

# From legacy limits to limitless possibilities

Explore how Cloud 2.0 is redefining enterprise networks in the age of AI



# Contents

The cloud at a crossroads	3
Legacy networks are straining under the weight of AI	4
Cloud 2.0: The network reset	6
Data centers as engines of the AI economy	8
Flexible, programmable networks are the future	10
Unlock the power of Cloud 2.0 for your business	12
Navigating Cloud 2.0 with Lumen	14

# The cloud at a crossroads



For the past 15 years, the cloud has been the backbone of our digital world, relying on centralized data centers and static networks. The truth is, what we now can call Cloud 1.0 wasn't built for modern day needs. With AI booming, data volumes exploding and remote data centers popping up everywhere, the old ways just aren't enough anymore.

Today, both enterprises and tech providers face a new reality. To stay ahead, they need to adapt—now. AI is reshaping industries and redefining what's possible, but the infrastructure supporting this revolution is struggling to keep up. The truth is, Cloud 1.0 wasn't built for the speed, scale and data

intensity that AI demands. It's time for a whole new approach.

Cloud 2.0 is the next generation of cloud infrastructure, designed from the ground up to support the speed, scale and intelligence needs of AI-driven businesses. Unlike legacy cloud models, Cloud 2.0 delivers programmable, distributed and ultra-high-performance connectivity, enabling organizations to innovate and grow in the AI economy.

**Keep reading to learn about the evolution of Cloud 2.0, what it means for your organization and how to adapt your network for tomorrow's demands.**

# Legacy networks are straining under the weight of AI

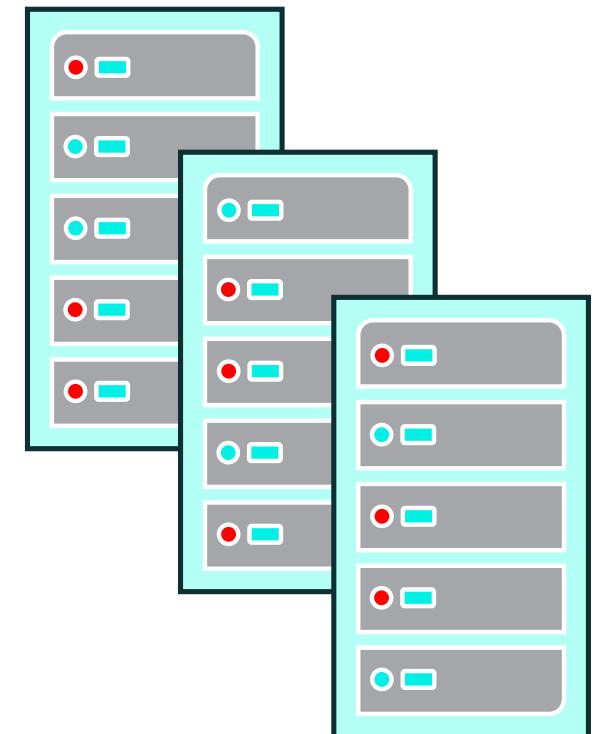
Legacy networks were built for a different era—think voice calls, basic internet and early virtual private networks (VPNs). Today, the demands of AI and modern cloud applications are stretching these old systems to their limits. Static architectures, limited bandwidth and patchwork solutions just can't keep up with the speed, scale and flexibility that businesses need now.

Many organizations tried to adapt by layering VPN tunnels and overlays on top of their existing networks, but these quick fixes only added complexity and made it harder to scale. Centralized

data centers and slow interconnects have become bottlenecks, especially as AI workloads require massive data movement and ultra-low latency.

There's also a common misconception that simply increasing internet bandwidth will solve these problems. In reality, the public internet isn't infinite, and it can't guarantee the performance, security or reliability that AI and Cloud 2.0 require.

**If your business is still relying on patchwork solutions, now's the time to rethink your approach.**



## Here's a look at why legacy networks can't keep up with modern business needs.

Legacy network element	Limitation	Impact on AI/Cloud 2.0
Static MPLS VPN meshes	Low bandwidth, rigid topology	Bottlenecks for data movement
Tunnel-over-internet solutions	Complex overlays, poor interoperability	Unreliable, hard to scale
Centralized data centers	Slow interconnects, single points of failure	Distributed AI is limited
Public internet reliance	No guarantees for latency or redundancy	Unfit for mission-critical AI
Reluctance to change	Skills gap, high migration costs	Delays transformation

**Table 1**

Legacy network limitations and impact on AI

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

# The network reset



Picture your AI workloads moving seamlessly across a programmable, distributed network with no more bottlenecks and no more delays. That's Cloud 2.0: a complete reimaging of how networks, data centers and cloud services work together to support the next big wave of AI digital transformation.

