INDUSTRY BRIEF · NOVEMBER 2026

Glass, Lasers and Labor

The Value of Fiber Optics in the AI Data Center Supply Chain

Where the money, the margin and the bottlenecks sit, from the glass preform to the GPU rack. Written for the people who sell, buy and build the physical layer of AI, with a close look at Europe.

Josh Leibovitch

International B2B sales and marketing · CRM · Email · Airtable · Sales enablement

Figures current as of November 3, 2026. Public sources only; see Section XI.

CONTENTS

What Is Inside

I. **Executive Overview** 1
- I.1 The Short Version 1
- I.2 Key Takeaways 1
- I.3 About This Brief 2

II. **Why Fiber, Why Now** 3
- II.1 The Spending Behind the Demand 3
- II.2 AI Changes the Traffic Pattern 3
- II.3 How Much More Fiber? 3
- II.4 Copper Runs Out of Road 4

III. **How Light Moves Through an AI Data Center** 5
- III.1 Four Networks in One Building 5
- III.2 The Parts of an Optical Link 5

IV. **The Value Chain, Link by Link** 7
- IV.1 Glass Preform and Raw Materials 7
- IV.2 Optical Fiber and Cable 7
- IV.3 Passive Connectivity 8
- IV.4 Lasers and Optical Chips 8
- IV.5 Transceivers 9
- IV.6 Switches, Co-Packaged Optics and Optical Switching 9
- IV.7 Data Center Interconnect and New Fiber Types 10
- IV.8 Installers and Integrators 10
- IV.9 The Buyers 11

V. **Where the Margin and the Power Sit** 12
- V.1 The Scorecard 12
- V.2 Scarcity Is Pricing Power 12
- V.3 Long-Term Agreements Are the New Contract 13

VI. **Five Forces Reshaping the Chain** 14
- VI.1 Optics Moves Into the Rack 14
- VI.2 Pluggables, Linear Optics and Co-Packaged Optics 14
- VI.3 Shortages at Both Ends of the Chain 15
- VI.4 Geopolitics and Supply Concentration 15
- VI.5 Labor and the Factory-Built Network 16

VII. **The Europe Angle** 17
- VII.1 Market Size and Growth 17
- VII.2 Power Redraws the Map 17
- VII.3 Policy: Sovereignty and AI Gigafactories 17

VII.4 Europe's Supplier Base 18

VIII. What This Means for a Seller 19

VIII.1 How I Read This Market 19

IX. What to Watch in 2027 and 2028 21

IX.1 The Watch List 21

X. Glossary of Terms 22

XI. Sources 24

SECTION I

Executive Overview

I.1 The Short Version

Every AI cluster runs on light. GPUs get the headlines, but the data they train on moves between chips, racks and buildings as pulses of light in glass fiber. In 2026 that physical layer became one of the tightest parts of the entire AI supply chain. This brief maps the chain from the glass preform to the hyperscaler that buys the finished network, and it shows where the value, the margin and the risk sit at each step.

The demand is real, and it is measurable. The four largest US hyperscalers guided to roughly $720 to $745 billion of capital spending for 2026 [1]. TrendForce projects that the AI optical transceiver market will grow from $16.5 billion in 2025 to $26 billion in 2026, more than 57 percent growth in one year [2]. Corning, the largest US fiber maker, grew Optical Communications sales 44 percent to $2.38 billion in the third quarter of 2026, and its Gen AI product sales more than doubled [3]. The money is moving into glass.

$26B
AI optical transceiver market forecast for 2026, up from $16.5B in 2025 [2]

+44%
Corning Optical Communications sales growth, Q3 2026 [3]

>30%
Gap between laser demand and supply reported by Lumentum [7]

13 GW
Europe's data center market by end of 2026, up 20% [15]

I.2 Key Takeaways

- **AI racks need far more fiber than cloud racks.** Estimates run from up to 10 times more fiber per AI data center to 36 times more fiber per GPU rack compared with a CPU rack [4,5].
- **The bottlenecks are upstream.** Glass preform capacity takes 18 to 24 months to add, and laser chip capacity is sold out, so supply sets the pace of the buildout, not demand [6,7].
- **Scarcity has shifted power to suppliers.** Buyers are signing multiyear, multibillion-dollar agreements to lock in fiber, cable and lasers [3,8,9].
- **Copper is losing ground inside the rack.** At 200G per link, passive copper becomes unusable at about 2 to 2.5 meters, so optics is moving into scale-up networks from 2027 onward [10,11].
- **Supply is concentrated.** Chinese-headquartered suppliers account for roughly 60 percent of datacom optical component revenue, which makes trade policy a live risk for every buyer [12].

- **Labor is the quiet constraint.** Skilled fiber technicians are scarce, and that shortage is pushing more of the work into factories as pre-terminated, plug-and-play assemblies [13,14].
- **Europe is growing fast and moving outward.** New AI capacity is heading to the Nordics, Southern Europe and Central Europe, where power is available, while the traditional hubs run out of space [15,16].

The lesson for anyone selling into this market is simple. The customer is not buying a cable or a module. The customer is buying a date: the day a cluster turns on and starts earning money. Whoever can promise capacity, deliver on time and cut skilled labor on site will win the deal, often before price comes up.

I.3 About This Brief

This brief is written for people who sell, buy or build the physical layer of AI data centers. It assumes no engineering background, and it explains each technical term the first time it appears and again in the glossary in Section X. Section II explains why demand exploded, Section III gives a short primer on how light moves through a data center, and Section IV walks the value chain one link at a time. Sections V through IX cover margin, the forces reshaping the chain, the European market, what it all means for a seller, and what to watch next.

Every figure comes from a public source current as of early November 2026, and each one carries a numbered reference to the source list in Section XI. Market sizes differ between research firms because each firm defines the market differently, so this brief names the firm behind each number. The supplier power ratings in Section V are the author's own assessment based on those sources. No company named in this brief reviewed it or paid for it.

SECTION II

Why Fiber, Why Now

II.1 The Spending Behind the Demand

As of late October 2026, Amazon expected about $220 billion of capital spending for the year, Alphabet $195 to $205 billion, Microsoft about $175 billion for calendar 2026 and Meta $135 to $145 billion [17]. Microsoft's figure came down on paper because of a lease accounting change, not because of a spending cut [17]. Together that is roughly three quarters of a trillion dollars in one year, and most of it is going into AI data centers. That figure does not include Oracle, the neoclouds or the sovereign AI programs in Europe and the Middle East.

Optical connectivity is taking a larger share of each of those dollars. LightCounting expects the top five cloud companies to put 3.1 percent of their capital spending into optical connectivity in 2026, up from 2.7 percent in 2025, with that share rising to 4.1 percent by 2031 [18]. A small change in share matters a great deal when the base is this large. A larger slice of a much larger budget is why fiber and laser suppliers are posting growth rates that look more like chip companies than cable companies.

II.2 AI Changes the Traffic Pattern

Traditional cloud data centers were built for north-south traffic. A request comes in from a user, a server answers it, and the response goes back out. AI training works differently, because thousands of GPUs must constantly exchange partial results with each other while they train a model. That GPU-to-GPU flow is called east-west traffic, and it now dominates the bandwidth of an AI data center [5].

