The Underappreciated Need to Enable AI and Data Center Growth:

INCREASED AND MORE STRATEGIC FIBER INTERCONNECTIONS

RVA LLC
Market Research & Consulting

Fiber Broadband ASSOCIATION

When fiber leads, the future follows.

I. THE BACKDROP: U.S. DATA CENTER GROWTH – THE UNPRECEDENTED NEED

Data Centers for the storage and computation of data began in earnest in the U.S. in about 2006 as the concept of performing computations and storing data off premises ("The Cloud") began to take shape. More recently, data centers have become synonymous with realizing the promise of generative artificial intelligence or AI.

It is generally understood that for the U.S. to continue to lead unabated in the AI revolution, the country must build data centers at an unprecedented, accelerated pace. **At least three general types of data centers are needed:**

1. Hyperscale training centers for AI development;
2. Hyperscale cloud centers for data storage and computation;
3. Colocation and edge centers for interconnection to users.

This paper focuses on the first two categories related to hyperscale data centers.

Most AI experts call for and forecast at least a 3X increase in cumulative hyperscale data center capacity in just the next few years. (See Figure 1 below.)

Based on data from Data Center Hawk, a company focused on analyzing data center real estate transactions, more capacity than the 3X trend has already been commissioned for the next three years. However, various constraints could prevent such a three-year construction plan from reaching full fruition by the end of 2027.

The number of new hyperscale centers to be constructed annually is forecast to reach 225 by 2032. (See Figure 2 below.) This assumes that hyperscale centers will continue to increase in size over time, as measured in megawatts per data center.

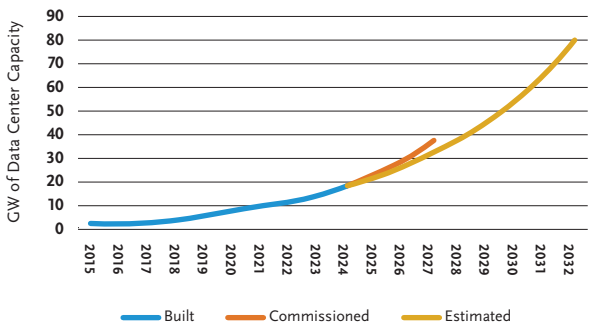

Figure 1:
U.S. Hyperscale Data Center Capacity

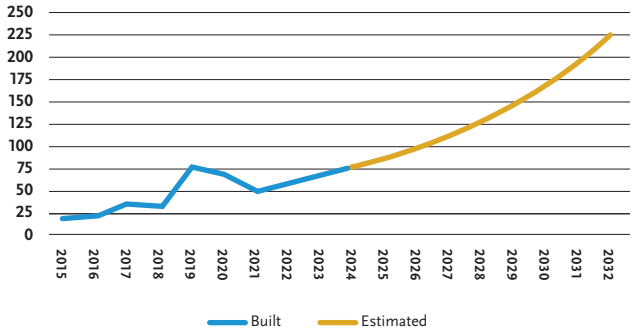

Figure 2:
New U.S. Hyperscale Data Centers Estimated

II. THE OVERALL CONSTRAINTS: THE FORGOTTEN DATA CENTER INTERCONNECTION CHALLENGE

There are many known constraints and challenges to building data centers. These usually focus on four areas: power/electric needs; suitable land for sites; water for cooling; and the availability of advanced microchips and other enabling products.

A fifth and potential critical constraint, seldom discussed, is the fiber-optic interconnection between data centers.

AI could not exist without fiber between data center sites and data centers and users. Given the laws of physics as known today, fiber optic cable is by far the superior type of interconnection — and will remain so far into the future. In a real sense, fiber is the foundation of AI. (Fiber also plays a critical role within the data center — but that is not the subject of this writing.)