You already know that AI demands a lot from your network: huge amounts of data moving quickly, with minimal delays and high redundancy. For example, the amount of time needed for AI training over a 10 Gbps connection would be nearly 222 hours—almost 9 days.<sup>1</sup>

Moving to Cloud 2.0 is like going from a single-lane highway to a multi-lane expressway where data can flow without bottlenecks. Both represent a transition from simple, linear infrastructure to a complex, interconnected and dynamic system. The single-lane highway is a basic, one-way road, much like early cloud models that relied on a single provider or basic data storage.

The expressway, on the other hand, is a sophisticated network that uses real-time data, AI and automation to adapt and optimize flow. This is analogous to the advanced capabilities of Cloud 2.0, such as hybrid and multi-cloud environments, edge computing and data-driven services.

## See how moving from Cloud 1.0 to Cloud 2.0 can provide the AI-readiness every business needs today.

Element	Cloud 1.0	Cloud 2.0
<b>Structure</b>	Centralized data centers; static, rigid network topologies	Distributed, programmable networks; flexible, stratified layers
<b>Data flow</b>	Slow interconnects; bottlenecks for large-scale data movement	Ultra-high bandwidth; low latency; seamless, high-speed data movement across regions
<b>Optimization</b>	Manual provisioning; patchwork overlays; limited automation	API-driven management; automated provisioning; real-time scaling and resource allocation
<b>Capabilities</b>	Built for voice, basic internet and early VPNs; limited support for AI	Purpose-built for AI workloads; supports massive data, real-time collaboration and innovation

**Table 2**

Comparison of Cloud 1.0 and Cloud 2.0 capabilities

As AI datasets put unprecedented pressure on networks, Cloud 2.0 provides a path to unlock the speed, scale and intelligence businesses need to thrive.

But achieving these breakthroughs is only possible with infrastructure built to support massive movement, ultra-low latency and real-time innovation.

# Data centers as engines of the AI economy

Cloud 2.0 is literally reshaping the physical footprint of digital infrastructure. AI workloads are growing more demanding and data center growth rates are exploding, with new facilities rapidly emerging in suburban and rural regions far beyond the usual tech hubs.

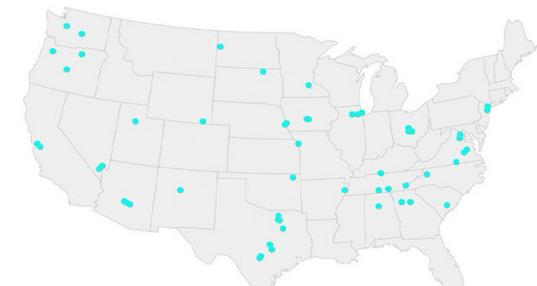
The U.S. data center market is projected to grow nearly **10x by 2030**, with over **1 billion square feet** of new capacity coming online.<sup>1</sup> Such massive growth goes beyond mere scalability, creating a structural shift in where and how digital infrastructure is built.

Tier 1 markets like **Northern Virginia, Phoenix, Chicago** and **Dallas** are densifying, with some regions expected to **double** or **triple** their power consumption. But the real story is diversification. Data centers are moving into adjacent suburban and rural areas, and places like **Reno, Nevada, Western Pennsylvania, North Carolina** and **West Texas** are emerging as major hubs.

What's driving this shift? Metro areas are space constricted and have higher expenses for power and water. As a result, data center operators are turning to new locations, often in more rural environments, where land and energy are cheaper.

