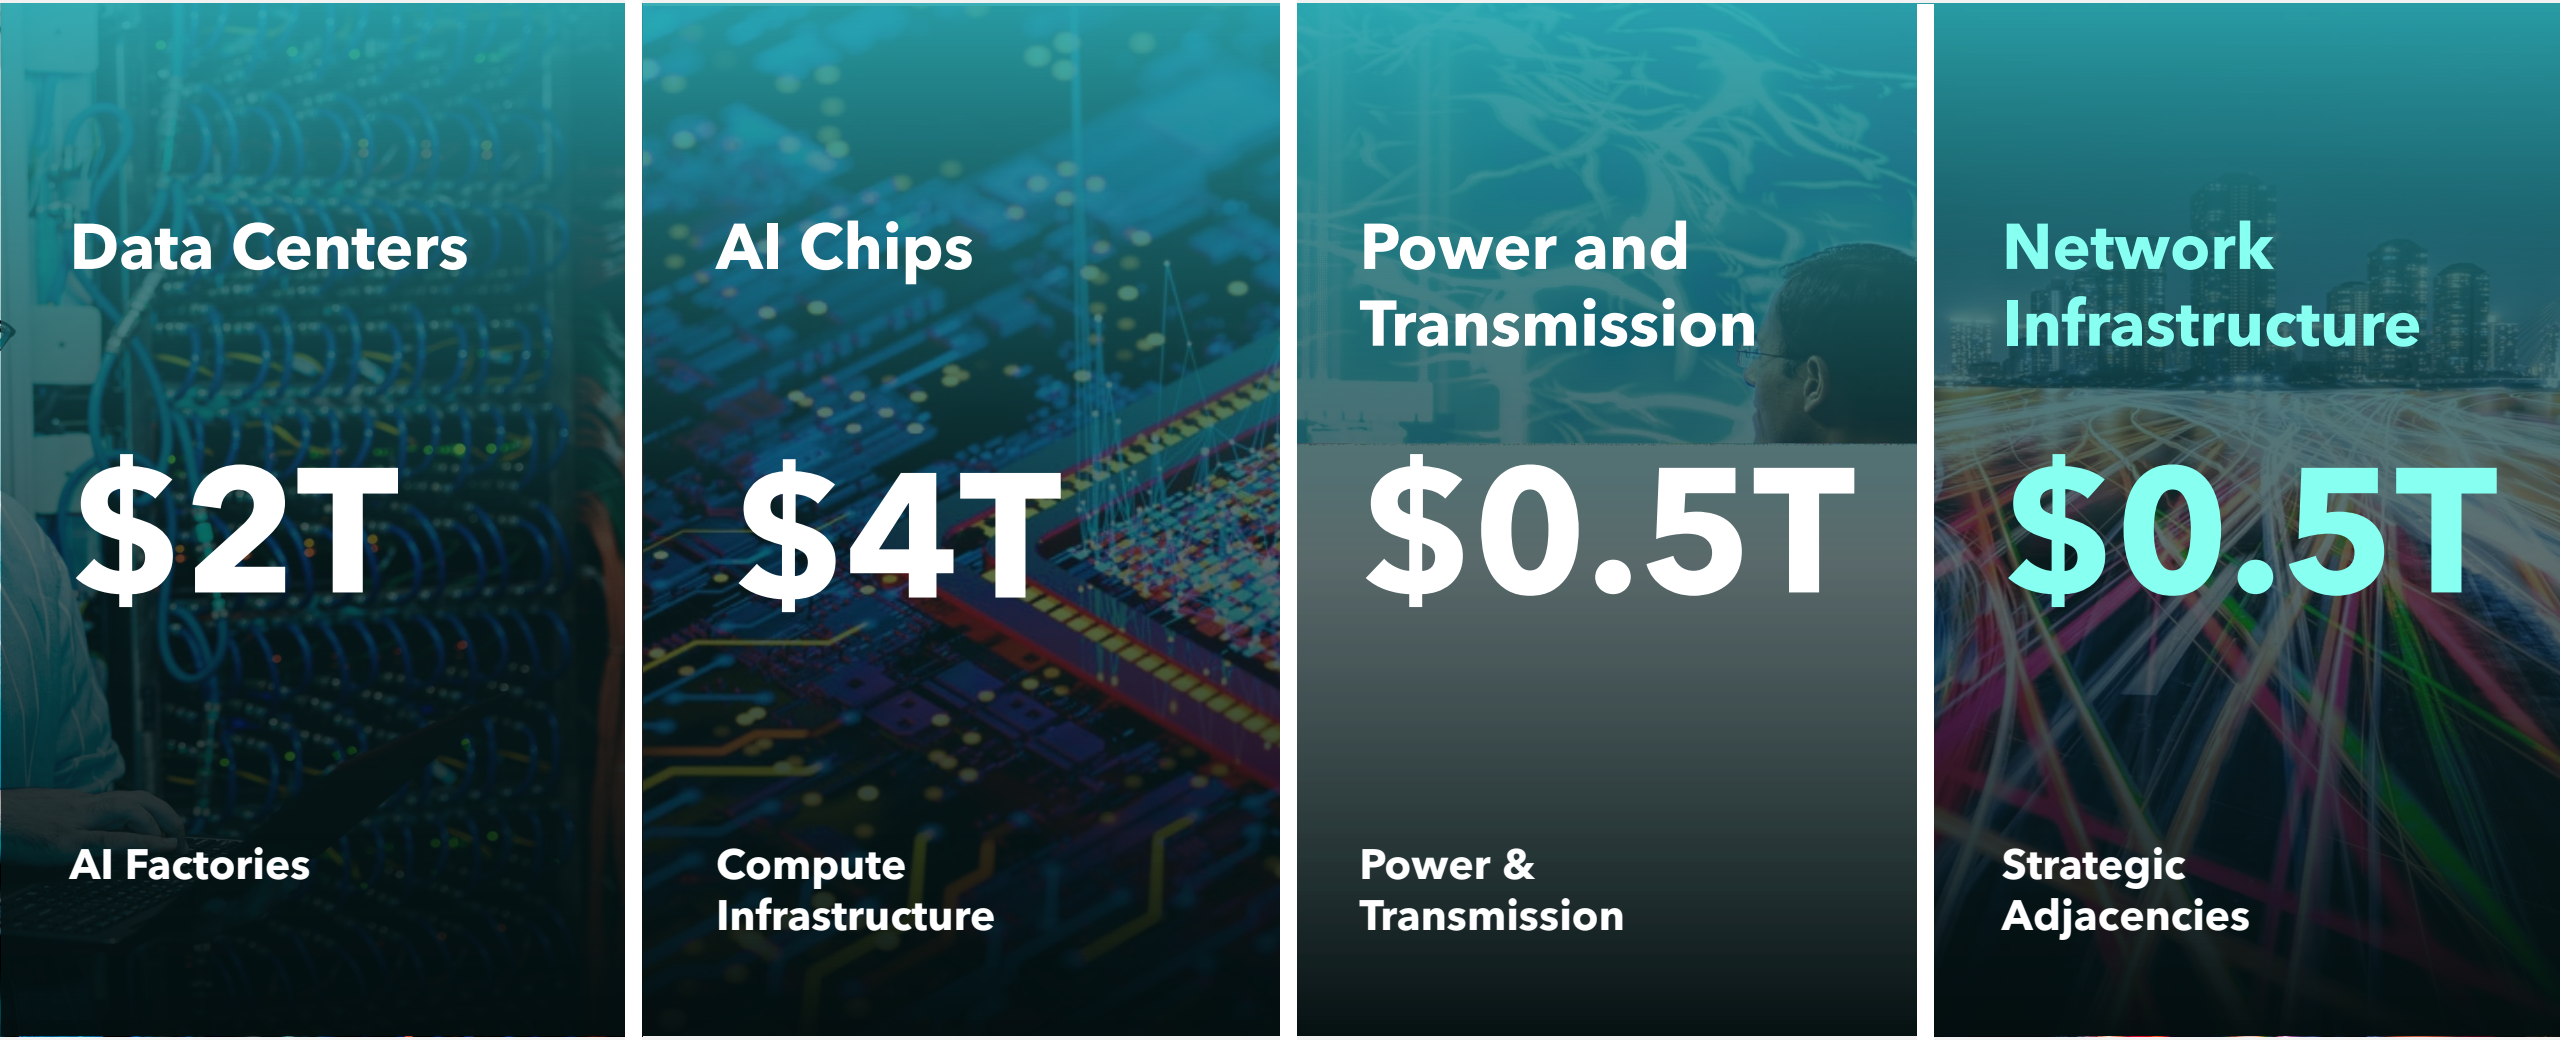

Brookfield_Building_the_Backbone_of_AI, Aug 2025

Without Bandwidth to DCs, GPU Sit Idle and Enterprises Can't Transform to AI

Enterprise network bandwidth has a significant impact on AI training cost...