Although the need for fiber addition and upgrade is now often understated, the importance of fiber availability and redundancy has been critical since the very beginning of the data center movement. The first U.S. hyperscale data center, opened in 2006 in The Dalles, Oregon, was sited based on both low-cost power from the hydroelectric plants along the Columbia River and the availability of fiber backbone that had been placed along the Columbia by Level 3 and others.

Other early hyperscale center sites such as those in Grant County Washington, opened in 2007, also credited both fiber availability and redundancy in addition to low-cost power. Beginning in 2000, Grant County had started building an extensive network of fiber in the county to support one of the largest experiments in fiber-to-the-home and fiber-to-the-business attempted to that point in time.

Figure 3: Locations of Long-haul Fiber Optic Lines Connecting Data Centers

III. THE INTERCONNECTION CHALLENGE: CAPACITY, SECURITY/ UPTIME, AND LATENCY

There are several components to the interconnection challenge.

A. High Capacity (enough "highway width" to transfer huge amounts of data at once)

For maximum capacity, there should be enough conduits in the interconnection routes, and there should be a high fiber count in each conduit. (New routes are being constructed deploying multiple cables with often fiber counts of 864 or 1,728 strands per cable. Older routes are also being refitted with higher fiber counts.) For the future, new fiber cables have been announced with 3,456 fiber strands in a single cable, and some with 6,912, and even 16,000 strands.

In addition to bundling more strands of fiber into a cable, another option to increase the capacity of fiber pathways is the use of multi-core fiber with several transmission cores within the same strand-of-fiber. Though not a widely used or fully mature technology yet, such an approach even further increases the number of unique optical pathways possible within the same amount of conduit space.

The latest form of dense wave division multiplexing (DWDM) is also usually used in the data center interconnection environment to send multiple wavelengths of light over an individual fiber (or individual fiber cores in the case of multi core fiber) to increase the amount of data that can be transported over each individual fiber.

DWDM capacity for long haul purposes has recently been upgraded 4-6 times, moving from 100 gigabits per second to 400-600 gigabits per second per wavelength. Some providers are already now planning to use advanced DWDM to get as much as 1,200 gigabits per second (1.2 terabits per second) for shorter runs.

While DWDM is an extraordinary tool to increase pathway capacity, it is limited at any point in time by current engineering and cost constraints — and eventually by the laws of physics. Shannon's law says that a maximum is eventually reached where it is not possible to pull usable information out of the noise caused by more channels and more bandwidth per channel.

B. High Security and Uptime (primarily eliminating accidental or malicious fiber cuts)

For maximum security, large data center operators generally now desire three to four redundant routes away from the data center, often from two different fiber providers (if one route is disrupted, there are alternative pathways). Also, wherever possible, additional security is desired, such as placing the fiber conduits deeper in the ground, sometimes encased in steel or concrete. Increasingly, fiber routes are also being protected with fiber optic sensing (a fiber acting as a continuous sensor to alert for potential disruptions, such as backhoe digging in the area, or seismic activity).

C. Low Latency (individual data packets travel quickly from one point to another to avoid delay)

In a perfect world, AI would prefer zero latency between data centers. Of course, zero latency for any method is impossible (excepting quantum entanglement communication, which most think is not possible). However, for those networking methods usable today (copper, wireless, satellite, and optical fiber), fiber has the lowest practical latency (in addition to more bandwidth and more security).

Long-haul fiber providers are constantly developing ways to lower latency further. Finding the most direct routes possible between data centers reduces total distance and thus latency is very important for every provider.

And, if possible, those direct routes should be controlled by a single entity to avoid switch-off points, which add latency. (In some cases, even DWDM is avoided to eliminate the small amount of latency it adds.)

Another method providers are perfecting is the concept of hollow core fiber for long distance transport to eventually improve latency by 20-30%. (In this case, the optical signal would travel in a free space core within the fiber, rather than through the glass itself, thus increasing speeds.)

