

## AI Acceleration and Fiber Optics

<https://fidepartners.com/>

Tel: +34 910 244 113

Cel: +1 617 359 8114

info@fidepartners.com

**April 2025**

AI impact on fiber network demand in the US

**Authors:** Jim Andrew, Nicolás Meneses, Lorenzo Morelli



### Inside

#### Page 03

AI adoption and fiber networks as required technology

#### Page 04

AI impacting directly the fiber business

#### Page 06

Growth perspectives for the fiber market and incentives for stakeholders

**AI solutions are substantially driving the demand for high-bandwidth, low-latency, and reliable connectivity**

As artificial intelligence (AI) rewrites the rules of modern computing, it is expected to create unprecedented demand for robust connectivity infrastructure. With training concentrated in massive data centers and inference workloads moving closer to users, fiber networks are emerging as the critical foundation enabling AI's expansion.

Despite lagging in residential fiber coverage, the US market stands at a pivotal inflection point. Fide Partners projects that AI's demand for bandwidth and ultra-low latency will accelerate fiber deployment nationwide, creating strategic opportunities for networks and investors.

# AI Acceleration: how advanced computing is reshaping the fiber infrastructure landscape

## Introduction

---

The AI revolution is rapidly transforming from an emerging technology to an omnipresent force across industries. What began as groundbreaking search tools and text chatbots has quickly evolved into a powerful lever fundamentally changing daily operations for individuals and corporations alike.

In the US market, major players like OpenAI, Google, Oracle, Meta, and Anthropic have demonstrated that delivering advanced AI's unprecedented capabilities demands significant infrastructure resources, including massive data centers and high bandwidth connectivity. Recently, competitors like DeepSeek have disrupted market dynamics by achieving results comparable to the most advanced models rapidly and at lower cost. However, even these "efficient" models still require substantial infrastructure support for their operation and continued development. While we expect future efficiency gains, the expansion of potential use cases will continue to drive infrastructure needs.

What remains clear is that regardless of how competition among AI players evolves in the short and medium term, robust infrastructure foundations will be essential to support the growth that the market demands.

The infrastructure requirements for AI are creating a profound shift in data center design and deployment. We are witnessing an evolution from megawatt to gigawatt-scale facilities while traditional air cooling is being replaced by liquid cooling techniques as the power densities within a rack explode.

This paradigm shift extends beyond just hyperscale data centers. The emergence of AI inference workloads (the deployment of trained models to respond to real-time queries) is creating new demand patterns across distributed locations. As organizations seek to reduce latency and optimize costs, robust connectivity is needed between distributed computing resources.

Several industries stand to benefit from AI's accelerated development, with fiber optic networks being a particularly compelling example. This whitepaper explores how AI adoption, computing requirements, and the unique positioning of the US market are creating significant opportunities for fiber network development in the coming years.

### Significant AI players rapidly entering the US markets



Founded in 2015, launched GPT-4 in 2023, which required ~16,000-25,000 GPUs for training



Major AI initiatives since 2016, with significant expansion in 2021-2023



Founded in 2021, launched Claude models in 2022-2023



Emerged in late 2024 as a disruptive Chinese AI lab claiming similar performance to leading models at much lower training costs



As of 2024, Grok, from xAI, used 100,000 H100 GPUs for training and was developed in 122 days. It is said that they will double its GPU size to 200,000 in the short term

### AI players driving infrastructure evolution

---

The emergence of AI pioneers has transformed data center requirements globally. From OpenAI's massive GPU clusters to DeepSeek's efficiency innovations, each advancement has accelerated infrastructure evolution.

This rapid growth has driven data center capacity from tens of megawatts to multi-gigawatt projections by 2030. This unprecedented scaling creates strategic challenges for AI providers, forcing trade-offs between capital efficiency and performance optimization.

The resulting architecture demands not only greater power capacity, but an entirely new approach to deployment strategy; balancing computational density, geographical distribution, and network performance to support AI's increasingly complex workload patterns.