Hours to egress 1 PB of training data

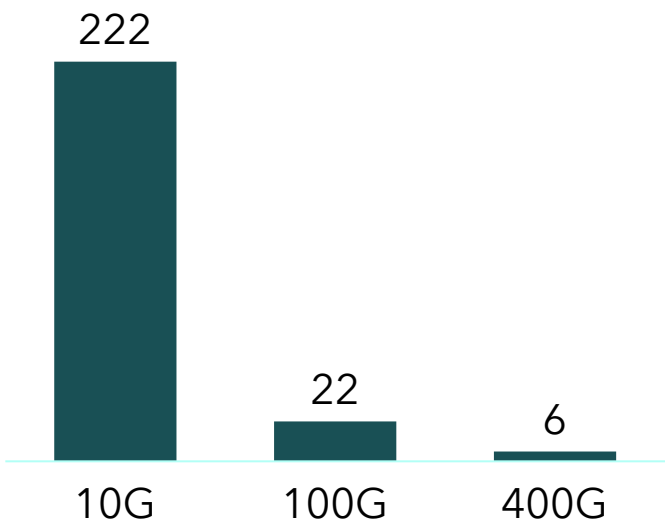

...and transferring 1 petabyte of data for training can reduce cost by **~40x** with 400G vs. 10G

