BAAC INNOVATION

DARK FIBER FOR AI INFRASTRUCTURE

WHITE PAPER

FIBER BEFORE WATTS

Sizing Dark Fiber Capacity to the Power and Compute Density of AI Data Centers

Prepared by BAAC Innovation Inc.

Dark Fiber & Long-Haul Infrastructure for AI, Cloud, and Neocloud Data Centers

EXECUTIVE SUMMARY

Every new wave of AI infrastructure investment starts with the same two questions: how much power can this site deliver, and how many GPUs can that power support? Far fewer teams ask the third question early enough to matter: how much fiber does that GPU footprint require to move data between racks, buildings, and campuses without becoming the bottleneck the rest of the build was designed to avoid?

Hyperscalers have been answering that third question for two decades. Their network teams size dark fiber routes, duct banks, and splice architecture years ahead of first power, because they have already lived through the cost of guessing wrong. Neocloud operators and AI-native data center developers, by contrast, are frequently building their first or second campus. Many are sizing fiber the way earlier data center generations sized it: reactively, after the power and mechanical plans are locked, and often assuming that lit services or a single wavelength system will absorb whatever growth comes next.

That assumption breaks down quickly in AI infrastructure. GPU-to-GPU training traffic between campuses scales with the number of accelerators in the cluster, not with the number of racks or square feet on a site plan, and it can move from hundreds of gigabits to tens of terabits per site as a build progresses through a handful of expansion phases.

This paper lays out why dark fiber capacity planning belongs in the same early design conversation as power and cooling, why the build-out should be pre-positioned on day one as a line item in the data center capital plan, why dense wavelength division multiplexing (DWDM) is necessary but not sufficient on its own for every hyperscale build, why banking spare fiber strands for future growth is cheaper than it looks, and how the right fiber strategy becomes a lever for energy and compute efficiency rather than just a connectivity checkbox. It also introduces the specific class of dark fiber these requirements demand — what BAAC calls Neofiber — and explains how it differs from the long-haul fiber built for the cloud era.

A second argument runs alongside the capacity case and is examined in Section 3: physical route control is now a strategic asset in its own right. Recent terrestrial and subsea incidents have shown repeatedly that contractual redundancy fails when both paths share the same ground, and the largest infrastructure buyers have responded by acquiring and routing their own physical plant rather than relying on service-level guarantees over someone else's.

A third, examined in Section 7, is architectural. The physical interconnect decision propagates upward: it determines which optics are appropriate, whether end-to-end jumbo frames can be guaranteed across the overlay, and whether MACsec — the layer-2 encryption NVIDIA recommends for data center

interconnect — is available at all. Capabilities inherited from the physical layer cannot be configured back in above it.

1. THE NEW BANDWIDTH REALITY: AI CAMPUSES VS. TRADITIONAL HYPERSCALE

Conventional enterprise and hyperscale data centers grew their network requirements in a fairly linear, predictable way: more racks, more servers, a proportional increase in north-south and east-west traffic. Capacity planners could forecast fiber needs years out with reasonable confidence because the underlying compute density per rack changed slowly.

AI training and inference clusters do not follow that curve. The unit of measure is no longer the server or the rack, but the scalable unit, or SU — a pod of tightly coupled GPUs, such as an NVIDIA GB300 NVL72 system, that must communicate with its peers at extremely high bandwidth to synchronize gradients during training or to serve large models during inference. As a campus adds SUs, inter-site bandwidth requirements do not creep upward — they step upward, often by an order of magnitude within a handful of build phases.

A single 24-fiber bundle that comfortably serves a first-phase AI campus can be fully saturated a year later, not because the site grew in square footage, but because it grew in GPUs per rack.

At the low end, a single-SU deployment may need only a few hundred gigabits of inter-site bandwidth — comfortably inside a first-generation build. By the time a site reaches the range now common among aggressive AI infrastructure buyers, requirements move into the multi-terabit range, and at true hyperscale AI campus scale, aggregate inter-site demand can exceed 100 Tbps. Few enterprise network plans written five years ago anticipated numbers in that range, which is exactly the point: AI bandwidth growth outpaces the planning cycles that most data center capital projects still use.