# AI adoption and fiber networks as a required technological solution

Organizations worldwide are increasingly integrating AI into their operations, leveraging its capacities to maximize efficiency, optimize decision-making and drive innovation. This expansion of AI adoption is not only transforming businesses but also increasing the demand for network infrastructure, which is essential for supporting AI-driven data processing connectivity. While different technologies such as coaxial cables and satellites are competing to keep up with AI adoption, fiber remains the best technology due to its low latency, high data capacity, and complete immunity to electromagnetic interference, unlike coaxial cables, which, despite their shielding, can still be affected in extreme environments.

Key industries such as healthcare, education, finance, manufacturing, marketing, and cybersecurity have significantly increased their adoption of AI in their business functions in the past few years. The finance sector has embraced AI for fintech solutions and banking automation, while healthcare is finding opportunities to leverage it to enhance diagnostics, personalized treatments and improve operational efficiency. Similarly, cybersecurity has seen advancements in the response time requirements for threat detection. In addition, the sales and marketing sectors seek ever-improving customer service experience.

As AI adoption grows, its applications vary widely, and so do their latency requirements. Some, like remote surgery and autonomous vehicles, demand ultra-low latency for real-time decision-making. AI in manufacturing also spans a range of latency needs. Robotic automation and real-time defect detection require minimal delays, while process optimization and supply chain forecasting operate on longer timeframes. Given these diverse needs, having a robust and low-latency infrastructure is critical for ensuring AI's optimal performance.

**Latency requirements according to AI use**












AI Application	Latency requirement	AI Application	Latency requirement
Telemedicine and remote surgery	Ultra low	Manufacturing	Ultra low to moderate
Autonomous vehicles	Ultra low	Customer Support	Moderate
Fraud detection and cybersecurity	Ultra low to moderate	Marketing	Moderate
Gaming	Low	Big Data Analytics	High

Source: Fide Partners analysis and research

Furthermore, AI applications require massive data processing that implies high speed, low latency networks, placing significant demands on network infrastructure. Despite the growing need for robust network infrastructure, the fiber optics industry faces competition from other technologies such as coaxial cables and satellite internet. When comparing these technologies, it is evident that fiber optics is the superior solution for keeping up with AI adoption.

As the figure below shows, fiber offers the lowest latency, meaning that its time delay in data transmission ranges between 5 and 20 ms (milliseconds), whilst coaxial cable latency speed can reach up to 30 ms and satellites can exceed hundreds of ms due to the vast distances involved.

**Comparison table by network technology**

NOT EXHAUSTIVE		Speed	Latency	Service quality
		Upload / Download	milliseconds	Highest ●
	<b>Fiber</b>	 <b>1 Gbps – +100Gbps</b>	5-20	
	<b>Coaxial (DOCSIS 4.0)</b>	 <b>≤ 10 Gbps</b>  <b>≤ 6 Gbps</b>	10-30	
	<b>Satellite (Starlink)</b>	 <b>≤ 150 Mbps</b>  <b>≤ 30 Mbps</b>	20-500	

Additionally, fiber optics provides superior speed, featuring symmetrical multiple Gbps upload/download capabilities, and is resistant to electromagnetic interference, unlike cable.

As AI continues to reshape industries, the demand for fast network infrastructure will only grow. As a result, fiber optics availability will be required as the superior solution, offering low latency, unmatched speed and resilience against interference.

Sources: Fide Partners analysis and research



# AI architecture impacting the fiber business

The explosive growth in AI adoption is generating demand for both training and inference capabilities across the US. The strategic shift of inference workloads toward Tier 2 and Tier 3 data centers creates a ripple effect across fiber connectivity, driving demand for high-capacity, ultra-low latency connections between increasingly distributed computing resources.

Understanding the distinction between training and inference is crucial. Training involves developing AI models through massive datasets, requiring substantial computing power concentrated in large data centers. Inference refers to deploying trained models to respond to real-time queries. While less computationally intensive, inference still requires significant processing power as applications grow more sophisticated.

AI model training necessitates robust connectivity between hyperscale data centers, with bandwidth requirements growing exponentially. DeepSeek's recent breakthrough (detailed in our article "The DeepSeek Moment") demonstrates how advanced architecture can improve efficiency and increase demand for fiber network capacity between foundational training models and inference applications.