Enterprise Cost before training can start ($K)

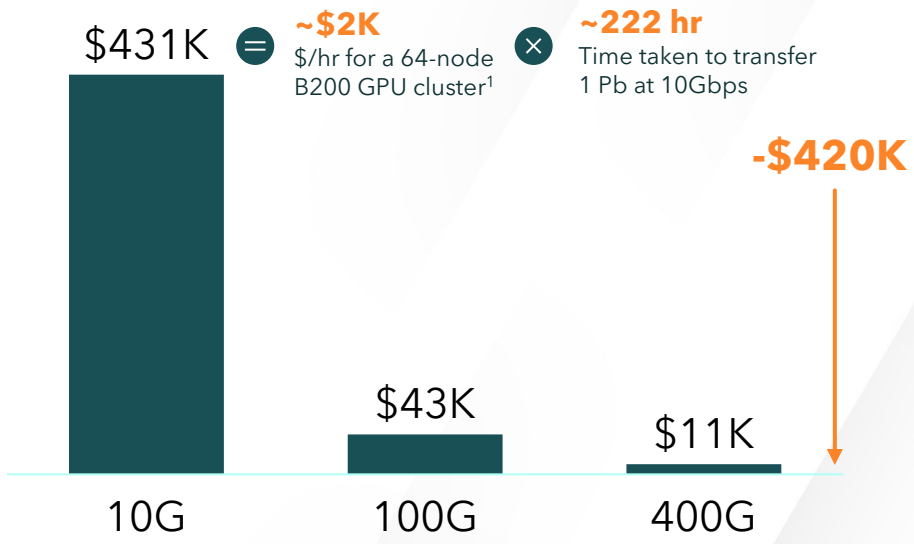

Bandwidth is a key driver for efficient GPU usage with AI.

At petabyte/exabyte scale, the network - not GPUs - sets training time and cost.