Three figures capture what AI now demands of the fiber layer, and each one sits well outside the envelope legacy long-haul was designed around:

- **1.6 Tbps** — per-link bandwidth needed for next-generation AI clusters, against 400 Gbps mainstream today.
- **Under 10 ms** — inter-data-center latency requirement for distributed training and inference workloads.
- **864+ strands** — strand counts on modern AI corridors, against 24–48 strands for legacy long-haul builds.

The conclusion those numbers point to is not that AI infrastructure is bandwidth-limited. It is that AI infrastructure is route-limited. Capacity can be added to glass that already exists; physical paths into power-rich, network-poor geography cannot.

2. NEOFIBER: A DIFFERENT CLASS OF DARK FIBER

Not all dark fiber is built to the same brief. The long-haul networks that carry most intercity traffic today were engineered for the cloud era, when connectivity rewarded density: modest strand counts run between markets where carrier networks already met, following population and existing interconnection rather than raw power availability. Those routes did exactly what they were designed to do, and they remain well suited to the workloads they were built for.

AI infrastructure inverts almost every input to that design. Its siting is driven by power, water, and land — which pushes campuses into corridors legacy long-haul largely skipped. Its bandwidth curve is steep and step-shaped rather than gradual. Its latency budget is tight enough that route geometry, not just route existence, determines whether a distributed training job is viable. And its growth trajectory is known in advance from the power roadmap, which makes provisioning to today's demand a self-inflicted cost.

What BAAC means by Neofiber

Neofiber is the class of dark fiber BAAC designs and builds for AI infrastructure: high-strand-count, low-latency, route-diverse long-haul engineered into the power-first sites AI occupies, rather than the already-networked metros legacy long-haul was built to connect.

In practice that means a route structure of four primary conduits, each carrying two 7-way microduct bundles sized for high fiber-count cable; physically independent corridors — including pipeline right-of-way conversions — rather than a second circuit in a shared duct bank; and capacity banked at construction for expansion phases the power roadmap has already signalled.

Everything that follows in this paper — the strand-count math, the day-one capex argument, the limits of DWDM, and the case for fiber banking — describes how a Neofiber route is sized and staged.

Design Attribute	Legacy Long-Haul (Cloud Era)	Neofiber (AI Era)
Typical strand count	24–48 strands	864+ strands
Route selection driver	Market density — connecting metros where networks already met	Power availability — reaching the power-first sites AI occupies
Per-link bandwidth target	100G–400G mainstream	1.6T for next-generation AI clusters
Latency requirement	Best-effort within regional SLAs	<10 ms inter-data-center for distributed training and inference
Diversity posture	Second circuit, often in the same duct bank	Physically independent corridors as a structural requirement

Design Attribute	Legacy Long-Haul (Cloud Era)	Neofiber (AI Era)
Capacity philosophy	Provision to current demand	Bank strands and duct for several expansion phases ahead

Figure 1: Legacy long-haul versus Neofiber design attributes. Strand counts, per-link bandwidth targets, and latency thresholds reflect current AI corridor practice and published cable and Ethernet roadmap references.

The distinction matters commercially as well as technically. An operator who procures fiber on cloud-era assumptions — a modest strand count on whichever route was already available — will typically hit a capacity or diversity wall within one or two expansion phases, at which point the remedy is a second civil construction project rather than a configuration change. An operator who procures on Neofiber assumptions pays a marginal premium once, at first construction, and converts every subsequent expansion into a splicing and optics exercise.

3. ROUTE SOVEREIGNTY: THE PHYSICAL PATH AS A STRATEGIC ASSET

The case for Neofiber is usually made on capacity. There is a second case, made increasingly often by the buyers with the most operating experience, that has nothing to do with bandwidth: whoever controls the physical path controls the risk. Over the past two years that argument has moved from a theoretical resilience discussion to a set of documented, expensive incidents on both land and seabed.

3.1 Logical redundancy is not physical diversity

The clearest recent illustration is terrestrial and North American. On September 19, 2025, a third-party contractor working near Argyle, Texas accidentally cut fiber serving the FAA's Dallas TRACON facility. Two cables were severed, and by American Airlines' account the cut took out both the primary and the secondary data paths supporting the area's radars, radio frequencies, and computer systems. Ground stops followed at Dallas/Fort Worth International and Dallas Love Field, roughly a fifth of flights were cancelled, and normal operations did not fully return for several days. The FAA also noted that a contractor had failed to confirm the system's redundancies were actually functioning.