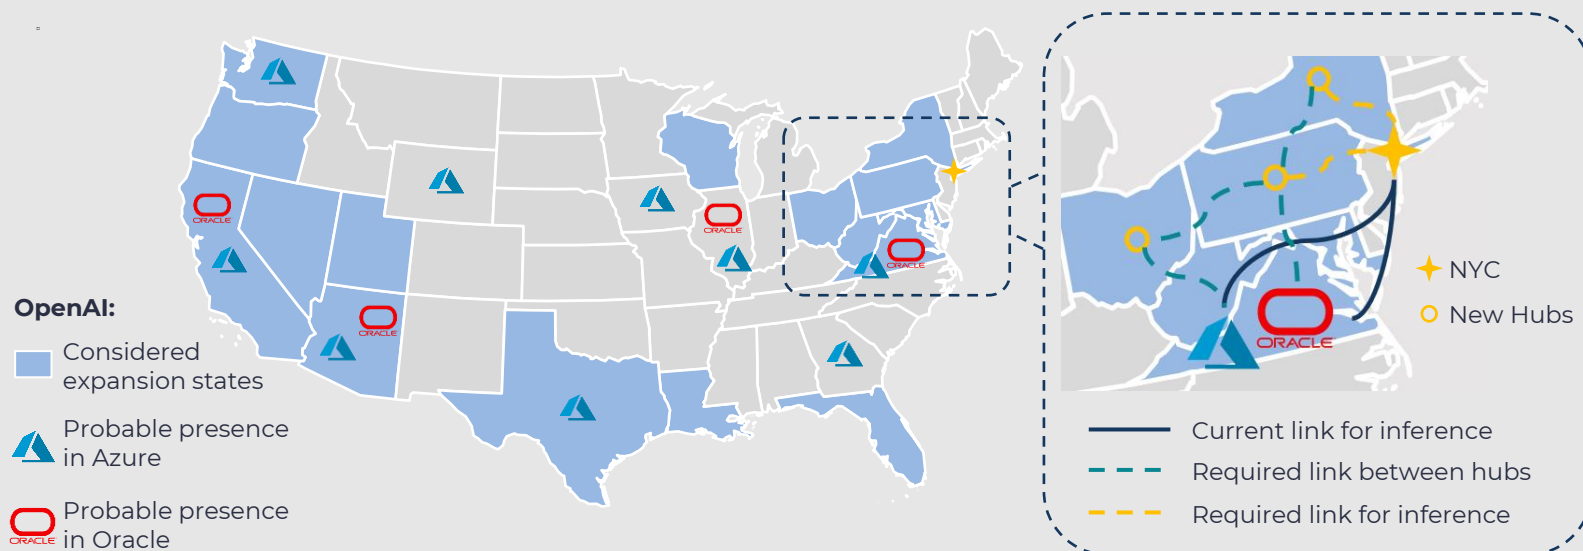
Training processes require massive backhaul capacity, high-bandwidth metro networks, may require low latency connectivity, and redundant fiber routes.

Inference workloads are increasingly migrating from centralized locations to edge-proximate data centers, driving demand for Fiber-to-the-Data-Center (FTTDC) and Fiber-to-the-Premises (FTTP) B2B services. As stated before, healthcare applications may require sub-millisecond latency for image-guided interventions and remote surgery, and financial services depend on real-time inference for fraud detection and algorithmic trading as two examples of many..

Most critically, inference instances maintain constant communication with training models, transferring user prompts, responses, and feedback data that continuously improves model performance. This creates a persistent need for high-capacity fiber connectivity between edge inference locations and central training facilities.

As AI capabilities continue to advance, the underlying fiber infrastructure must evolve in parallel. The distributed nature of modern AI architectures creates a dynamic ecosystem that relies on robust fiber connectivity. This symbiotic relationship between AI advancement and fiber deployment will shape the digital landscape for years to come, presenting significant opportunities for infrastructure providers who can strategically address these emerging connectivity needs.