In the real world, hyperscale data center owners would prefer to have at least one fairly direct and low latency connection from each data center (with about 5 ms of latency), particularly for backup of data. (Data center data is always backed up for redundancy in at least one other data center.) Secondary paths sometimes have 7-10 ms. Unfortunately, finding new, more direct routes for right-of-way is far from easy, and the procurement and construction of such routes is expensive.

IV. THE INTERCONNECTION NEED: HOW MUCH STRATEGIC GROWTH IS REQUIRED

In summary, hyperscale operators would like high bandwidth routes, with low latency (accomplished with the most direct routes possible between locations with limited switching) and very high uptime (accomplished with good physical and cyber security and route redundancy). With the birth of AI, the need for continued improvement in all these categories has accelerated.

Hyperscale data center operators often use the assistance of other fiber operators but would often like to control the interconnection using “dark fiber” (fiber they lease and light themselves) for long distance transport. More preferably, they would place fiber in their own conduit with controlled access, and in some cases, even own and operate their own routes.

The following estimates of future need for interconnects are broken into three categories and are made with the important caveat that exact data for model inputs is not available. Every effort was made to use the best estimates possible.

A. New Short Interconnections/ Laterals To Nearby Access Points

For each new data center built, most operators desire connections in different directions to at least three different Internet connection points (Point of Presence/PoP or Internet Exchange Carrier/IXC). (Note: Hyperscale owners would not connect to the general Internet, but rather, use private "dark fiber" in or near Internet backbone routes.)

To estimate the route miles needed for short Internet connections for new hyperscale data centers, AI was used. (AI, ironically, being the ultimate basis for this paper.) Poisson calculations were utilized to estimate the average miles for short term connections to the Internet backbone/dark fiber.

Based on recent information, RVA estimates there are currently about 5,381 connection points to the Internet in the U.S and an estimated 2,152 unique randomly distributed access points to the Internet backbone (access points estimated to be in the same mile based on data center clustering were eliminated).

Calculations were made using Poisson random distribution analysis. With the fact that there are 3,380,000 square miles in the U.S., the average straight-line distance from any random square mile in the U.S to three different access points was calculated. It was determined that, at present, there are 19.8 miles from a random new hyperscale center to the first access point, 29.8 miles to the second closest access point, and 37.2 miles to the third closest access point. This sums to a total of 86.6 route miles. Then, assuming actual routes would not be in straight lines (because such easements would not be available), we conservatively added 15% to this figure, and one mile for on-campus interconnections, for a total of 101.5 route miles for lateral interconnections per existing data center. This analysis could be conservative for new hyperscale data centers now being built in even more rural areas via ongoing efforts to locate new land and power.

Some long-haul providers have estimated about 50 miles per connection to such new, more rural data centers. At three connections per center, that would place lateral connections at 150 route miles, or 200 miles if four connections per new center. On the other hand, no doubt some new hyperscale centers will be located near others (clustering). Thus, for this paper's analysis, we have used a forecast of 135 route miles of connectivity per new data center.

It should also be noted that In Line Amplifier (ILA) sites are usually required about every 50 miles to amplify/boost the optical signals. Such sites require power, buildings, access, and land area (often over an acre).

B. Upgrading Existing Long-Haul Interconnections

To maintain enough capacity to serve all the new and upgraded data centers, a significant effort to increase the capacity on a number of existing long-haul or intercity routes is also required. This will mean upgrading the fiber optic transport capacity over and above just increasing the DWDM optical capacity that each fiber can transport.

Such capacity increase is often accomplished by increasing the fiber counts on existing routes, as noted earlier. Fiber cables have become more and more 'dense" over the years, with smaller diameter fiber and more fiber packed into the same diameter cable. Therefore, more fiber can be packed into the same conduit size.

Existing routes generally have excess conduit. If a provider has enough such duct, the operator can pull ("overpull") a new high-capacity and low loss fiber through an empty conduit. If not, a process is used to replace existing fiber cables with denser fiber cables within existing conduit. All this requires extensive capital expenditure, of course, and often requires new permissions and permitting (even for work on existing rights-of-way).