The engineering lesson generalizes well beyond aviation. A protected circuit, a diverse-path SLA, and a second wavelength are all logical constructs; if both paths ultimately traverse the same duct bank, the same right-of-way, or the same contractor's excavation zone, the redundancy exists on paper and not in the ground. Days later that same month, mass 911 outages across parts of Mississippi, Louisiana, and Alabama were likewise attributed by AT&T to fiber cuts made by third parties. Neither incident involved a sophisticated adversary. Both involved ordinary excavation.

For an AI campus, the equivalent failure is not a delayed flight but an interrupted training run: a synchronized job spanning two sites does not gracefully degrade when the interconnect drops, and the compute is billed whether or not the network was available to use it.

3.2 Subsea and geopolitical exposure

The same pattern is visible internationally, at higher stakes. In November 2024, the BCS East-West Interlink and the C-Lion1 cable were disrupted in the Baltic Sea within roughly a day of each other, prompting hybrid-warfare and sabotage accusations from European and NATO officials; on December 25 of that year the Estlink 2 power cable between Finland and Estonia failed along with four telecommunications lines, and Finnish authorities seized a tanker on suspicion its anchor caused the damage. NATO launched a dedicated Baltic Sea patrol mission, Baltic Sentry, in January 2025. Around Taiwan, authorities have recorded a comparable run of incidents — at least eleven cases of undersea cable damage since 2023 by Taiwanese and European tallies — and Taiwan's Ministry of Digital Affairs found that anchor dragging accounted for the largest share of nearshore faults, with foreign flag-of-convenience vessels emerging as a distinct category of threat.

It is worth stating the base rate plainly, because it cuts against alarmism while strengthening the engineering argument. Cable faults are routine: industry bodies count roughly 100 to 200 globally each year, and United Nations telecommunications officials attribute around 80 percent of recorded damage to fishing and anchoring rather than deliberate action. Most physical-layer failure is mundane. That is precisely why route diversity is an engineering baseline rather than a geopolitical precaution — the network has to survive the ordinary case, and the ordinary case is frequent.

3.3 What the most experienced buyers are doing about it

The hyperscalers, who have absorbed two decades of these failures, have responded by buying the physical layer outright rather than renting resilience from it. Google has deployed multiple wholly owned private subsea cables, and Meta announced Project Waterworth in February 2025 — a roughly 50,000 km, 24-fiber-pair system spanning five continents, of which Meta will be sole owner. Reporting on the project notes that its routing was chosen specifically to avoid the Red Sea, the South China Sea, and the Strait of Malacca as concentrated single points of failure. Sector-wide, TeleGeography projects roughly $13 billion of subsea cable investment across 2025–2027, close to double the prior three-year period, with AI infrastructure demand cited as the primary driver.

Meta's own network investment leadership has framed the rationale bluntly: data centers without the connectivity that links them are simply very expensive warehouses. That is the same conclusion this paper reaches from the strand-count math, arrived at from the opposite direction.

The implication for neocloud and AI-native operators is not that they should be laying their own transoceanic cable. It is that the buyers furthest up the learning curve treat physical route ownership and route independence as strategic assets, not procurement line items — and that an operator without that scale can obtain the same property by contracting a carrier-neutral route built on genuinely independent corridors.

3.4 The right-of-way concentration problem

There is a structural reason genuine diversity is hard to buy in North America, and it is worth understanding before signing a diverse-path agreement. The national long-haul network is built overwhelmingly on two categories of right-of-way: railroad corridors and state or federal highway corridors. A large share of it was placed during the 1999–2001 build cycle, in small-count duct structures that are now fully occupied — which means the incumbent carrier frequently cannot augment its own route without new construction, whatever the customer is willing to pay.

The consequence is that two carriers' "diverse" routes between the same pair of markets often occupy the same rail corridor or the same highway shoulder, sometimes the same duct bank. Diversity purchased on paper resolves, on inspection, to shared physical exposure — the Dallas failure mode at national scale. Sophisticated buyers have noticed, and hyperscale customers are actively seeking routes on rights-of-way diverse from the legacy national providers, along with new entrants offering vacant conduit that can scale.

This is the structural gap BAAC's model addresses. Pipeline, utility, and industrial corridors are physically separate from rail and highway alignments, generally deeper-buried, and carry fewer intercepts and splice points across a long-haul run. Where the corridor owner controls the right-of-way directly, permitting to place future cable is materially simpler than on Department of Transportation or railroad land — which turns the banking strategy described in Section 8 from an aspiration into something an operator can actually execute later.