## Case Study: OpenAI's planned expansion in light of the Stargate project



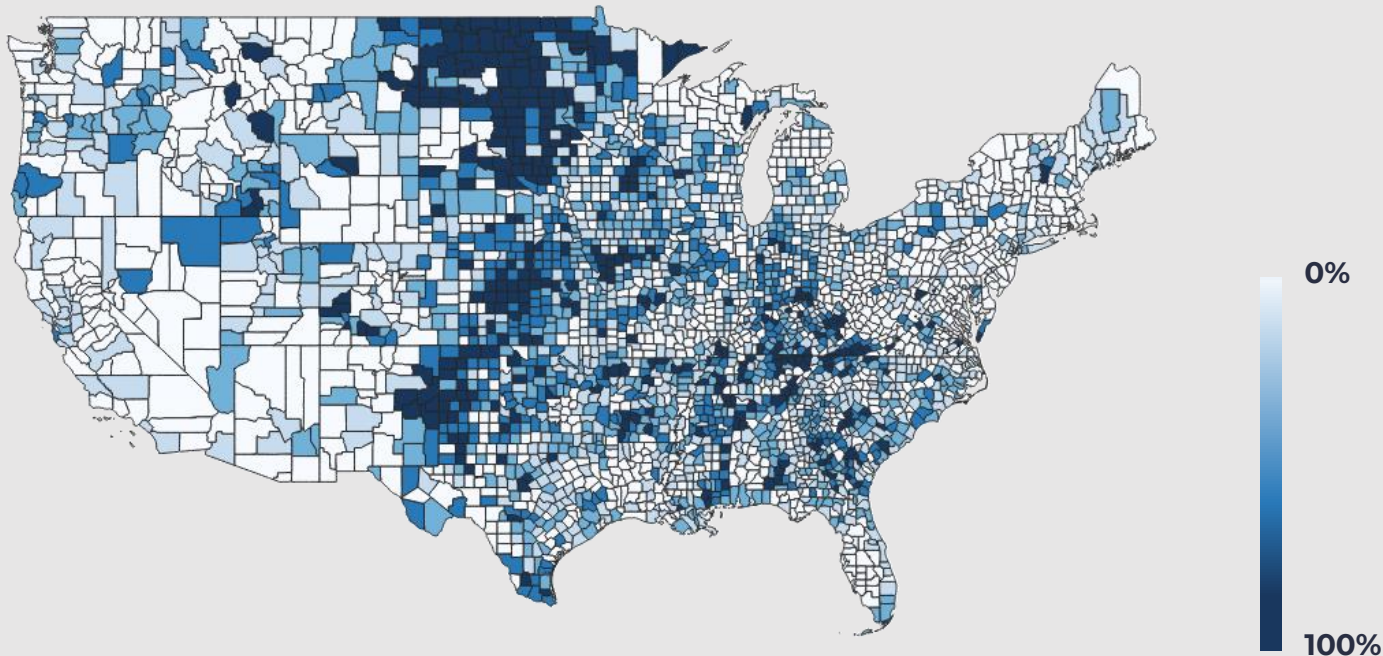
OpenAI's Stargate initiative represents a prime example of this strategic shift. The company is evaluating 16 states for new data center campuses, with construction already underway in Abilene, Texas.

This expansion exemplifies the evolving fiber requirements for AI infrastructure, with OpenAI creating a nationwide network of interconnected AI facilities. The initiative will necessitate substantial fiber deployment between these new facilities and existing hyperscale data centers, illustrating the symbiotic relationship between AI advancement and fiber infrastructure expansion. Our NYC illustration shows how these secondary hubs will directly impact the requirement of new metro and backhaul-capable networks.

Sources: Companies' websites, press releases, Fide Partners research and analysis

# Enterprise fiber market outlook in the US

**B2B fiber connections over total FBB connections in the US**  
 June 2024, % of total connections



Sources: Fide Partners research and analysis, FCC

The US enterprise fiber market shows significant growth potential, particularly in the data center interconnect space, given the low penetration in relevant regions. Connectivity demand, driven by the rise of AI, may accelerate fiber expansion across the country. In fact, investors have shown appetite for this growing trend as it has been reflected in the growing dark fiber market.