East-west traffic multiplies the number of links. Each GPU server connects to far more switch ports than a typical CPU server, and each of those links runs at 400G or 800G. CommScope notes that a single NVIDIA DGX SuperPOD of 32 GPU servers needs 384 fiber links at 400G for its switch fabric and storage, plus 64 copper links for management [19]. Every GPU added to a cluster adds fiber, and every faster generation adds more of it.

II.3 How Much More Fiber?

The answer depends on who is counting and what they compare. Industry analysis indicates that AI-optimized data centers can need up to 10 times more fiber than conventional data centers [4]. HFCL cites estimates of roughly 36 times more fiber per GPU rack than a traditional CPU rack, and it reports that NVIDIA's 72-GPU Blackwell nodes need 16 times more fiber than traditional cloud switch racks [5]. Corning's outlook has fiber density moving above 1,000 fibers per rack, with the most demanding designs approaching 5,000 [20].

At the cluster level the numbers get large quickly. One industry source told DatacenterDynamics that a single AI compute cluster with multiple compute racks can

require more than 8,000 fiber connections, and that one AI rack can need as many as 144 high-speed connections [21]. A rack can also generate more passive connectivity revenue even when the number of racks grows slowly, because each rack needs more trunks, connectors and patching positions [20]. The exact multiple matters less than the direction, and the direction is up at every level of the building.

II.4 Copper Runs Out of Road

Copper still does a lot of work at short distances, and the industry wants to keep it there as long as it can. Inside NVIDIA's rack-scale systems, the NVLink scale-up network runs on a passive copper backplane. At 200G per link, the signal in passive copper degrades to the point of being unusable at about 2 to 2.5 meters [10]. That physical limit is why the GPUs and switches in these systems are packed so tightly together.

Larger AI pods break that limit. To join eight racks into one 576-GPU domain for Vera Rubin Ultra, NVIDIA keeps copper inside each rack and adds optical links between racks [11,22]. NVIDIA's Feynman generation, expected in mid-to-late 2028, is planned with either copper or co-packaged optical NVLink connections [11]. Each step up in speed shortens copper's reach, and each lost meter hands more of the network to fiber.

SECTION III

How Light Moves Through an AI Data Center

III.1 Four Networks in One Building

An AI data center is really several networks stacked together. Each one connects a different set of things, works over a different distance, and uses a different mix of copper and optics. The table below gives the short version, and the paragraphs that follow explain why each network matters to the fiber supply chain.

Network	What it connects	Typical medium today	Where it is heading
Scale-up	GPUs inside a rack or pod, working as one giant GPU	Copper inside the rack; optics between racks	Optics between racks from 2027, co-packaged optics later
Scale-out	Racks and pods across a hall or building	Fiber with pluggable transceivers	Faster modules, lower-power linear and co-packaged optics
Scale-across	Separate data center campuses acting as one cluster	Coherent optics over long-distance fiber	Rapid growth in 800G coherent links
Front end	Users, storage and the wider cloud	Fiber, with copper for management	Steady growth

Sources: Ciena [23]; Counterpoint Research [22]; The Register [11].

The scale-out network is the largest fiber consumer today. It is the back-end fabric that ties thousands of GPUs together through layers of switches, and every link in it needs a transceiver at each end and fiber in between. The front-end network is the older cloud network that connects the cluster to users and storage, and it grows at a steadier pace.

Scale-up and scale-across are the newer opportunities. Scale-up is moving from copper to fiber as pods grow beyond one rack, which opens a market that barely used optics before. Scale-across lets one training cluster span data center campuses built wherever power is available, so the network between buildings must behave like the network inside one [23]. Both trends add fiber in places where there was little or none.

III.2 The Parts of an Optical Link

Every optical link has the same basic parts, whether it runs two meters or two hundred kilometers. A laser creates light, and a modulator imprints data onto it. In most links today those parts sit inside a pluggable transceiver, a module the size of a pack of gum that plugs

into the front of a switch or network card. The light travels through fiber, passes through connectors, patch panels and cassettes, and reaches a photodetector at the far end that turns it back into an electrical signal.

Component	What it does	Examples of makers
Laser and optical chip	Creates the light and imprints data onto it	Lumentum, Coherent, Broadcom, Mitsubishi Electric
Transceiver	Converts electrical signals to light and back	Innolight, Eoptolink, Coherent
Optical fiber	Carries the light	Corning, Prysmian, YOFC, Fujikura, Furukawa
Cable and assemblies	Bundles, protects and pre-terminates fiber	Corning, Prysmian, CommScope, Panduit
Connectors and hardware	Joins, routes and manages fibers	US Conec, Amphenol, Molex, Huber+Suhner, R&M
Switch	Routes traffic between links	NVIDIA, Arista, Cisco, Broadcom-based systems

Sources: company references throughout Section IV. The list is illustrative, not complete.

Fiber comes in two main types. Single-mode fiber carries light over long distances and is the standard for links between buildings and for most new AI fabrics. Multimode fiber works over shorter distances and pairs with cheaper, lower-power transceivers. CommScope estimates that in one AI cluster with 768 transceivers, multimode can save up to 1.5 kilowatts compared with single-mode [19]. Each part of the link is made by a different layer of the supply chain, and the next section walks through them in order.

SECTION IV

The Value Chain, Link by Link

IV.1 Glass Preform and Raw Materials

Optical fiber starts as a preform, a rod of ultra-pure glass that is heated and drawn into strands thinner than a human hair. Preform is the least visible part of the chain and, in 2026, the most important. CRU calls preform tightness the defining constraint on the fiber market, and it notes that producers without their own preform capability struggle to secure stable supply [24]. Adding preform capacity means installing deposition equipment and qualifying new processes, which typically takes 18 to 24 months [6].

The price data show how hard the squeeze has been. CRU's global single-mode fiber price index stood at 107.9 in January 2026 and reached 263.0 in March 2026, the sharpest move since CRU began tracking it in 2017 [25]. TrendForce forecasts a global supply shortfall of 180 million fiber-kilometers in 2026 and expects spot prices for finished fiber products to reach 2.8 to more than 5 times their 2025 levels, depending on fiber type [6]. TrendForce expects the shortage to last through 2028 [6].

AI is not the only buyer competing for that glass. Telecom and broadband buildouts continue, and fiber-guided military drones have become a meaningful new source of demand, especially for bend-insensitive fiber [26]. Some producers have shifted capacity toward the higher-margin bend-insensitive grades, which tightens supply of standard fiber at the same time [26]. Whoever controls preform controls the pace of everything downstream.

IV.2 Optical Fiber and Cable

A small number of large companies draw most of the world's fiber, and the leaders are posting record results. Corning's Optical Communications segment grew 44 percent to $2.38 billion in the third quarter of 2026, its Enterprise Networks business grew 84 percent, and its Gen AI product sales grew more than 100 percent [3]. In the second quarter, the segment earned $438 million, or 21 percent of sales, and net income rose 77 percent [27]. Corning says that profitability reflects sales of high-density fiber, cable and connectivity systems rather than higher prices for standard fiber [28].