4. FROM POWER CAPACITY TO FIBER STRAND COUNT

Power capacity and fiber capacity are two expressions of the same underlying variable: how much AI compute a site is meant to run. As data center power density increases — measured in megawatts per building or per hall — the GPU count that power supports increases with it, and so does the inter-site bandwidth those GPUs require for training and inference traffic. A fiber plan that is not derived from the power roadmap is, in practice, a guess.

That relationship has not always held. Earlier generations of power-intensive facilities — cryptocurrency mining being the clearest example — consumed enormous power while generating very little network traffic, so megawatts and bandwidth were effectively decoupled. AI has re-coupled them. As sites move from energy-intensive, low-bandwidth workloads through AI training, which is energy-intensive with low to moderate bandwidth, to AI inference, which is both energy-intensive and bandwidth-intensive, megawatts deployed converge on network demand. Industry estimates now put roughly two-thirds of national compute on inference against one-third on training, with the inference share growing — and inference is the more network-hungry of the two, absorbing spikes in user requests alongside backup, storage, and workload distribution traffic.

4.1 A planning heuristic: megawatts to fiber pairs

The practical consequence is that a fiber requirement can be estimated directly from a power figure, early enough to influence the route design. A widely used planning heuristic sizes roughly one fiber pair of production traffic per 20 MW of IT load. That figure alone understates the requirement, because no operator runs a campus on a single path. Dual-path redundancy is the floor; most hyperscalers require triversity — three physically separate paths — and quadversity is increasingly requested at gigawatt scale.

Triversity has a specific arithmetic consequence worth internalizing. Where one production pair is protected by two additional pairs allocated for failover, every three fiber pairs of raw supply support only one pair of production traffic. Raw fiber demand is therefore two to four times the production figure, depending on the diversity posture the operator's customers or workloads require.

Campus IT Load	Production Fiber Pairs (@ ~20 MW / FP)	Raw Pairs — Dual Path (2×)	Raw Pairs — Triversity (3×)	Raw Pairs — Quadversity (4×)
100 MW	5	10	15	20
250 MW	12.5	25	37	50
500 MW	25	50	75	100
1 GW	50	100	150	200
2.5 GW	125	250	375	500
5 GW	250	500	750	1,000

Figure 2: Fiber pair requirements derived from campus IT load at the three common diversity postures. Production pairs are calculated at approximately one pair per 20 MW; raw supply is the production figure multiplied by the number of physically separate paths required.

Plotted, the same figures make the planning point directly: a gigawatt-class campus operating to a hyperscaler's standard diversity posture needs raw fiber supply in the range of 150 pairs — 300 strands — before any allowance for growth, spare capacity, or the pass-through traffic a route may also carry.

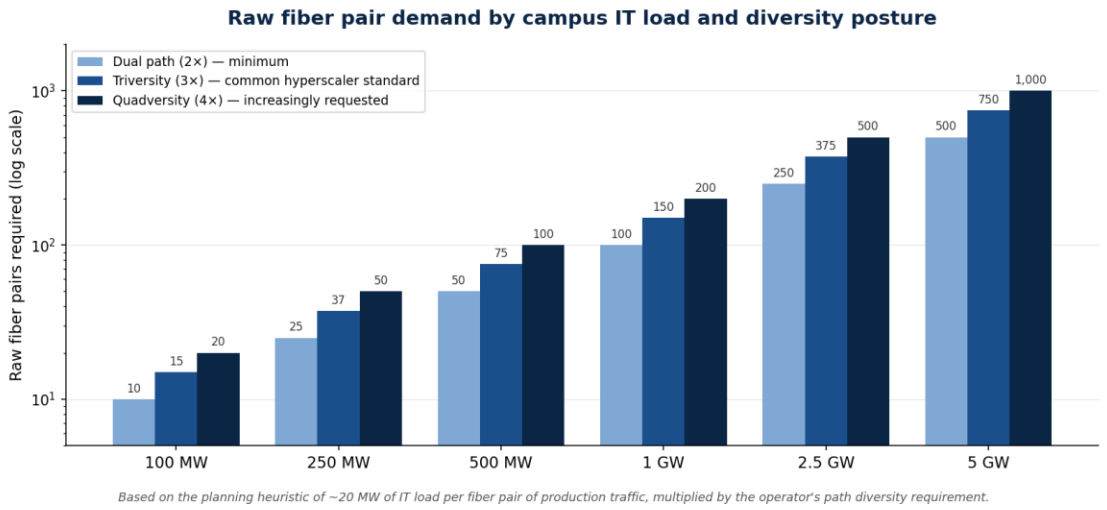


Figure 3: Raw fiber pair demand by campus IT load and diversity posture, logarithmic scale.