C. New Long-Haul Interconnections

In addition to new connections from new data centers and upgraded current routes, there is also a need for new long-haul routes to increase capacity and achieve lower latency. This is true for some routes between existing clusters of hyperscale data centers — but it is even more compelling for connections between new clusters of hyperscale data centers.

A recent Corning presentation shows an estimate of new data center clusters in orange, with the implication of more potential routes. Of course, many of these paths would already be available, but some will require new, or more direct, routes for the best performance. (See Figure 4 below.)

Constructing new routes is often a very difficult challenge — finding continuous rights-of-way, dealing with permitting and, of course, funding and actualizing the construction. Developing such routes is further complicated by the need for reasonable access (periodic entry points to the conduit, with appropriate means to get to these points).

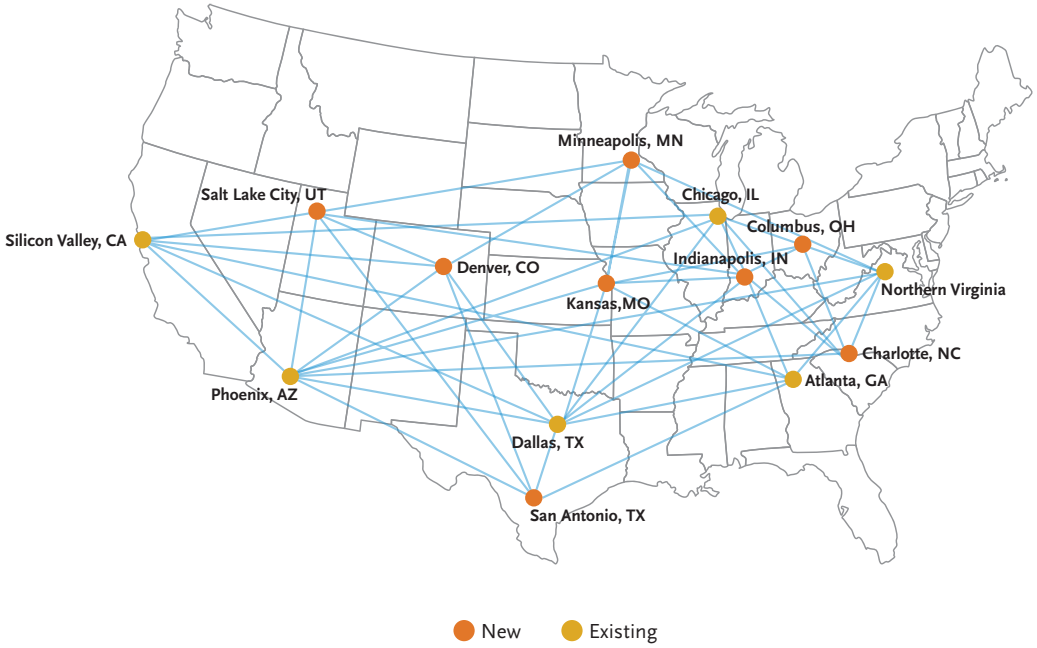

Figure 4: Potential Data Center Connections in United States

Corning Inc., Investor Event March 18, 2025

D. The Resulting Need and Forecast

What do all these types of potential growth mean in terms of the overall need for upgraded and new interconnect capacity? Such a need is usually considered in terms of route miles (distance miles from point to point on a pathway — though these paths are often covered by multiple carriers, multiple conduits and multiple fiber strands per conduit). Another relevant measure is fiber miles (the sum in miles of all the individual fiber optic strands within these paths).

While exact data for long haul and data center interconnect routes is not available, RVA estimates there are currently a total of 95 thousand such unique long-haul route miles and 159 million long-haul fiber miles in the U.S. (See Appendix.)

As discussed, there is a clear need to rapidly expand both route miles and fiber miles to meet the new needs of this new world of AI. It is expected that route miles will increase about 2X from 95 thousand route miles to about 187 thousand route miles by 2029. This would mean that route miles increase by 92 thousand miles (92,000) over the forecast period.