Cloud 2.0 requires enterprise networks built for exabyte-class moves; 400G+ is starting point and need for 800/1.6T obvious.

Source: Multi-source data prepared by 4MC Partners for Lumen analysis, RBC 2025 AI Market Analysis

AI is Driving Explosive Data Center Growth

~1B

Sq. ft. U.S. DC footprint by 2030[1] up from 120M built 2020-2024

DC footprint growing

10x 2025-2030[2]

DC Market All Segments 2024

Data center projects[2] in 2024

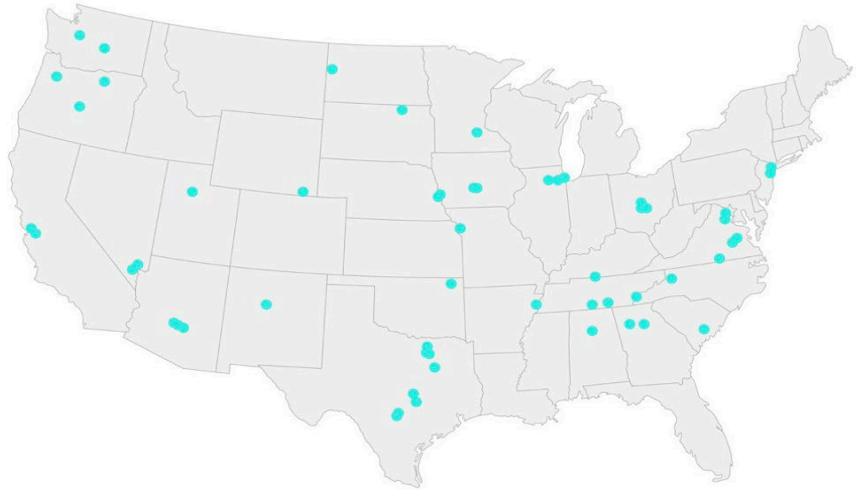

DC Market All Segments 2030

Data center projects[2] in 2030

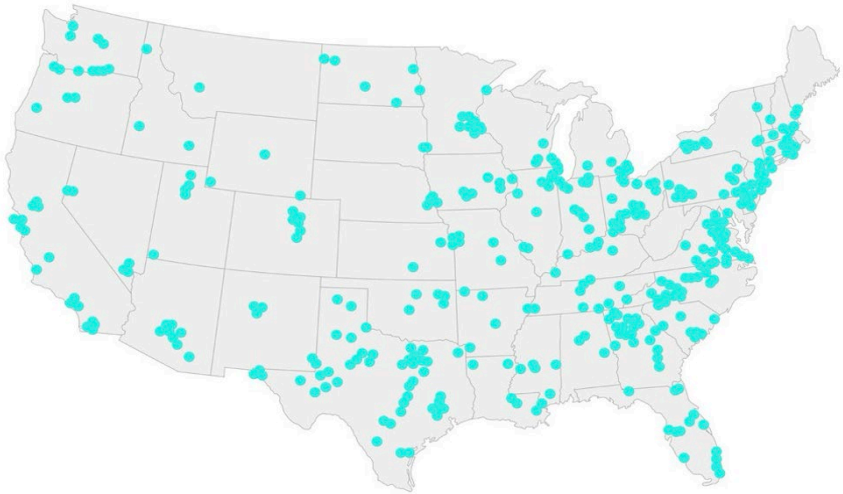

1. Of which ~700M more speculative (from ecosystem primers) 2. Projects may include multiple data center buildings on one campus
Source: Multi-source data prepared by 4MC Partners for Lumen analysis

A New Geography Will Support AI and Cloud Workloads

Three structural shifts

Densification:
Tier 1 hubs (Northern Virginia, Dallas, Phoenix, etc.) are doubling or tripling, driving extreme pressure on power and bandwidth.

Diversification:
Growth is spilling into new AI regions across rural regions (e.g., Western Pennsylvania, North Carolina, West Texas) where latency and scarce dark fiber become critical.

Disaggregation:
The “flat” WAN is breaking apart into three distinct strata – cloud access edge, data center interconnect (DCI), and hyperscale back-end fabrics.