This heuristic is deliberately coarse. Its value is that it can be applied at the moment a site's power roadmap is drafted — before rack layouts, before workload mix is settled, and critically, before route selection is finalized. It converts a number the developer already knows into a number the fiber design requires.

4.2 Strand count against GPU scalable units

The megawatt heuristic and the scalable-unit model answer different questions, and they should not be added together or read as competing estimates. The heuristic in Section 4.1 sizes the route portfolio — total fiber a campus should hold across all paths, covering east-west and north-south traffic alike, at the coarse resolution available when only a power figure exists. The scalable-unit model below sizes one thing precisely: the synchronous GPU-to-GPU interconnect between campuses running a distributed training job. Because that traffic class is the most demanding of the several a route carries, the scalable-unit figures at hyperscale run several times higher than the heuristic implies for the same megawatt figure. Where both apply, plan to the larger of the two.

The table below translates GPU scalable-unit counts into approximate inter-site bandwidth, minimum strand counts using 400GbE DWDM optics, and the practical fiber bundle size needed to support each stage of growth, from a single-SU pilot through a 500+ SU hyperscale AI campus footprint.

SUs / Site	GPUs / Site	Est. Inter-Site Bandwidth	Min. Strands (400GbE DWDM)	Practical Fiber Bundle
1	72	~200 Gbps	2 strands	12F minimum
2–4	144–288	~400–800 Gbps	4–8 strands	12–24F
5–8	360–576	~1–1.6 Tbps	10–16 strands	24F (saturates)
9–16	648–1,152	~1.8–3.2 Tbps	18–32 strands	48F / DWDM upsell

SUs / Site	GPUs / Site	Est. Inter-Site Bandwidth	Min. Strands (400GbE DWDM)	Practical Fiber Bundle
32–64	2,304–4,608	~6–13 Tbps	64–128 strands	96F or 144F multi-tube
128	9,216	~26 Tbps	~200+ strands	288F multi-tube
500+	36,000+	~100+ Tbps	1,000+ strands	864F+ / multi-route

Figure 4: Illustrative fiber strand and bundle requirements by GPU scalable-unit count, based on GB300 NVL72 inter-site bandwidth modeling. Figures assume AllReduce-style gradient synchronization for large-model training, using a conservative estimate of approximately 200 Gbps of inter-site bandwidth per scalable-unit (SU) pair. Strand counts are calculated against 400GbE optics, today's mainstream rate; as 800G and 1.6T coherent optics reach volume deployment, the same bandwidth will be carried on proportionally fewer strands, so these counts should be read as a conservative upper bound on the optics side and a floor on the civil side. Fiber bundle sizes are per diverse route, not summed across routes.

A clarification before reading the thresholds: the strand counts below describe fiber lit for inter-site traffic at each stage, not the cable an operator should place in the ground. Those are different numbers by design. Lighting 24 strands today on a route built with 864-strand cable is the strategy this paper recommends; building a 24-strand route because 24 strands are currently needed is the failure mode it warns against.

Three thresholds in this table matter more than the rest, because they are the points where an under-planned fiber build forces an expensive mid-project correction:

- **5–8 SUs per site:** a standard 24-fiber bundle running 400GbE DWDM begins to saturate. This is the first upgrade trigger, and it arrives sooner than most first-time builders expect.
- **9–16 SUs per site (~3 Tbps):** the next meaningful threshold, where operators typically move to a 48F build or multiply existing strand capacity through DWDM rather than pulling a second bundle.
- **32–64+ SUs per site:** multi-route territory, where a single bundle — however large — stops being a sufficient answer on its own.

Plotted on a logarithmic scale, the shape of the problem becomes clearer than any table can convey: the curve is not a gradual ramp with occasional steps, but a series of order-of-magnitude jumps punctuated by the three upgrade triggers above.