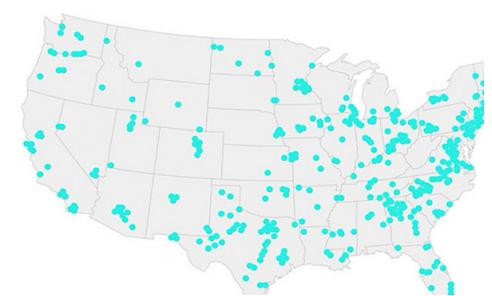
## DC market all segments 2024

*Data center projects<sup>1</sup> in 2024*



## DC market all segments 2030

*Data center projects<sup>1</sup> in 2030*



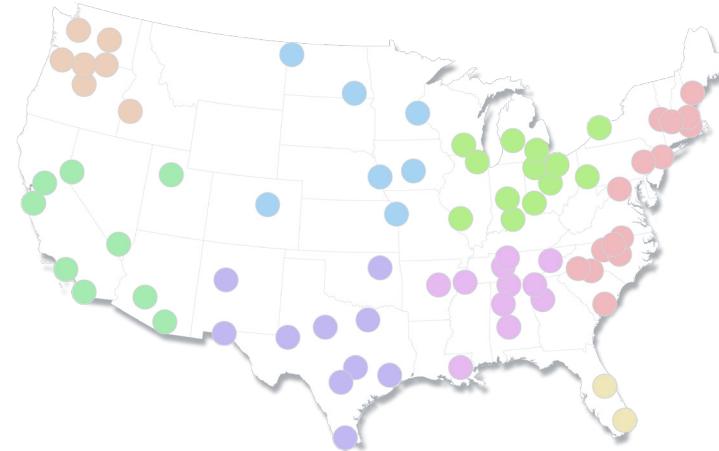
**Figure 1**

U.S. data center projects in 2024 and 2030 by location.

## Disaggregation: Breaking the old model

This new wave of expansion goes beyond simply building more data centers in new places to completely rethinking how everything connects. Instead of relying on one big, central hub to handle all the traffic, Cloud 2.0 is creating lots of smaller, local clusters that work together.

Older facilities that were built decades ago just aren't designed for today's lightning-fast data needs. Cloud 2.0 calls for modern transport hubs and direct connections built specifically for the massive scale and speed that AI demands. With data center operators spreading out and upgrading the way data moves, businesses can unlock new levels of performance, reliability and flexibility.



**Figure 2**

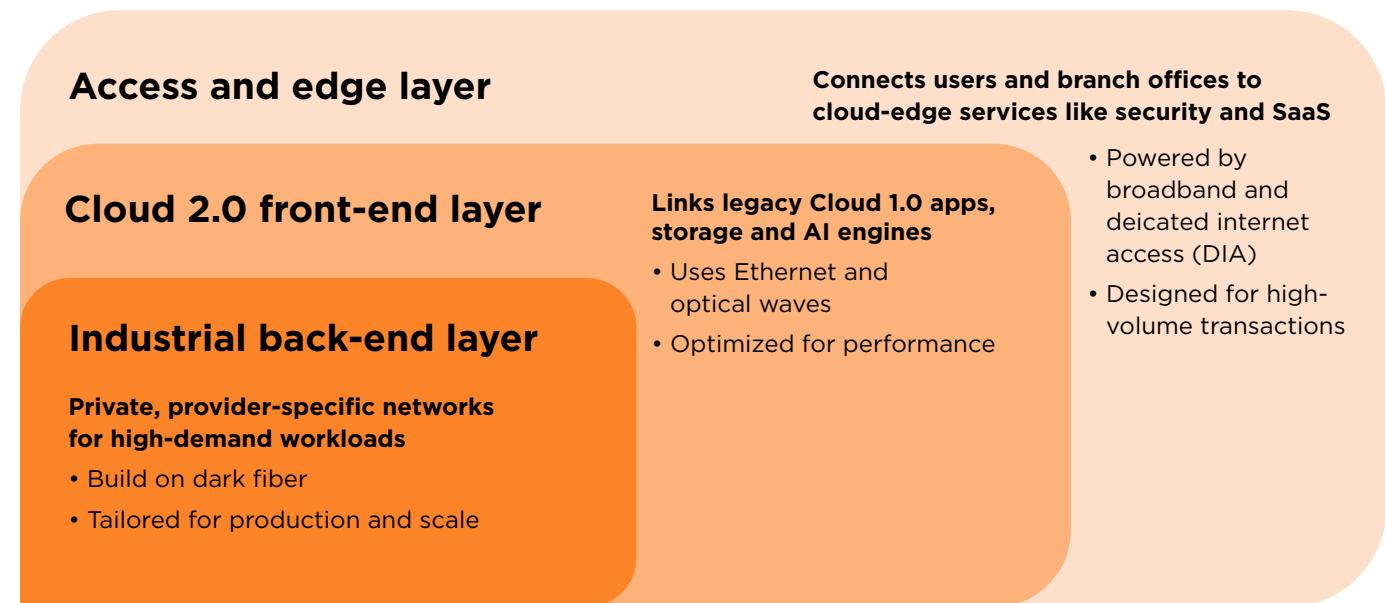
New AI regions are emerging with where resources are available to support massive workloads.

As data centers and network architecture break away from old, centralized models, the next step is clear: businesses need networks that are as adaptable and dynamic as the environments they support.

# Flexible, programmable networks are the future

For years, enterprise networks followed a “flat” model—one big system trying to connect everything the same way. It worked fine for email and web traffic, but AI demands push it past its limits. We can now envision and implement new architecture.