Likewise, RVA estimates the number of needed fiber miles (miles including multiple strands of fiber) will more than double (2.3X) from 159 million miles in 2024 to 373 million miles by 2029. This would mean that long haul fiber miles increase by 214 million miles (214,000,000) over the period. (See Figure 5 below.)

This estimate of additional fiber miles includes 109 million new fiber miles for upgrades on existing routes, 66 million new fiber miles for connections to new data centers, and 38 million new fiber miles for new long-haul routes. (See Appendix.) It is anticipated that the annual new addition of fiber miles will increase from about 30 million in 2025 to nearly 60 million in 2029. (See Figure 6 below.)

Many types of providers and construction companies are building this infrastructure. While many of the providers involved are large scale providers such as Lumen, AT&T, and Verizon, even Tier 3 telecom companies are sometimes being called upon in part because of their detailed knowledge of local power availability, easements, land ownership etc. "The rise of AI and cloud-based services is reshaping data center architecture, blending hyperscaler infrastructure with distributed edge facilities to support real-time demands," said Justin Suhr, Sr. Director of Construction and Engineering at Great Plains Communications. "This shift is accelerating hyperscaler growth and driving an increased need for high-capacity fiber networks. To meet these needs, providers must respond by increasing bandwidth and expanding connectivity and interconnection solutions."

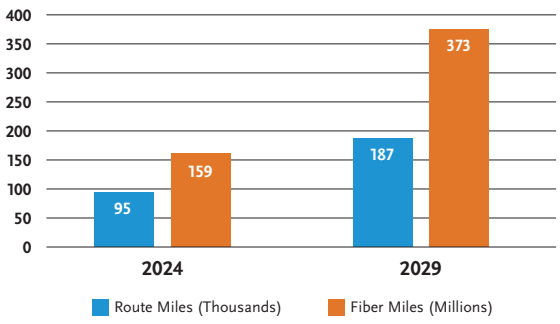

Figure 5:
U.S. Long Haul/Data Center Interconnect Changes

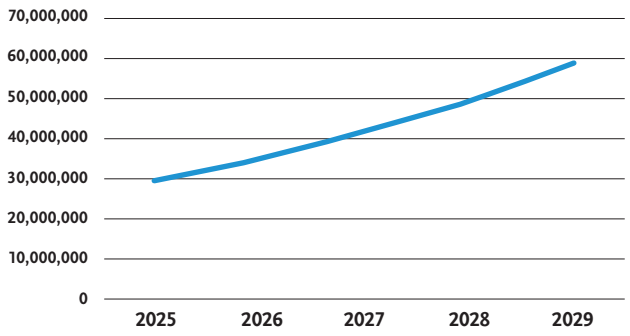

Figure 6:
Approximate Annual Growth of Fiber Miles

V. KEY FINDINGS AND RECOMMENDATIONS

In many ways, the fiber optic backbone interconnection fabric is the backbone of AI. This paper is designed to provide more visibility into the somewhat undisclosed and mysterious needs and challenges of making the "Cloud" and "AI" function and thrive.

The need for high quality, low latency fiber optic interconnects for data centers is critical for the successful AI vision to reach fruition. This will require upgrading and increasing fiber optic interconnections between hyperscale data centers, with the required capacity, latency, and security/redundancy. While exact numbers are difficult to project, this paper forecasts an estimated 2X increase in fiber route-miles construction is needed by 2029 alone, with a 2.3X increase in fiber miles in the same period. Meeting this need will be far from easy.

For existing routes, success will take sufficient capital, permitting relief, construction resources, etc. For new routes, success will take finding new rights-of-way, sufficient capital, permitting relief, and construction assets.

In March 2025, the Fiber Broadband Association submitted a separate filing Request for Information (RFI) to the National Science Foundation suggesting policy actions to be taken to support this effort (highlighting needs for permitting relief, capital gains treatment, and access to public lands). Providers who are heavily involved in this space (such as Lumen) have also submitted RFIs making similar points.