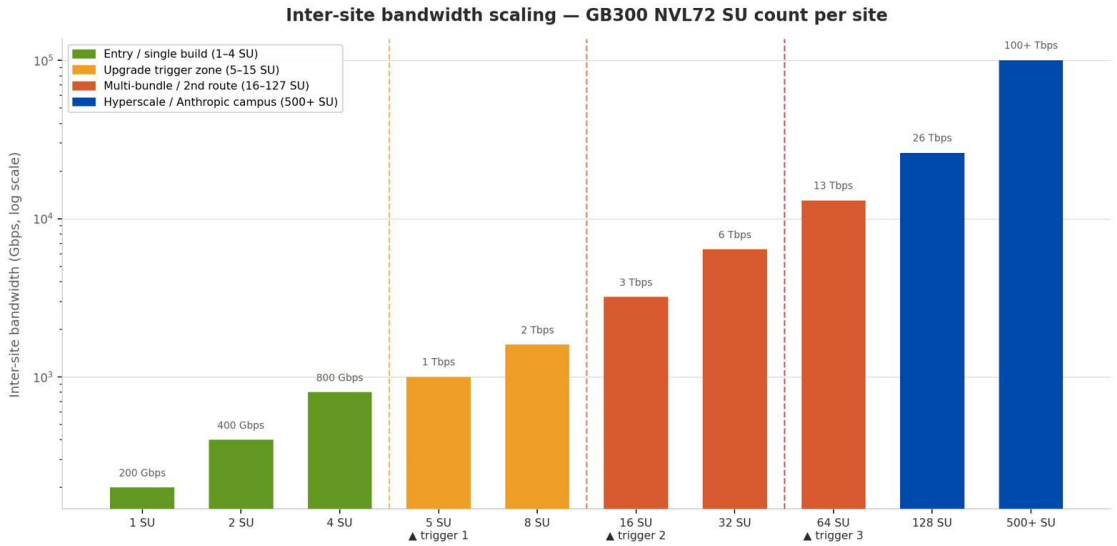


Figure 5: Inter-site bandwidth scaling by scalable-unit count per site, logarithmic scale. The three dashed lines mark the upgrade triggers where the installed fiber plant, rather than the optics, becomes the constraint.

Route diversity matters as much as raw strand count. A single physical route, however many strands it carries, remains a single point of failure — one excavation, one fire, one right-of-way dispute can take an entire campus offline regardless of how much fiber is in that conduit. This is the physical expression of the diversity multipliers in Figure 2: dual path is the floor at any scale, triversity becomes the practical standard as planned IT load passes roughly 50 MW, and at gigawatt-class campuses a mesh of four or more diverse long-haul routes is a more realistic target than any single corridor, however large.

These figures are planning-level estimates intended to size a route portfolio early, not a substitute for a site-specific traffic and power delivery study. Actual requirements will vary with rack-level power density, cooling architecture, and the specific mix of training versus inference workloads a given campus is built to run.

4.3 Two traffic flows, two route requirements

Sizing capacity is only half the design problem; the other half is direction. AI campuses generate two distinct classes of traffic, and they impose different requirements on the route portfolio.

- **East-west, machine to machine.** Data center to data center traffic — distributed compute jobs, redundant storage, replication, and metadata access. Operators overwhelmingly want this carried entirely on a private backbone, never crossing a security boundary onto the public internet. This is the traffic that the scalable-unit math in Section 4.2 sizes, and it is one to two orders of magnitude larger than end-user traffic flows.
- **North-south, data center to edge.** Application delivery, authentication, security, and content distribution reaching actual users. Operators want this on the private backbone for as much of its

path as possible, handing off to the public internet or directly to the network serving the user as geographically close to that user as they can manage.

The second flow is what makes interconnection geography a route-design input rather than an afterthought. Peering is heavily concentrated: a small number of exchange points account for a disproportionate share of network interconnection in any given region, and the largest carry hundreds of peered networks each. AI campuses sited for power are frequently a long way from those points. A route portfolio therefore has to satisfy two different criteria at once — short, high-capacity, diverse paths between campuses for east-west traffic, and multiple redundant paths toward the region's major interconnection hubs for north-south traffic.

Both requirements point the same direction on strand count and diversity, which is convenient. They do not necessarily point the same direction on route geography, which is why the portfolio has to be designed rather than assembled one circuit at a time.