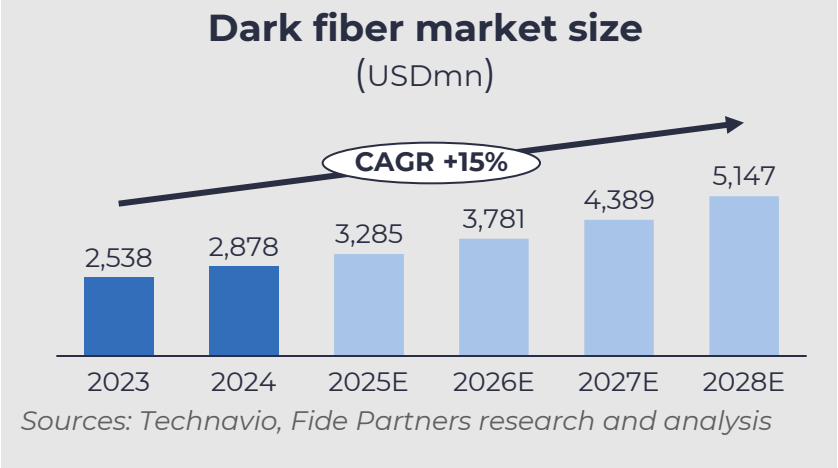
Current market analysis indicates that only about 460 counties in the US (~14%) have over 75% fiber B2B connections. This uneven deployment landscape is particularly evident in business districts across the Northeast, parts of California, and metropolitan areas in Texas and Florida.

On the other hand, with the growing adoption of AI solutions, businesses will require symmetrical speeds exceeding 1 Gbps to support inference and training model demands. This will drive fiber penetration across key markets, representing a significant opportunity for providers targeting this customer cluster.

As a result, the market is witnessing growing institutional interest in enterprise fiber infrastructure, particularly focusing on data center interconnectivity within major metropolitan areas and primary colocation hubs. These investments predominantly target dark fiber assets and advanced connectivity solutions such as multi-gigabit wavelength services and carrier-grade Ethernet.

In line with those investment trends, there has been a significant expansion of dark fiber networks in the US designed to meet the high-capacity, low-latency needs of AI-driven enterprises and data centers, as seen in the graph below.

Some players like Summit IG, DF&I and Bandwidth IG, which specialize in deploying dark fiber networks for FTTDC and high-speed data center interconnectivity, are addressing specific data center intensive markets with their network solutions.




We also expect an increase in investments in backbone and metro networks for enterprise end customers in the near term. The convergence of AI-driven demand and the competitive advantage of fiber's symmetrical capabilities positions the enterprise fiber market for significant growth.

# Growth perspectives for the fiber market and incentives for stakeholders

The mentioned explosive growth of AI applications and the current positioning of fiber networks in the US indicate substantial market expansion in the coming years. As data traffic volumes surge and new inference hubs emerge, the fiber market is poised for significant growth across multiple dimensions. This growth trajectory is further supported by government initiatives and evolving cost structures that create compelling incentives for industry stakeholders.

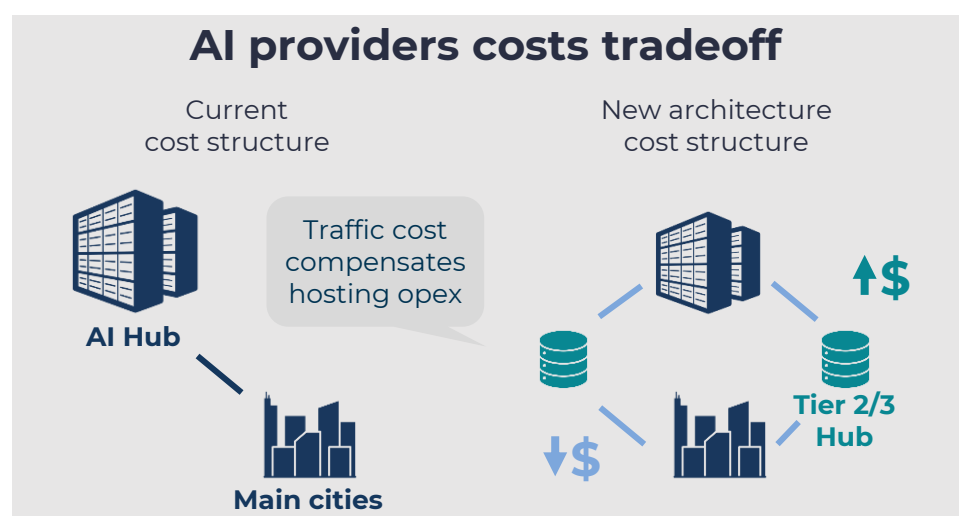
Data traffic volumes have reached unprecedented levels, with total internet traffic in the US growing at a CAGR of ~21% since 2020. While recent reports suggest that overall broadband data consumption has grown more steadily, AI-driven applications are expected to be a key driver of future network traffic expansion. Our estimates suggest that the market should expect to almost triple current levels by 2030, with much of this growth anticipated to come from AI-related data transfers between training and inference locations. This will require the high bandwidth and low latency that only fiber can reliably provide, ultimately increasing demand for backbone and metro fiber networks across the country.