Cloud 2.0 is redrawing the network map, new AI regions will emerge with new corridors

Analysis of AI traffic and medoids

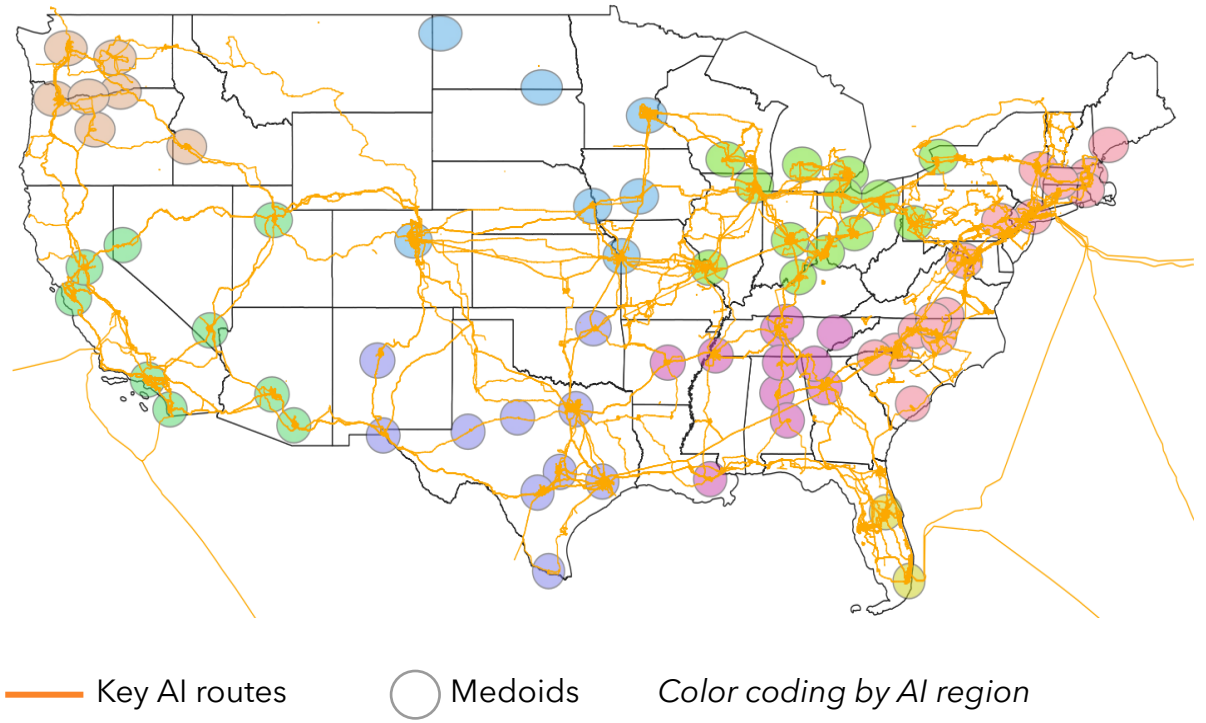

Source: Multi-source data prepared by 4MC Partners for Lumen analysis

Physical Network: the Backbone for the AI Economy

Legend

- Data Center
- Hyperscaler
- Enterprise Premise
- **Fabric Port enabled**
- Point-to-point connectivity (<100G)
- Point-to-point connectivity (100G+)
- Lumen Connectivity Fabric (400G+ enabled)