5. WHY FIBER BUILD-OUT BELONGS ON DAY ONE

The economics of fiber construction are dominated by civil work, not glass. Trenching, boring, permitting, and restoration make up the majority of the cost of a fiber route; the incremental cost of adding extra ducts or a higher fiber count cable to an already-open trench is comparatively small. That asymmetry is the entire argument for over-building conduit and duct structure at first construction, rather than provisioning exactly what current demand requires.

Once a route is buried and restored, going back to add capacity means repeating nearly the full civil cost a second time — new permits, new traffic control, new restoration — to gain a fraction of the value achieved the first time the ground was open. For an AI campus where the power roadmap already signals two, three, or four future expansion phases, that second mobilization is a predictable, avoidable cost rather than an unforeseen one.

Lead times compound the argument, and they now extend well past the fiber itself. Cable delivery lead times have lengthened to as much as a year under AI-driven demand, and new glass preform manufacturing capacity takes roughly 18 to 24 months to bring online — meaning the industry cannot respond quickly to a demand spike even in principle. New long-haul routes typically require three to five years to permit — a timeline that sits awkwardly against AI capacity roadmaps measured in quarters.

Amplifier sites are the constraint most often missed. In-line amplifier huts built two decades ago were designed for lower power density and are now largely full, and unsuitable for current-generation DWDM amplification equipment. Modern deployments draw on the order of 5 kW — roughly 100 amps DC — per rack, and because these sites sit in remote locations, power augments routinely take twelve months or longer. An operator who has secured fiber but not amplifier land and power has not secured a usable route.

A fiber decision deferred to a later phase is therefore not simply the same decision made later; it is a decision made against a tighter supply market, a longer permitting queue, an amplifier siting problem, and a shorter runway. The counter-move is to reserve cable and complete amplifier site planning ahead of firm commitments, trading upfront capital exposure for schedule certainty.

This is why dark fiber deserves a line item in the same capital planning conversation as power delivery and mechanical cooling, not a follow-on request submitted once a bandwidth ceiling is already in sight. A standard design of four primary conduits per route, each carrying two 7-way microduct bundles sized for high fiber-count cable, lets an operator light only what current demand requires while banking the rest — turning future capacity into a design decision made once, instead of a construction project repeated every time the cluster grows.

6. WHY DWDM ALONE DOESN'T SOLVE THIS FOR ALL HYPERSCALE BUILDS

Dense wavelength division multiplexing is one of the most effective tools available for extracting more capacity from existing fiber: by sending many independent wavelengths, or colors of light, down a single strand pair, a single fiber pair can carry many times its native single-wavelength capacity without any new construction. For a large share of AI infrastructure growth, DWDM is exactly the right first response to a bandwidth ceiling.

It is not, however, a substitute for physical strand count at every scale, for three reasons that matter specifically to AI campus planning:

- **DWDM has a ceiling too.** A given fiber pair supports a finite number of usable wavelengths at a given per-wavelength rate. Once that ceiling is reached, the only paths forward are a new fiber pair, a new route, or a more expensive optical line system upgrade — DWDM defers the need for more strands, it does not eliminate it.
- **Physical route diversity cannot be multiplexed into existence.** AI clusters at meaningful scale need more than one physical path between sites for resilience; no amount of wavelength multiplexing on a single fiber pair protects a training job from a single backhoe strike on that route. The Dallas TRACON outage described in Section 3 is the canonical demonstration — a single excavation took out both the primary and secondary paths at once.
- **Cost structure still matters at the extremes.** Managed wavelength and carrier Ethernet services scale in discrete, often expensive steps as bandwidth grows past a few hundred gigabits, while dark fiber's cost is tied to distance rather than bandwidth — a difference in economics, not just engineering, that becomes more pronounced the larger the cluster gets.

The practical conclusion is to design for DWDM from day one on every route, while sizing physical duct and strand capacity for the multi-route, multi-terabit endpoint the power roadmap implies — not for whatever wavelength count looks sufficient today.

7. THE PHYSICAL LAYER DETERMINES THE ARCHITECTURE ABOVE IT

Everything argued so far concerns capacity and route selection. There is a further consequence of the fiber decision that tends to surface only after a network is in service: the physical interconnect choice constrains what the network architecture above it is permitted to do. NVIDIA's own data center interconnect guidance is explicit that the two dominant physical-layer strategies for connecting data centers are dark fiber interconnect and DWDM interconnect, and that dark fiber is what provides network redundancy through path diversity on each route — the same conclusion this paper reaches from the civil-construction and resilience side.