For fiber network operators, this evolution presents significant opportunities. The BEAD program, with its \$42.5 billion allocation for broadband infrastructure, represents just one of several federal and state government initiatives promoting expanded connectivity. While primarily targeting consumer access, these programs establish the foundational infrastructure that can be leveraged for enterprise connections and data center connectivity, creating economies of scope for providers who strategically align their deployment plans.

On the other hand, AI providers face a strategic tradeoff that benefits the fiber market: Deploy inference capacity in secondary markets or bear the cost of transporting massive data volumes to centralized hubs. As data volumes grow, distributed deployment becomes increasingly economical. While establishing facilities in Tier 2/3 markets requires investment, these locations offer lower power and real estate costs, creating a favorable cost structure. This economic reality accelerates fiber deployment and creates strategic growth opportunities for providers who can optimize routes between distributed AI infrastructure.

With data traffic growing exponentially, new inference hubs emerging in secondary markets, government support programs providing foundational funding, and the inherent advantages of fiber for AI workloads becoming increasingly apparent, the US fiber market is well-positioned for growth over the next five years. This convergence of factors creates a uniquely favorable environment for strategic investments in fiber infrastructure to support the next generation of AI-driven innovation.





# Conclusions: AI and fiber networks in a symbiotic bond for future growth

The convergence of AI advancement and fiber infrastructure development represents more than just parallel growth trajectories: it signifies a fundamental reshaping of the US digital backbone. As inference workloads migrate to secondary markets and edge locations, we are witnessing the birth of a new distributed computing architecture that demands high-performance connectivity.

This evolution is driving measurable network traffic growth, with data volumes expected to triple by 2030. The fiber networks being deployed today aren't merely supporting general connectivity needs but specifically addressing the unique AI workloads that require high-bandwidth, ultra-low latency connections.

For stakeholders across the ecosystem, from fiber providers and data center operators to AI developers and enterprise customers, this moment presents a rare alignment of technological requirements, economic incentives, and policy support.

As we look toward 2030, the most successful organizations will be those that approach fiber deployment not as isolated infrastructure projects, but as integral components of a cohesive AI strategy that bridges centralized and distributed computing resources.



## Direction:

### Boston:

50 Milk Street,  
15th Floor.  
Boston, MA  
02109

### London:

Aldwych House  
London, WC2B 4HN  
United Kingdom

### Madrid:

C/Don Ramón de la Cruz, 6, 1º  
28001 - Madrid  
Spain

### Bogota:

Carrera 11A #98-50  
Ofc. 704, Edificio Punto99  
110221, Bogota  
Colombia

### Mexico

## Contact:



### Jim Andrew

Partner

**Mail:** [jim.andrew@fidepartners.com](mailto:jim.andrew@fidepartners.com)

**LinkedIn:** [Jim Andrew](#)



### Nicolás Meneses

Senior Consultant

**Mail:** [nicolas.meneses@fidepartners.com](mailto:nicolas.meneses@fidepartners.com)

**LinkedIn:** [Nicolás Meneses](#)



### Lorenzo Morelli Avilán

Associate Consultant

**Mail:** [lorenzo.morelli@fidepartners.com](mailto:lorenzo.morelli@fidepartners.com)

**LinkedIn:** [Lorenzo Morelli](#)