The leaders are using their position to lock in customers for years. Corning announced agreements with Meta, NVIDIA and Amazon covering fiber, cable and connectivity for AI data centers, and it signed four more multibillion-dollar agreements in the third quarter, including Verizon and AT&T [3,27]. Prysmian, the Italian cable maker, plans to invest €1.25 billion through 2031 to more than double its US fiber capacity, and it signed a deal worth up to €5.5 billion over up to ten years to supply optical cable to Molex for use inside data centers [8]. Prysmian expects its new data center agreements to add more than €10 billion in cumulative value through 2035 [8].

Corning also expects demand per GPU to keep rising. The company estimates that optical content per GPU in its enterprise business could rise 1.3 to 1.5 times by 2028, driven by larger clusters, higher bandwidth and the move to optics in scale-up networks [27]. Meanwhile, buyers without long-term agreements are waiting. One US broadband provider was quoted 52 weeks for standard loose-tube cable, compared with a normal lead time of 8 to 12 weeks [29]. In this market, a signed agreement is worth more than a low price.

IV.3 Passive Connectivity

Passive connectivity is where loose fiber becomes a working system. It includes trunk cables, cassettes, patch panels, fiber harnesses and the connectors that join them. None of these parts uses power, which is why the industry calls them passive, but each one has to fit thousands of fibers into a fixed amount of rack and panel space. As fiber counts climb, that density problem becomes an engineering problem.

New connector designs are answering it. Very small form factor connectors such as US Conec's MMC deliver more than three times the port density of the standard MPO connector [30]. Pre-terminated trunks, modular panels, high-density cassettes and factory-assembled harnesses are capturing a larger share of connectivity value than field termination, especially in standardized rack-scale designs [20]. The work is moving from the job site to the factory, and the value is moving with it.

The connectivity layer is also where the next generation of optics will plug in. In 2026 Molex acquired Teramount, an Israel-based developer of detachable fiber-to-chip connections built for high-volume co-packaged optics [31]. That deal shows where connector makers expect the market to go: from the front panel of a switch to the edge of the chip itself. A part of the chain that buyers once treated as a commodity is turning into precision engineering.

IV.4 Lasers and Optical Chips

The laser is the narrowest point in the optical supply chain. High-speed data center transceivers rely on lasers built from indium phosphide, a compound semiconductor that emits light efficiently at the wavelengths fiber is designed for. The key device is the electro-absorption modulated laser, or EML, which generates light and imprints data on it in one chip. For long-haul coherent links, indium phosphide remains the only commercially proven substrate for high-performance EML chips [9].

Demand has outrun every expansion plan. Lumentum grew fiscal third-quarter 2026 revenue 90 percent to a record $808 million, said its components and laser chips are effectively sold out, and put the gap between supply and demand at more than 30 percent [7]. Lumentum has raised EML output eight times since fiscal 2023 and bought a fifth indium phosphide fab in Greensboro, North Carolina, but that site is not expected to contribute until early 2028 [7,32]. Coherent, the other major Western laser maker, says assembly and test capacity is not its bottleneck, which puts the ceiling on its data center revenue squarely on laser output [33].

The raw material adds another layer of risk. Indium is recovered mainly as a byproduct of zinc refining, which makes it hard to produce more quickly. China accounts for roughly 70 percent of global refined indium production and has tightened export controls on indium and indium phosphide since 2025 [9]. To secure supply, NVIDIA committed $2 billion each to two key suppliers in March 2026 and signed long-term supply agreements [9]. When the world's most valuable chip company invests directly in its laser suppliers, the laser is the bottleneck.

IV.5 Transceivers

Transceivers are the highest-volume active product in the optical chain. TrendForce projects the AI optical transceiver market will grow from $16.5 billion in 2025 to $26 billion in 2026 [2]. TrendForce also estimates that modules running at 800G and above will rise from 19.5 percent of shipments in 2024 to more than 60 percent in 2026 [34]. The 1.6T generation is entering mass production now, and 3.2T modules are expected to enter production in 2028 [2,12].

This layer is highly concentrated, and most of it sits in China. Counterpoint Research estimates that Zhongji Innolight holds about 27 percent of global data center transceiver revenue and Coherent about 17 percent, followed by Eoptolink [35]. Cignal AI attributes $2.6 billion of datacom optical component revenue to Innolight in the first quarter of 2026, or 34 percent of a $7.7 billion worldwide market, and estimates that Chinese-headquartered suppliers hold about 60 percent of that market [12]. The supply chain still runs both ways, because Chinese module makers depend on US-made signal processors from Broadcom and Marvell and on lasers from Lumentum, Coherent and Mitsubishi Electric [36].

TrendForce identifies component shortages, not module assembly, as the main limit on transceiver output, and it notes that buyers are moving to long-term agreements instead of spot purchases [2]. Power is the other pressure. Semtech says 200G linear pluggable optics target about 10 watts per module, compared with 23 to 25 watts for modules that use a digital signal processor [37]. A cluster with tens of thousands of modules turns those watts into megawatts, so transceiver makers now compete on power as hard as they compete on price.

IV.6 Switches, Co-Packaged Optics and Optical Switching

Co-packaged optics, or CPO, moves the optical engine off the front panel and places it right next to the switch chip. That shortens the electrical path from inches to millimeters and removes much of the power a pluggable module burns. NVIDIA moved its Spectrum-X Ethernet Photonics CPO switch into full production in 2026 with a named five-vendor supply chain [38]. NVIDIA has framed the platform's advantage over networks built on traditional transceivers as 5 times better power efficiency [38].

NVIDIA is not alone. Broadcom's 51.2T Bailly CPO switch has entered volume manufacturing, and TrendForce reports that it cuts power by up to 70 percent compared with pluggable transceivers and has completed validation with hyperscale operators such as Meta [39]. TrendForce names optical engines, silicon photonics chips and advanced packaging as the bottlenecks, and it expects the CPO switch market to hit its volume ramp in 2027 and 2028 [39]. It

also expects vendors that control silicon photonics design, optical engine manufacturing and packaging in-house to lead that ramp [39].

Optical circuit switching is a third path. Google's Apollo system uses tiny MEMS mirrors to connect fiber directly to fiber, which avoids converting light to electricity and back at every hop [34]. TrendForce reports that one such switch draws about 100 watts, compared with about 3,000 watts for a traditional switch [34]. Each of these designs moves value away from the pluggable module and toward whoever owns the switch, the package and the photonics.

IV.7 Data Center Interconnect and New Fiber Types

Data center interconnect, or DCI, links buildings and campuses. Scale-across has turned it into a growth market, because operators now build AI campuses wherever power is available and then connect them into one training cluster [23]. Dell'Oro expects a sharp increase in shipments of 800 ZR+ coherent modules in 2026, driven by scale-across networks, and Cignal AI reports that 800ZR and ZR+ optics began large-scale rollouts in support of these architectures [40]. Coherent optics, line systems and high-count outdoor cable all benefit.