Exemplar **Cloud 1.0**

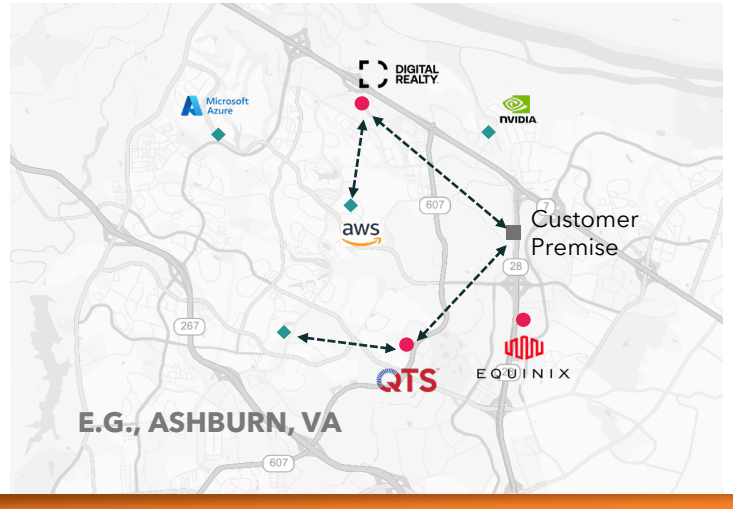

Vision for **Cloud 2.0**

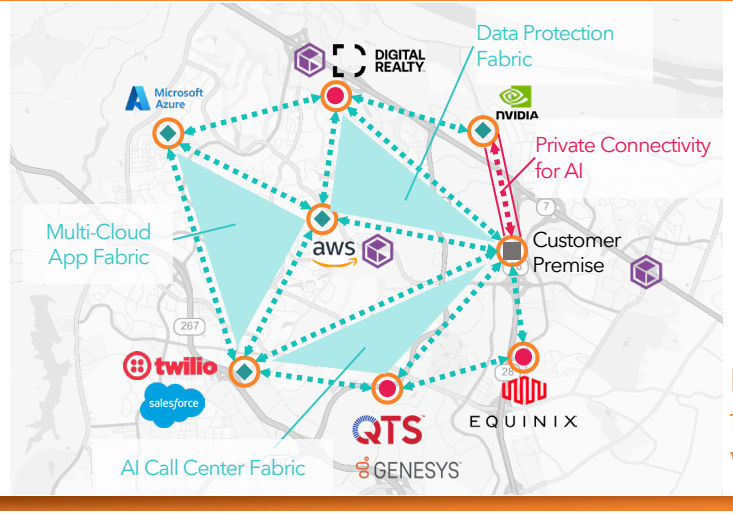

Illustrative, future state vision

	Exemplar **Cloud 1.0**	*Vision for* **Cloud 2.0**
Topology & Speed	**Hub-and-spoke, Tier-1 metros, point-to-point circuits**	**Any-to-any fabric across many metros, distributed data center interconnect (DCI) and virtual meet-me rooms (MMRs)**
Location pattern	**Centralized in Tier-1 metros, rural areas were underserved and far from metro on-ramps**	**Explosive geographic sprawl with new corridors and clusters beyond metros (in rural areas with land and power)**
Scale of data	**Workloads in gigabytes, data moved north-south or user to data center**	**Exabyte-scale transfers between "AI factories", often requiring 400G and 1.6T links**
Meet-me rooms	**Connectivity to carrier-neutral facilities (CNF), public IP access and peering dominated**	**Public cloud direct connect with distributed DCI**

Illustrative future state vision

Enterprises can **design and control** their network without owning, managing and operating assets