Cloud 2.0 introduces a smarter approach: stratified networks. Instead of one-size-fits-all, it’s now possible to create purpose-built layers:



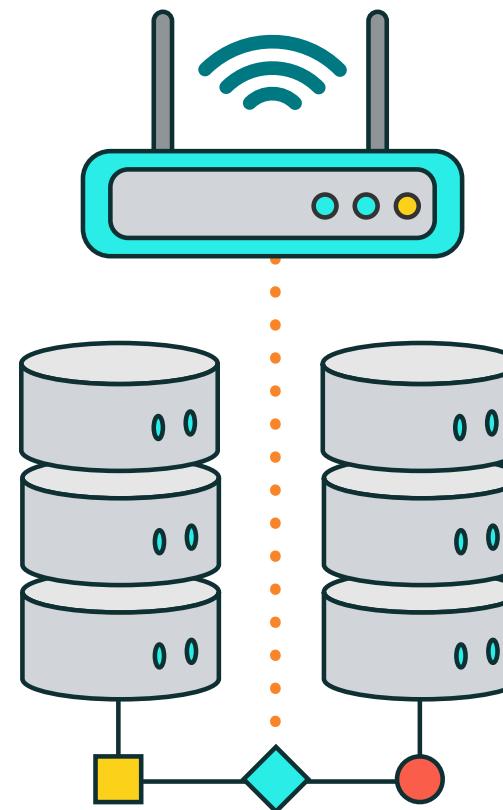
**Figure 3**

The stratified approach to Cloud 2.0 includes three programmable layers, each designed for a specific type of task. This gives enterprises flexibility to align network performance with workload needs.

Each layer is optimized for its job, and the network's programmable architecture lets you instantly adjust bandwidth, latency and redundancy. APIs provide the controls, but it's the flexible, software-driven infrastructure that makes these real-time changes possible.

This shift is both technical and strategic, giving enterprises flexibility to align network performance with workload needs, especially as AI demands grow. And it's a firm departure from today's swamp-like internet, with its lack of bandwidth, redundancy and latency guarantees.

Businesses that plan strategically can take advantage of this growth, ensuring they have the infrastructure needed to scale and innovate.



# Unlocking the power of Cloud 2.0 for your business

As data centers expand and networks evolve, the way businesses operate is changing fast. Cloud 2.0 transforms how and where your business runs, connects and grows. Let's break down the practical implications and steps you can take to stay ahead in the Cloud 2.0 era.

- **Think big when it comes to bandwidth.**

Cloud 2.0 isn't shy about data. To keep up, you'll need connections that can handle massive volumes—think 400G and beyond.

- **Think smart when it comes to programmability.**

Forget static networks. The future is API-driven, letting you scale up and down as needed or reroute traffic with just a few clicks.

- **Go distributed.**

Centralized data centers are out and flexible, everywhere connectivity is in. You don't have to think about where you can connect. Instead, you are free to think about what you need and the infrastructure that follows.

- **Prioritize security and control.**

With more data moving faster, you need precise, real-time control over who accesses what and how that access is provided.

- **Invest in dark fiber and wavelength services.**

These aren't just buzzwords. They're the backbone for high-performance AI workloads.

- **Automate everything you can.**

From provisioning to scaling, automation isn't a luxury anymore. It's a necessity for agility and speed.

## **Fueling a new kind of economy**

Cloud 2.0 represents much more than a technology upgrade. It signals the start of a new era for businesses. As networks become faster, smarter and more distributed, organizations are empowered to innovate at unprecedented speed, thanks to the ability to connect across new markets and uncover fresh opportunities.

Imagine a healthcare provider expanding into rural areas. With Cloud 2.0 infrastructure, they can securely connect remote clinics to AI-powered diagnostic tools, enabling faster patient care and real-time collaboration between specialists, no matter where they're located. This kind of connectivity and

innovation simply wasn't possible with legacy networks, but it's becoming routine in the Cloud 2.0 economy.