New fiber types are moving out of the lab. Hollow-core fiber carries light through air instead of solid glass, and Microsoft says its version delivers up to 47 percent faster data transmission and about 33 percent lower latency than standard single-mode fiber [41]. Microsoft announced a two-year plan in November 2024 to install 15,000 kilometers of hollow-core fiber, with Corning and Heraeus Covantics as manufacturing partners [42]. Huber+Suhner has set up special processes at its plant in Herisau, Switzerland, to build hollow-core fiber into multi-fiber cables for Microsoft [43].

Lower latency changes where data centers can be built. If light travels faster, two sites can sit farther apart and still meet the same latency budget, which gives operators more freedom to chase cheap power and land [43]. Hollow-core fiber is still a small and expensive product, and standardization is unsettled because several designs are competing [42]. It is worth watching because it could turn latency into a site selection tool.

IV.8 Installers and Integrators

Every fiber in an AI data center has to be pulled, terminated, tested and labeled by a person. That makes labor a real limit on how fast new capacity turns on. CommScope has described a vast shortage of people who can assemble these networks, and fiber splicers are being pulled between data center work and broadband buildouts at the same time [13]. Back in 2024, the Fiber Broadband Association estimated that an additional 58,000 workers would be needed just to deliver government-funded fiber projects in the US [44].

The largest buyers are now training their own workforce. In April 2026, Meta and CBRE launched LevelUp, a free four-week program designed to turn candidates with no prior experience into fiber technicians ready for data center construction sites [45]. Google and AWS

run their own fiber training programs as well [44]. When the customer has to build its own labor pool, labor has become part of the supply chain.

The industry's other answer is to move skilled work off the job site. CommScope describes a two-phase approach in which splicing happens offsite, so crews on site can simply plug in pre-built assemblies [13]. Panduit's experts expect more factory-terminated, plug-and-play fiber systems because they install faster and need less skilled labor [14]. The supplier that saves a customer hundreds of technician hours has a real advantage over one that only saves a few dollars per meter.

IV.9 The Buyers

The buyers sit at the end of the chain and drive everything above them. Hyperscalers such as Amazon, Microsoft, Google and Meta buy the most and increasingly design their own networks. Neoclouds, the specialist GPU cloud providers, are growing quickly, and CBRE reports that new signings for European colocation capacity for AI purposes quadrupled in the first half of 2026 [15]. Colocation operators, telecom carriers and government-backed AI programs round out the list.

These buyers do not purchase fiber as one decision. A typical AI data center project involves network architects who set the design, procurement teams who manage supply risk, construction and program managers who own the schedule, and operations teams who live with the result. Each group cares about something different, from bandwidth and power to lead time and how easy a cassette is to swap at 3 a.m. Anyone selling into this market needs a message for each of them.

SECTION V

Where the Margin and the Power Sit

V.1 The Scorecard

The table below rates each link in the chain on who leads it, what limits it and how much pricing power its suppliers hold in late 2026. The ratings are the author's assessment, built from the evidence in Section IV. They are meant as a quick map for a seller or buyer, not as investment guidance.

Link	Leading players	Main constraint	Supplier power	Lead-time risk
Preform and glass	Corning, Prysmian, YOFC, Fujikura, Furukawa	18 to 24 months to add capacity	Very high	High through 2028
Fiber and cable	Corning, Prysmian, YOFC, Fujikura, Sumitomo	Preform supply	High	High
Passive connectivity	Corning, CommScope, Panduit, US Conec, Amphenol, Molex, Huber+Suhner	Density engineering and assembly capacity	Rising	Medium
Lasers and optical chips	Lumentum, Coherent, Broadcom, Mitsubishi Electric	Indium phosphide wafer capacity	Very high	Very high
Transceivers	Innolight, Eoptolink, Coherent	Lasers, signal processors, trade policy	Medium to high	High
Switches and CPO	NVIDIA, Broadcom, Arista, Cisco	Optical engines and advanced packaging	High	Medium to high
Interconnect (DCI)	Ciena, Cisco, Marvell, Nokia	Coherent chips and line systems	High	Medium
Installers	Regional contractors and integrators	Skilled labor	Rising	High

Author's assessment based on sources [2,3,6,7,12,13,39,40].

V.2 Scarcity Is Pricing Power

Margin in this chain follows scarcity. Lumentum, which sells the scarcest product, reported a non-GAAP gross margin of 47.9 percent and an operating margin of 32.2 percent in its fiscal third quarter of 2026 [7]. Coherent reported a non-GAAP gross margin of 40.2 percent in its fiscal fourth quarter, with revenue up 34 percent to $2.05 billion [46]. Corning's Optical

Communications segment reached record profitability at 21 percent net margin in the second quarter of 2026 [27].

The pattern is clear when the links are lined up. The parts of the chain that take years to expand, such as preform and indium phosphide lasers, hold the most pricing power. The parts that depend on assembly scale, such as transceivers, compete harder and carry more exposure to trade policy. A seller who understands where a product sits on that spectrum knows whether to lead with price, with availability or with engineering.

V.3 Long-Term Agreements Are the New Contract

The purchase order is giving way to the long-term agreement. Verizon and Corning signed a supply agreement in September 2026 covering 129 million fiber-kilometers of high-density fiber and connectivity for delivery between 2027 and 2032 [6]. Prysmian's Molex deal runs up to ten years [8]. TrendForce reports that some hyperscalers are locking in CPO supply through long-term deals and taking strategic stakes in upstream silicon photonics suppliers [39].

These agreements change who gets served first. Market reports describe a market in which long-term contract customers get priority over spot buyers, with extended lead times and elevated prices for everyone else [26]. That pushes smaller operators, integrators and enterprise buyers to the back of the line. The relationship has become as important as the product, and that favors sellers who can plan with a customer years ahead instead of quoting one order at a time.

SECTION VI

Five Forces Reshaping the Chain

VI.1 Optics Moves Into the Rack

Scale-up networking is the largest new market for fiber, and it is opening now. NVIDIA's plan for Vera Rubin Ultra NVL576 uses copper inside each rack and optical connections between racks, and its Feynman generation is planned with co-packaged optical NVLink as an option [11]. Jensen Huang told GTC that NVIDIA will scale up with both copper and optics, and said the ecosystem needs a lot more capacity [11]. Copper will not disappear, but it will hold less of the network with each generation.

The suppliers are already planning around this shift. Lumentum expects scale-up CPO volume shipments to begin in the second half of 2027, followed by first customer deployments in 2028, while Coherent expects scale-out CPO revenue to begin in the second half of 2026 and scale-up revenue in 2027 [47]. In 2025 Corning said that replacing copper in scale-up networks with optical links could create a market two to three times larger than its existing $2 billion enterprise fiber business [48]. The rack itself is becoming fiber territory.

VI.2 Pluggables, Linear Optics and Co-Packaged Optics

Three ways of building an optical link now compete, and all three will coexist for years. The traditional pluggable module carries its own digital signal processor, which makes it easy to swap and interoperate but costly in power. Linear pluggable optics remove that processor and rely on the host chip instead, which saves power but demands tighter engineering. Co-packaged optics go further and build the optics into the switch package itself.