In addition to policy relief, new creativity must be used to find new routes. For example, power and fiber providers should work collaboratively to find options for joint easements. Undergrounding DC power results in narrower and less obtrusive co-easements, which would allow both fiber and power infrastructure. Such easements may also allow developing partnerships for the use of fiber sensing to improve security for multiple assets in the right of way.

VI. ADDENDUM: FIBER'S OVERALL ROLE IN THE AI REVOLUTION

While every new technology is accompanied by some concerns and uncertainty, nearly all would agree there is much to be excited about regarding Artificial Intelligence and the impact it could have on nearly every aspect of human life. Fiber optic networks are indeed critical to the advancement of this vision — at both ends of the network.

At the front end, as discussed extensively in this paper, fiber is central to enabling the increased data storage and deep computations necessary for both today's "cloud" and tomorrow's "AI".

At the back end of the network, fiber will continue to bring paradigm shifting technologies to individuals across the economic and geographic spectrum. In the past two decades, the Internet, upgraded with more fiber in the backbone and, increasingly in the last mile, assisted the world in getting through COVID by enabling increased remote work, remote telemedicine, and advanced communication. Now, as the age of AI dawns, fiber is already enabling the integration of its most promising applications in both urban and rural homes. To improve telemedicine, AI will be soon analyzing video selfies and denser data generated from portable medical devices in the home. A new age of technology designed to lengthen aging in place while protecting personal privacy will be enabled. Better home security capable of identifying human threats on property is imminent ... and much more.

APPENDIX:

ESTIMATES OF U.S. LONG HAUL FIBER STATUS

	MULTIPLIER	ESTIMATED U.S. MILES	
ESTIMATED 2024 STATUS			
Unique Routes Miles		95,000	
Carrier Miles (Assumes multiple carriers per route mile)	3.5	332,500	
Conduit Miles (Assumes multiple conduit per carrier)	3	997,500	
Occupied Conduit Miles (Conduit with fiber per carrier)	40%	399,000	
2024 Average Fiber Miles (Average fiber strands x occupied conduit)	400	**159,600,000**	
UPGRADED CURRENT ROUTES 2025-2029			
Conduit Newly Occupied	4%	39,900	
New Fiber Miles in Newly Occupied Conduit	1728	68,947,200	
Existing Conduit Miles Upgraded with Higher Fiber Count	3%	29,925	
Net Additional Fiber Miles from Upgrades (Old out new in)	1328	39,740,400	
Total New Fiber Miles from Upgraded Routes		**108,687,600**	
NEW LATERALS TO NEW HYPERSCALE DATA CENTERS 2025-2029			
New Data Centers	573		
New Unique Route Miles	135	77,355	
Carrier Miles (Possible multiple carriers per route mile)	1	77,355	
Conduit Miles (Assumes multiple conduit per carrier)	3	232,065	
Occupied Conduit Miles (Conduit with fiber per carrier)	33%	76,581	
Fiber Miles (Fiber strands x occupied conduit)	864	**66,166,373**	
NEW LONG HAUL ROUTES 2025-2029			
New Route Miles Added	1	15,000	
Carrier Miles (Multiple carriers per route mile)	1.5	22,500	
Conduit Miles (Multiple conduit per carrier)	3	67,500	
Occupied Conduit Miles (Conduit with fiber per carrier)	33%	22,275	
Average Fiber Miles (New fiber strands x occupied conduit)	1728	**38,491,200**	
TOTAL NEW MILES FROM UPGRADED & NEW ROUTES 2025-2029			
New Estimated Unique Route Miles		92,355	
New Estimated Fiber Miles		**213,345,173**	
ESTIMATED 2029 STATUS			**CHANGE VS 2024**
2029 Estimated Unique Route Miles		187,355	**197%**
2029 Estimated Fiber Miles		372,945,173	**234%**