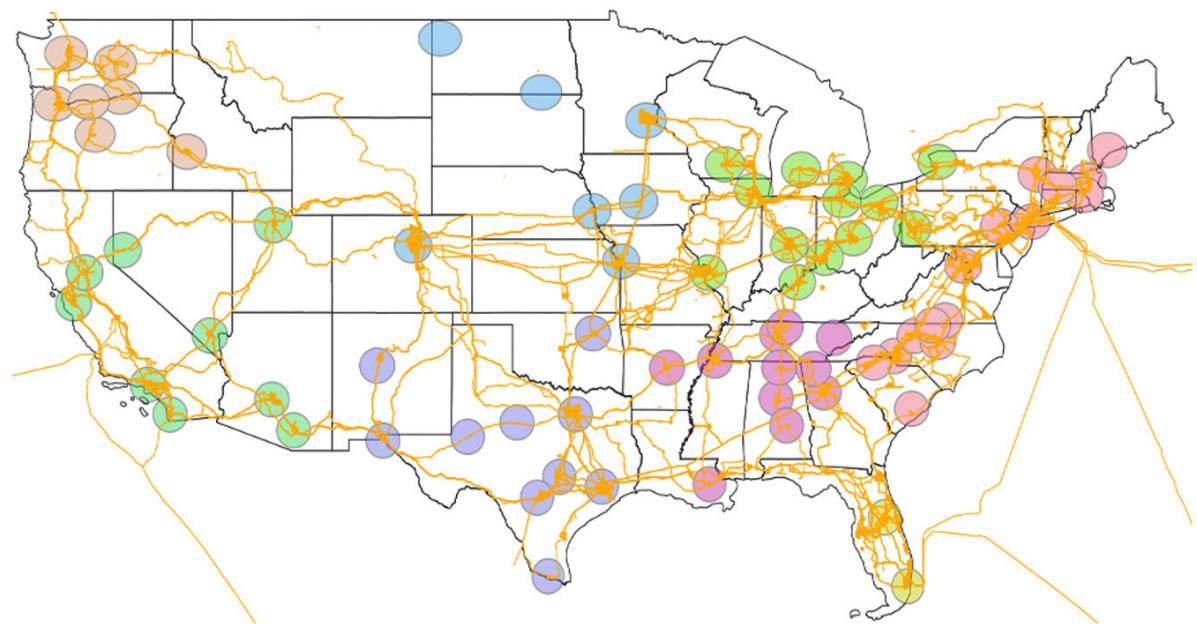
The bottom line? Companies that embrace Cloud 2.0 early will avoid the bottlenecks of legacy systems and set themselves up as leaders in the AI-driven economy, while those who wait risk falling behind.

All of this may seem like a monumental task, but navigating this transformation doesn't have to be overwhelming. By choosing the right partner—one who truly understands the Cloud 2.0 landscape—you're not just keeping pace with change, you're positioning your business to anticipate it.

# Navigate the shift to Cloud 2.0 with Lumen

Cloud 2.0 is the next frontier, and Lumen is your guide.

We're not just talking about faster networks. We're talking about smarter infrastructure, purpose-built for the demands of AI, distributed workloads and real-time innovation. Lumen has invested deeply in expanding one of the world's most advanced fiber networks, and we plan to grow our U.S. intercity footprint strategically **from 17 million to 47 million fiber miles by 2028.**



**Figure 4**

Lumen next-gen fiber builds are planned as part of an overall 3.9x increase in route miles, including eight new AI regions.

But our focus isn't solely on fiber. Flexibility is equally important. Lumen API-first platforms let enterprises design, provision and manage networks in real time. Whether you're spinning up bandwidth for AI training, connecting to a new data center in a rural market or securing edge transactions, Lumen gives you the tools to move fast and stay in control.

We also understand that innovation doesn't happen in isolation. That's why we've built deep integrations with hyperscalers, cloud marketplaces and data center operators, so your workloads can move seamlessly across environments without friction or delay.

In the Cloud 2.0 economy, speed, agility and intelligence are everything. Lumen helps you navigate this shift with confidence, whether you're modernizing legacy infrastructure, scaling AI

operations or building the next generation of digital experiences.

The future of enterprise networking is programmable, distributed and built for AI. And with Lumen, you're not just keeping up—you're ready to lead the way.

**So, what's your next move? Don't let legacy systems hold you back. Start designing your network for Cloud 2.0 today, because the future won't wait, and neither should you. For more insights, read the complete white paper, [Cloud 2.0: Because AI won't run on yesterday's internet.](#)**

<sup>1</sup> Multi-source data prepared by 4MC Partners for Lumen analysis.

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.

## Why Lumen?

As you tackle the challenges and possibilities of Cloud 2.0, Lumen supports you with more than just fast connections—we deliver the strategic edge you need to lead in the AI era. With one of the world's most advanced fiber infrastructures and API-driven network solutions, we make it easy to deploy, scale and secure mission-critical workloads with speed and confidence. Our unified management and direct access to leading cloud providers speed innovation and streamline integration so your organization can stay ahead of the AI-driven curve.

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. ©2025 Lumen Technologies. All Rights Reserved.