Approach	Where the signal processing sits	Power	Serviceability	Status in 2026
Retimed pluggable	Inside the module	Highest; about 23 to 25 W for a 200G-per-lane module	Hot-swappable from the front panel	Mainstream at 800G; 1.6T ramping
Linear pluggable (LPO)	In the host switch or accelerator	About 10 W for a 200G-per-lane module	Hot-swappable, with tighter host requirements	Shipping; standards work under way
Co-packaged optics (CPO)	Integrated with the switch chip	Up to 70% lower than pluggables in Broadcom's design	Optics no longer a front-panel module	Volume production started; ramp in 2027 to 2028

Sources: Semtech via IT Brief [37]; TrendForce [39]; StorageReview [38].

Forecasts for these segments vary widely, which is a warning about all market numbers in this space. One firm valued the LPO market at about $2.07 billion in 2025, while another valued it at $380 million for the same year, because each defines the market differently [49,50]. For a seller, the safer signal is what the leading buyers are actually qualifying and installing. Right now they are buying all three.

VI.3 Shortages at Both Ends of the Chain

The optical chain is squeezed at its two most basic inputs: glass and lasers. Preform capacity takes 18 to 24 months to build, and TrendForce expects the fiber shortage to last through 2028 [6]. Lumentum's new indium phosphide fab will not contribute until early 2028, and the company has said its supply gap is more than 30 percent [7]. Neither constraint can be fixed with money in the short term.

That timing shapes every deal in 2027. Buyers who secured capacity early will build on schedule, and buyers who did not will pay more and wait longer. Some fiber prices stabilized after the first-quarter spike, but they settled at several times their 2025 levels rather than falling back [6,51]. Anyone planning a project for 2027 or 2028 should treat optical supply the way they treat power: as something to reserve before the design is finished.

VI.4 Geopolitics and Supply Concentration

The concentration of transceiver supply in China has made trade policy a design input. On July 22, 2026, the FCC adopted a rule allowing it to deny authorization to devices containing logic-bearing components from companies on its Covered List, and the rule took effect on October 13, 2026 [52]. The major Chinese optical module makers, including Innolight and Eoptolink, were left off the Covered List, so as of mid-October the rule does not restrict imports of their transceivers [52]. A transceiver-specific restriction remains under consideration, and reporting tied to Morgan Stanley meetings in Washington suggests any such rule would phase in starting with the 3.2T generation [52].

The exposure runs in both directions. Innolight generated 62 percent of its revenue from the United States in the first quarter of 2026, and Chinese module makers depend on US signal processors and lasers [36]. On the raw material side, China controls most refined indium and has tightened export controls on it [9]. Optical transceiver modules were explicitly excluded when the current US tariffs were put in place [12].

For buyers, country of origin has become a line in the specification. Qualifying a new transceiver supplier requires validation for interoperability, power, thermal performance and reliability, which takes time that no buyer wants to spend in a hurry [53]. Coherent says the deliberations over US import restrictions have already changed customer conversations [33]. A supplier with diversified manufacturing and a clean compliance story now has an advantage that does not show up on a datasheet.

VI.5 Labor and the Factory-Built Network

The last force is the simplest. There are not enough skilled people to splice and terminate the fiber that AI data centers need, and training programs take time to scale [13,45]. The industry is responding by building more of the network in factories and shipping it as pre-terminated trunks, cassettes and harnesses [14,20]. That shift moves value toward suppliers with engineering and assembly capacity, and away from field labor.

The trend also changes how projects are planned. A factory-built network has to be designed earlier and in more detail, because lengths, fiber counts and connector types are fixed before the parts ship. That gives suppliers who work with the customer during design a large advantage over suppliers who only respond to a bill of materials. In a market short on both glass and people, the winning supplier saves both.

SECTION VII

The Europe Angle

VII.1 Market Size and Growth

Europe is the second-largest data center market in the world after the United States, and it is growing fast [16]. CBRE expects Europe's total data center market to grow 20 percent in 2026 to reach 13 GW of capacity [15]. The hyperscaler self-build segment alone is expected to grow 22 percent to 4.3 GW, the seventeenth consecutive year of double-digit growth for that segment [15]. New signings for European colocation capacity for AI purposes quadrupled in the first half of 2026 [15].

Other measures point the same way. The European Data Centre Association reports that European IT power capacity rose from 10,539 MW in 2023 to 14,784 MW in 2025, ahead of earlier forecasts [54]. The association forecasts €176 billion of cumulative investment between 2026 and 2031 [55]. Different firms measure capacity differently, but every measure shows the same trend: Europe is building at a pace it has never seen.

VII.2 Power Redraws the Map

Power, not land or money, decides where Europe builds. Combined live capacity in the FLAP-D markets of Frankfurt, London, Amsterdam, Paris and Dublin has reached 3.8 GW, more than double its 2019 level [56]. Vacancy across those markets stands at just 6.4 percent, and Frankfurt is the tightest at 3.1 percent [57]. MUFG reports that grid connection delays are driving more than half of new AI growth into secondary markets [16].

The Nordics and Southern Europe are the main beneficiaries. AI training workloads favor places with abundant power, including the Nordics and parts of Southern Europe, and the industry is spreading into Iberia and Central and Eastern Europe as well [16,54]. The EUDCA expects grid readiness and permitting to shape how quickly Europe can add capacity over the next several years [54]. The map of European AI is being drawn by the power grid.

This shift matters directly for fiber. Large AI campuses far from the traditional hubs still need to connect back to users, to cloud regions and to each other, which drives demand for long-haul cable, coherent optics and scale-across links [23]. New regions may also have thinner benches of experienced fiber crews than the established hubs, which makes pre-terminated systems and on-site support more valuable there. That second point is the author's inference from the labor trends in Section IV.8, not a published finding.

VII.3 Policy: Sovereignty and AI Gigafactories

European policy is adding public money and a sovereignty agenda to private demand. On July 30, 2026, the European Commission opened a call for tenders to support at least seven AI Gigafactories, targeting more than €30 billion of investment, with the call open until

November 12, 2026 [58]. Selected Gigafactories are expected to begin operations within 18 months of signature [58]. Each one will combine advanced AI processors, high-speed connectivity and energy-efficient data centers [58].

The Cloud and AI Development Act sits alongside that program and is projected to triple the EU's data center capacity within five to seven years [58]. Both initiatives emphasize building advanced AI on European infrastructure under European rules. That emphasis will shape procurement, and it gives suppliers with European manufacturing a story to tell.

Europe is also feeling the fiber shortage directly. CRU data show European prices for standard single-mode fiber rose about 136 percent between January and March 2026 [26]. For European buyers, the global shortage is not a distant headline. It is a line item on the next project budget.

VII.4 Europe's Supplier Base

Europe has real strength in the passive and fiber layers of the chain. Prysmian, based in Milan, describes itself as the only major domestic fiber and optical cable player in Europe [8]. In Germany, Reichle & De-Massari, Rosenberger and Huber+Suhner hold strong positions in connectors, patch cords and distribution [59]. Huber+Suhner is also building hollow-core fiber cables for Microsoft at its Swiss plant [43], and Nexans in France remains a major cable maker.