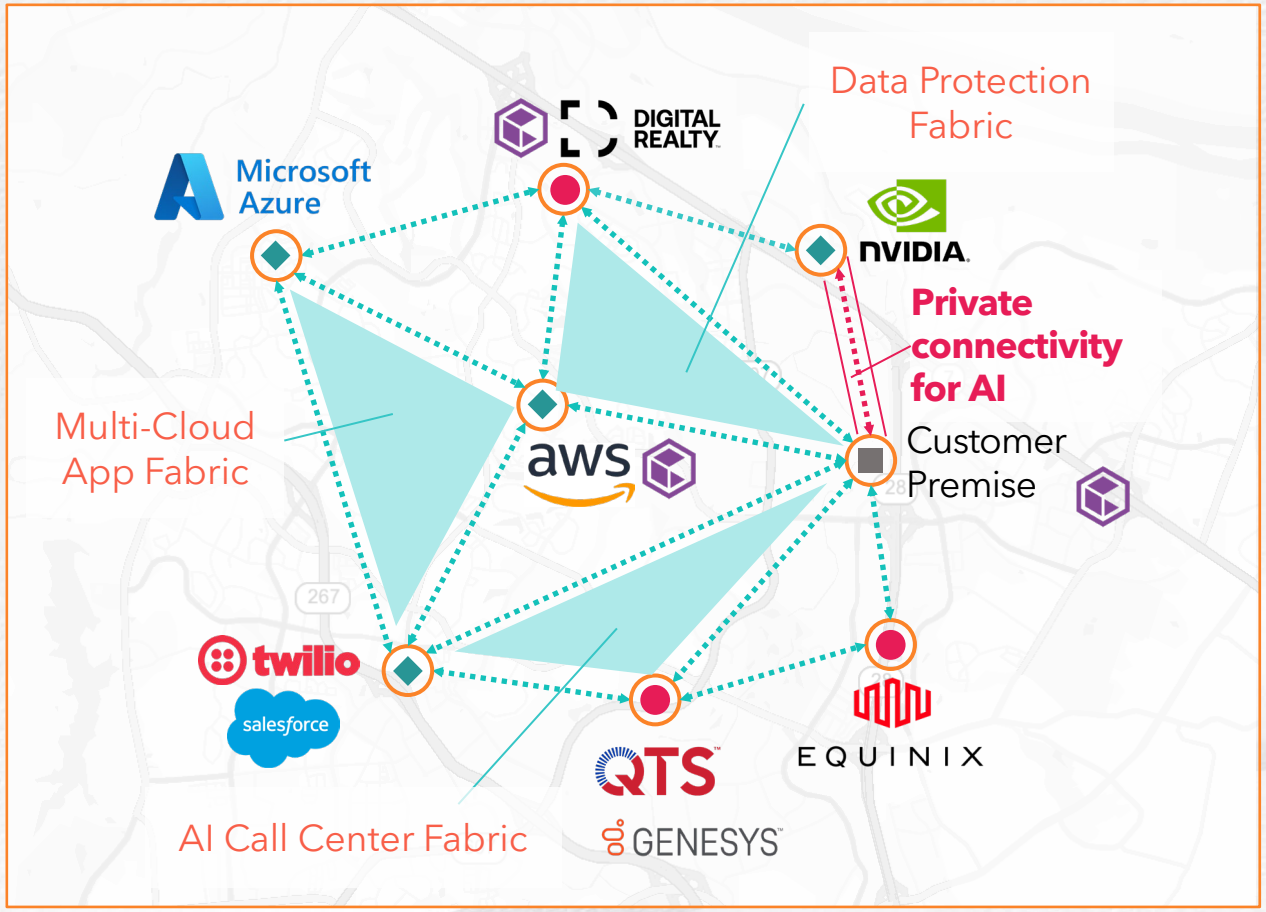

Illustrative, future state vision

Legend: Network components

Data Center · Hyperscaler · Enterprise Premise · Fabric port enabled

Lumen Connectivity Fabric (400G+ enabled) · Private AI connectivity · Fabric

Lumen Drives Significant Benefit for Enterprises

- Programmable fabric connecting clouds, data centers, and AI regions to transform enterprise operations
- High-speed, secure and low latency connectivity allowing firms to participate in the AI economy
- API-led activation across marketplaces, with activation and delivery within minutes, helping firms to save time
- Build-in safe, redundant, and reliable connectivity, reducing business disruptions and IT bottlenecks
- Single Platform to design, price, order, provision and monitor services without owning everything

Enterprises will be defined as leaders or laggards:

Cloud 2.0 Won't Wait

Enterprises are going AI-native: Cloud-first isn't enough, inference, privacy, multimodal workflows demand expansion across multi-cloud. First movers define Industries.

Networks need to change fast Exabyte transfers and real-time inference will overwhelm today's overlays – Cloud 2.0 change will happen 10x faster than Cloud 1.0

Programmable fabrics to be enabled DCI-centric, on-demand connectivity flexes with AI workloads through self-serve, secure, zero-touch platforms

Customer experience must be digitized From static catalogs, manual provisioning to API marketplaces, self-serve fabrics, click-to-buy and zero-touch turn-up


GNE
Global NaaS Event