The active layers are a different story. The largest laser makers are based in the US and Japan, and most transceiver volume comes from Chinese suppliers [12,35,36]. That leaves European AI projects dependent on imported lasers and modules even when the fiber and connectivity are made locally. For a seller, that gap is both a risk to manage with customers and an opening for any supplier that can offer a more local or diversified supply story.

SECTION VIII

What This Means for a Seller

VIII.1 How I Read This Market

I sell technical video hardware to multinational buyers across Europe, and I have spent years building the CRM, email and sales systems behind that work. I am not an optical engineer. I read this market as a seller: who the buyer is, what keeps them up at night, and what makes them sign. The seven points below are my conclusions from the research in this brief.

1. Sell the date, not the part

When glass and lasers are scarce, the customer's biggest risk is a cluster that sits idle waiting for parts. Long-term contract customers get served first, and spot buyers wait [26]. A seller who can show allocation, a credible delivery schedule and a backup plan is selling the one thing the customer cannot get elsewhere. Price still matters, but it comes after certainty.

2. Sell labor hours

Skilled fiber technicians are scarce, and the biggest buyers are training their own [45]. Pre-terminated, plug-and-play systems cut the hours a project needs on site [14]. A proposal that counts the technician hours it saves speaks directly to the construction and program managers who own the schedule.

3. Sell watts

Power is the hard limit on AI data centers, and nowhere more than in Europe. Every watt saved in the network is a watt that can go to GPUs. Linear optics, co-packaged optics and the right choice of fiber type all save measurable power [19,37,39]. A seller who can translate a product into kilowatts saved per hall is speaking the language of the people who approve the budget.

4. Build a message for each member of the buying committee

Network architects want bandwidth, reach and a roadmap to the next speed. Procurement wants supply security, a clean country-of-origin story and contract terms that hold up. Construction managers want fewer surprises on site, and operations teams want parts that are easy to find, swap and test. One pitch will not land with all four, so the sales materials need a version for each.

5. Get in at the design stage

Fiber counts, pathways, connector types and panel layouts are decided early, and factory-built networks lock them in before anything ships. A supplier that helps the customer design the network shapes the bill of materials. A supplier that waits for the request for quote competes on price for a specification written around someone else.

6. Know the supply chain story cold

Customers now ask where a product is made, where its lasers come from and how exposed it is to trade rules [52]. Sellers need clear, honest answers before the customer asks. A compliance question answered well builds trust, and a vague answer can end the deal.

7. In Europe, follow the power

New AI capacity is moving to the Nordics, Iberia and Central and Eastern Europe because that is where grid connections are available [16]. The AI Gigafactories program will add publicly backed projects with tight timelines and a strong preference for European infrastructure [58]. A seller covering Europe should map where power is being connected and build relationships there before the tenders go out.

The physical layer used to be a line item that nobody outside engineering noticed. In 2026 it decides when an AI cluster turns on. That makes the people who sell it more important to the customer than they have been in a generation.

SECTION IX

What to Watch in 2027 and 2028

IX.1 The Watch List

The next two years will decide whether the optical supply chain catches up with AI demand. The items below are the signals that matter most for pricing, availability and where the value moves. Each one has a date attached, which makes it easy to check progress against the forecast.

What to watch	Why it matters	Expected timing
New glass preform capacity	The first real relief for fiber prices and lead times	Late 2027 at the earliest, given 18 to 24 month build times [6]
Fiber shortage	TrendForce expects the supply gap to persist	Through 2028 [6]
Lumentum Greensboro laser fab	Adds indium phosphide laser capacity in the US	Early 2028 [7]
CPO switch volume ramp	Shifts value toward switch and packaging owners	2027 to 2028 [39]
Optical scale-up	Opens a new fiber market inside AI pods	Scale-up CPO shipments from second half of 2027; Feynman in mid-to-late 2028 [11,47]
3.2T transceivers	The next speed generation, and a possible trigger for US trade rules	Production expected in 2028 [12,52]
FCC transceiver rule	Could restrict new Chinese module models	Pending as of November 2026 [52]
EU AI Gigafactory awards	Publicly backed European AI capacity	Tenders close November 12, 2026; operations within 18 months of signature [58]

Timing reflects company statements and analyst forecasts as of early November 2026 and may change.

Two questions sit under all of these dates. The first is whether AI spending keeps growing at its current pace, because the entire optical buildout depends on it. The second is whether new capacity arrives before the next speed generation raises demand again. If both answers are yes, the suppliers of glass, lasers and skilled labor will keep the upper hand through the end of the decade.

SECTION X

Glossary of Terms

These are the terms used in this brief, in plain English. Each definition is written for a reader who works with this industry but does not design optical networks. Where a term has a common abbreviation, the abbreviation follows in parentheses.

Term	Meaning
800G, 1.6T and 3.2T	Transceiver speed classes: 800 gigabits, 1.6 terabits and 3.2 terabits per second. Each generation roughly doubles the data one module can carry.
800ZR and ZR+	Standards for 800G coherent pluggable modules that plug straight into routers and switches to link data centers over long distances.
Back-end network	The network that ties GPUs together for training. It carries most of the east-west traffic in an AI data center.
Bend-insensitive fiber	Fiber designed to lose little signal when bent tightly. It is used inside dense data centers and in fiber-guided drones.
Co-packaged optics (CPO)	A design that places the optical engine inside the same package as the switch or processor chip instead of in a front-panel module.
Coherent optics	Optical technology that encodes data in the phase and amplitude of light. It is used for long-distance links between data centers.
Colocation	A data center business that leases space, power and cooling to many customers in a shared facility.
Data center interconnect (DCI)	The links that connect separate data center buildings, campuses or regions.
Digital signal processor (DSP)	A chip inside a traditional transceiver that cleans up and retimes the signal. It improves reliability but uses significant power.
East-west traffic	Data moving between servers inside a data center, such as GPUs exchanging results during training.
EML	Electro-absorption modulated laser. A laser chip, usually made from indium phosphide, that both generates light and imprints data on it.
Fiber-kilometer	One kilometer of a single optical fiber. A cable with 144 fibers that runs one kilometer contains 144 fiber-kilometers.
FLAP-D	Europe's five traditional data center hubs: Frankfurt, London, Amsterdam, Paris and Dublin.
Front-end network	The network that connects an AI cluster to users, storage and the wider cloud.

Term	Meaning
Hollow-core fiber	Fiber that guides light through an air-filled core instead of solid glass, which lowers latency.
Hyperscaler	One of the very large cloud and platform companies, such as Amazon, Microsoft, Google and Meta.
Indium phosphide (InP)	A compound semiconductor used to make the lasers in high-speed optical transceivers.
Linear pluggable optics (LPO)	A pluggable transceiver without its own DSP. It relies on the host chip for signal processing and uses less power.
Long-term agreement (LTA)	A multiyear supply contract that reserves capacity, often at agreed prices, in exchange for a volume commitment.
MPO	Multi-fiber push-on connector. The standard connector for joining many fibers at once in data centers.
Multimode fiber	Fiber for shorter links that pairs with lower-cost, lower-power transceivers.
Neocloud	A specialist cloud provider focused on renting GPU capacity for AI workloads.
NVLink	NVIDIA's high-speed link for connecting GPUs into a single scale-up system.
Optical circuit switch (OCS)	A switch that uses tiny mirrors to connect fibers directly, without converting light to electricity.
Pluggable transceiver	A removable module that converts electrical signals to light and back, plugged into the front of a switch or network card.
Pre-terminated assembly	A cable, trunk or harness that arrives from the factory with connectors attached and tested, ready to plug in.
Preform	The rod of ultra-pure glass from which optical fiber is drawn.
Scale-across	Linking separate data center campuses so they operate as one AI training cluster.
Scale-out	Linking racks and pods across a building through layers of switches.
Scale-up	Linking GPUs within a rack or pod so they act as one very large processor.
Silicon photonics	Optical components built on silicon chips with semiconductor manufacturing methods.
Single-mode fiber	Fiber that carries light over long distances. It is the standard for links between buildings and most new AI fabrics.
Spot market	Buying at current market prices without a long-term contract.
VSFF connector	Very small form factor connector, such as MMC, that fits more fibers into the same panel space than MPO.

SECTION XI

Sources

Sources are numbered in the order they first appear in the text. All were accessed between October and early November 2026. Where a source reports another firm's data, such as CRU or LightCounting, the original firm is named in the text.

1. Yield Theory, "Hyperscaler Capex 2026: $720B to $745B Guidance by Company". https://www.yieldtheory.app/research/hyperscaler-ai-capex-tracker-2026
2. TrendForce, "Global AI Optical Transceiver Market to Reach US$26 Billion in 2026," April 20, 2026. https://www.trendforce.com/presscenter/news/20260420-13017.html
3. Corning Incorporated, Third-Quarter 2026 results, Form 8-K Exhibit 99, October 27, 2026. https://www.sec.gov/Archives/edgar/data/0000024741/000002474126000273/glw-20261027xex99xq32026.htm
4. Electronic Specifier, "Why Fibre Solutions in AI Data Centres Matter". https://www.electronicspecifier.com/products/cables-and-connecting/why-fibre-solutions-in-ai-data-centres-matter/
5. HFCL, "How the AI Infrastructure Boom Is Reshaping Fiber Optic Network Design," July 20, 2026. https://hfcl.com/insights/blogs/how-the-ai-infrastructure-boom-is-reshaping-fiber-optic-network-design
6. TrendForce, "Optical Fiber Spot Prices Surge Severalfold in 2026," October 19, 2026. https://www.trendforce.com/presscenter/news/20261019-13275.html
7. MarketBeat, "Lumentum Q3 Earnings Call Highlights," May 5, 2026. https://www.marketbeat.com/instant-alerts/lumentum-q3-earnings-call-highlights-2026-05-05/
8. Prysmian, "Prysmian to Accelerate Data Center Growth in Digital Solutions," July 20, 2026. https://www.prysmian.com/en/media/press-releases/digital-solutions-20-07-2026
9. TrendForce News, "InP Shortage Emerges as AI Optical Interconnect Bottleneck," August 6, 2026. https://trendforce.com/news/2026/08/06/news-inp-shortage-emerges-as-ai-optical-interconnect-bottleneck/
10. HPCwire, "Inside Nvidia's Silicon Photonics Roadmap," April 20, 2026. https://www.hpcwire.com/2026/04/20/inside-nvidias-silicon-photonics-roadmap/
11. The Register, "How Nvidia Learned to Embrace the Light in Its Quest for Scale," April 5, 2026. https://www.theregister.com/2026/04/05/nvidia_optical_scale_up/
12. Cignal AI, "FCC Ban on New Chinese Optical Modules," August 2026. https://cignal.ai/2026/08/fcc-ban-on-new-chinese-optical-modules/
13. Fierce Network, "The Labor Shortage Is Changing the Way Data Center Networks Are Built". https://www.fierce-network.com/cloud/labor-shortage-changing-way-data-center-networks-are-built
14. Panduit, "Broadband 2026: Panduit Experts Predict What's Next". https://www.panduit.com/latam/es/about/blog/broadband-2026-panduit-experts-predict-whats-next.html
15. DatacenterDynamics, "CBRE: Europe Data Center Market to Grow by Nearly a Quarter in 2026". https://www.datacenterdynamics.com/en/news/cbre-europe-data-center-market-to-grow-by-nearly-a-quarter-in-2026/
16. MUFG, "AI Weekly: Europe Data Center Market," July 31, 2026. https://www.mufgamericas.com/sites/default/files/document/mufgamericas_com/2026-07/AI_Weekly_7_31_Europe_Data_Center_Market.pdf
17. BuilderMuse, "Hyperscaler Data Center Capex 2026" (guidance as of October 26, 2026). https://buildermuse.com/commercial/hyperscaler-data-center-capex-2026/

18. CN Biz Insight, "China's Optical Transceiver Industry 2026" (citing LightCounting). https://cnbizinsight.com/chinas-optical-transceiver-industry-2026-the-backbone-of-the-ai-computing-era/
19. IT Brief Australia, "CommScope Highlights Key Cabling Considerations for AI Data Centers". https://itbrief.com.au/story/commscope-highlights-key-cabling-considerations-for-ai-data-centers
20. GII Research, "AI Data Center High-Density Fiber Connectivity" report summary. https://www.giiresearch.com/report/ksi2148218-ai-data-center-high-density-fiber-connectivity.html
21. DatacenterDynamics, "When AI Grows New Limbs: The Fiber Scale-Out". https://datacenterdynamics.com/en/marketwatch/when-ai-grows-new-limbs-the-fiber-scale-out
22. Counterpoint Research, "From Blackwell to Feynman: Analyzing NVIDIA's Optics Roadmap". https://counterpointresearch.com/jp/insights/from-blackwell-to-feynman-analyzing-nvidias-optics-roadmap
23. Ciena, "What Does DCI Really Mean Today? From Scale Across to Submarine," 2026. https://www.ciena.com/insights/blog/2026/what-does-dci-really-mean-today-from-scale-across-to-submarine
24. CRU Group, "Optical Fibre Prices Rise as Preform Availability Tightens," 2026. https://www.crugroup.com/en/communities/thought-leadership/2026/optical-fibre-prices-rise-as-preform-availability-tightens/
25. DYS Fiber, "Fiber Optic Cable Pricing: Cost Drivers" (citing CRU price index). https://www.dysfiber.com/blog/sourcing/fiber-optic-cable-pricing
26. Hengtong, "Optical Fiber Shortage 2026: What Buyers Need to Know" (citing CRU). https://www.hengtongglobal.com/info/optical-fiber-shortage-2026-buyers-guide-103466005.html
27. MarketBeat, "Corning Q2 Earnings Call Highlights," July 28, 2026. https://www.marketbeat.com/instant-alerts/corning-q2-earnings-call-highlights-2026-07-28/
28. Converge Digest, Corning second-quarter 2026 results coverage. https://convergedigest.com/?p=114226
29. Glory Optical, "FTTH Fiber Shortage 2026". https://www.gloryoptics.com/info/ftth-fiber-optic-supply-shortage-2026-how-bea-103451274.html
30. T&S Communications, MMC connector solution (US Conec VSFF). https://www.china-tscom.com/mmc-cables/
31. I-Connect007, TrendForce release with industry briefs including Molex's acquisition of Teramount, 2026. https://iconnect007.com/article/148744/googles-highspeed-interconnect-architecture-to-push-800g-optical-transceiver-share-past-60-by-2026/148741/aep
32. MarketBeat, "Lumentum Touts AI Optics Boom, New Fab and 1.6T Transceiver Shipments," March 20, 2026. https://www.marketbeat.com/instant-alerts/lumentum-touts-ai-optics-boom-2b-quarterly-target-new-fab-and-16t-transceiver-shipments-2026-03-20/
33. Futurum Group, "Coherent Q4 FY 2026 Earnings: 1.6T Transceivers Ramp," August 14, 2026. https://futurumgroup.com/insights/coherent-q4-fy-2026-earnings-1-6t-transceivers-ramp-cpo-revenue-approaches/
34. TrendForce, "Google's High-Speed Interconnect Architecture to Push 800G+ Optical Transceiver Share Past 60% by 2026," February 10, 2026. https://www.trendforce.com/presscenter/news/20260210-12919.html
35. Counterpoint Research, "Who Wins and Loses in the Proposed FCC Ban on Chinese Transceivers?" August 4, 2026. https://counterpointresearch.com/en/insights/fcc-china-ban-ai-component
36. Yahoo Finance, "New Report Warns China Transceiver Ban Would Hurt U.S. AI Giants," 2026. https://finance.yahoo.com/technology/ai/articles/report-warns-china-transceiver-ban-093946021.html
37. IT Brief UK, "Semtech Sets 200G LPO Targets for Data Centre Links," August 2026. https://itbrief.co.uk/story/semtech-sets-200g-lpo-targets-for-data-centre-links

38. StorageReview, "NVIDIA Spectrum-X Ethernet Photonics Enters Full Production," August 15, 2026. https://www.storagereview.com/it/news/nvidia-spectrum-x-ethernet-photonics-enters-full-production-with-4x-fewer-lasers-and-a-five-vendor-cpo-supply-chain

39. TrendForce, CPO switch volume ramp and supply bottlenecks, July 27, 2026. https://trendforce.com/presscenter/news/20260727-13151.html

40. Lightwave, "Cloud and Data Center Interconnect (DCI) Creates a New Coherent Optics Wave". https://www.lightwaveonline.com/home/article/55358759/cloud-and-data-center-interconnect-dci-creates-a-new-coherent-optics-wave

41. Tom's Hardware, "Microsoft Expects to Commercialize MicroLED Datacenter Cables by Late 2027". https://www.tomshardware.com/tech-industry/microsoft-expects-to-commercialize-microled-datacenter-cables-by-late-2027

42. K. Borzycki, "Recent Progress in Development of Hollow-Core Fibers," Preprints.org, February 2026. https://www.preprints.org/manuscript/202602.1976

43. TelecomTV, "Huber+Suhner Strengthens Collaboration with Microsoft for Deployment of Hollow Core Fiber," April 21, 2026. https://www.telecomtv.com/content/access-evolution/huber-suhner-strengthens-collaboration-with-microsoft-for-deployment-of-hollow-core-fiber-55338/amp/

44. Fierce Network, "Meta Is Training Up a New Fiber Workforce for Its Data Center Boom". https://www.fierce-network.com/cloud/meta-training-new-fiber-workforce-its-data-center-boom

45. ConstructConnect News, "Meta, CBRE Launch Program to Train Jobsite Fiber Technicians," April 2026. https://news.constructconnect.com/meta-cbre-launch-program-to-train-jobsite-fiber-technicians-from-the-ground-up

46. Coherent Corp., Fourth Quarter and Full Year Fiscal 2026 results, Form 8-K Exhibit 99.1, August 12, 2026. https://www.sec.gov/Archives/edgar/data/0000820318/000119312526346860/d128030dex991.htm

47. OnMSFT, Lumentum and Coherent CPO timing updates, 2026. https://onmsft.com/?p=60157

48. Converge Digest, "Corning's AI and DCI Push Targets Multi-Billion-Dollar Fiber Expansion," August 8, 2025. https://convergedigest.com/?p=104328

49. MarketsandMarkets, "Linear Pluggable Optics (LPO) Market," July 2026. https://www.marketsandmarkets.com/Market-Reports/linear-pluggable-optics-lpo-market-73870478.html

50. Meticulous Research, "Linear Pluggable Optics (LPO) Market (2026 to 2036)". https://www.meticulousresearch.com/product/linear-pluggable-optics-lpo-market-7085

51. ZerolineAI, "Fiber Optic Cable Prices in 2026" (citing CRU, August to October 2026). https://connect.zerolineai.com/guides/fiber-optic-cable-prices/

52. Roc Telecom, "What Changed on October 13, and What Didn't: The FCC's Chinese Optical Module Rules". https://roctelecom.com/insights/fcc-optical-module-rules-update/

53. MLQ.ai, "FCC Weighs Ban on New Chinese Optical Transceivers Used in AI Data Centers," 2026. https://mlq.ai/news/fcc-weighs-ban-on-new-chinese-optical-transceivers-used-in-ai-data-centers/

54. DataCentreNews UK, "AI Fuels Europe Data Centre Boom Amid Power Crunch," February 5, 2026. https://datacentrenews.uk/story/ai-fuels-europe-data-centre-boom-amid-power-crunch

55. Data Center Knowledge, "Europe's Data Center Market Enters a Pivotal Phase Amid Structural Challenges" (EUDCA report). https://www.datacenterknowledge.com/data-center-site-selection/europe-s-data-center-market-enters-a-pivotal-phase-amid-structural-challenges

56. JLL, "AI Boom Redraws Europe's Data Centre Map as Power Access Overtakes Location". https://www.jll.com/en-de/newsroom/ai-boom-redraws-europes-data-centre-map-as-power-access

57. Data Centre Review, "AI Boom Pushes Hyperscale Data Centres Away from Europe's Major Hubs," August 2026. https://datacentrereview.com/2026/08/ai-boom-pushes-hyperscale-data-centres-away-from-europes-major-hubs/

58. European Commission, Digital Skills and Jobs Platform, "EU AI Gigafactories Call Open Until 12 November," 2026. https://digital-skills-jobs.europa.eu/en/latest/news/eu-ai-gigafactories-call-open-until-12-november-targeting-over-eu30-billion-investment

59. IndexBox, "Germany Fiber Optic Connectivity Market Analysis". https://www.indexbox.io/store/germany-fiber-optic-connectivity-market-analysis-forecast-size-trends-and-insights/