7.1 Termination, optics, and distance

Once a physical strategy is chosen, it propagates into equipment decisions. DCI links can be terminated directly on the border leaf switches that form part of an EVPN or VXLAN fabric, or on an additional layer outside the fabric attached to those border leaves; the right choice depends on existing network design, routing architecture, and physical constraints. The architecture then determines which pluggable optics are appropriate on the border devices, largely as a function of how far the signal must travel.

Deployment Scenario	Typical Optic Class	Design Note
Provider equipment in same facility or on same floor	SR4 / DR4 / FR4	Use the simplest, most cost-effective option; distance is not the constraint
Provider PoP not in close proximity to the campus	LR4 / IR4	Long-distance optics required to reach the interconnection point
Dark fiber interconnect between campuses	LR4 / IR4 / ER4 / ZR4	Optic class selected against actual route distance
Dark fiber with multiplexing for capacity growth	Colored or gray optics over CWDM / DWDM	Joins 100G connections over a single fiber pair; choice depends on xWDM equipment, cost, and timeline

Figure 6: Optic class selection by interconnect scenario, following NVIDIA DCI guidance. Where a dark fiber route requires more capacity than the installed optics provide, a CWDM or DWDM layer can join 100G connections over a single fiber pair, with colored or gray optics selected according to the xWDM equipment in use and the operator's cost and schedule constraints.

The operationally important point is that a campus which owns or holds an IRU on its fiber selects optics against the actual distance of its own route. A campus buying a managed service inherits whatever the provider's point of presence location dictates — which is frequently a longer optical path than the geography requires, because the traffic must reach the provider's PoP before it reaches its destination.

7.2 MTU, overlay overhead, and third-party dependencies

Overlay technologies such as EVPN and VXLAN add encapsulation overhead to the data plane. The VXLAN header itself is only eight bytes, but the encapsulated frame preserves the original layer 2, layer 3, and layer 4 headers inside the new tunnel frame and adds its own destination addressing and UDP header — for a total overhead of 50 bytes. Avoiding host-side and network-side fragmentation therefore requires increasing the end-to-end MTU, which local data centers normally accomplish with jumbo frames.

That requirement becomes a dependency the moment traffic crosses someone else's domain. Traffic traversing a third-party layer 2 or layer 3 service must have the same end-to-end jumbo frame capability, or the interconnect experiences bottlenecks and payload mismatch problems. NVIDIA's guidance is blunt on this point: whatever the differences between ELINE and ELAN service types, jumbo-frame support is the single most important criterion when selecting among third-party vendors.

For an AI campus, MTU is not a minor tuning parameter. It sits directly on the path of the synchronization traffic described in Section 4, and a fragmentation problem introduced by a provider's frame-size ceiling costs GPU throughput in exactly the way described in Section 9. On owned or IRU fiber, MTU is a design decision. On a leased service, it is a procurement question — and one that has to be re-asked of every provider on every segment.

7.3 Encryption and the layer-2 boundary

Security compounds the same asymmetry. Confidentiality on a DCI link can be implemented with IPsec at layer 3, typically in firewall hardware rather than switching hardware, or with MACsec at layer 2. NVIDIA identifies MACsec as its recommended implementation for DCI security — but MACsec-encrypted traffic cannot cross a layer 2 boundary and cannot be routed across an IP network. It works on DCIs built over dark fiber or CWDM/DWDM infrastructure.

The implication is worth stating plainly, because it is rarely surfaced during procurement: the vendor-recommended encryption posture for data center interconnect is only available to operators who control the physical layer. An operator whose inter-site traffic rides a routed third-party service is architecturally excluded from it and must fall back to a layer 3 approach with different performance and hardware implications. For AI operators moving proprietary model weights, training checkpoints, and customer data between campuses, that is a material difference in available security architecture, decided years earlier by a fiber procurement choice.

Capability	Third-Party Ethernet Service (ELINE / ELAN)	Carrier DWDM Wavelength	Owned / IRU Dark Fiber (Neofiber)
Physical path diversity	Provider-defined; not visible to customer	Limited to provider's route set	Customer-selected per route

Capability	Third-Party Ethernet Service (ELINE / ELAN)	Carrier DWDM Wavelength	Owned / IRU Dark Fiber (Neofiber)
End-to-end jumbo frame / MTU control	Dependent on provider support	Generally supported	Set by the operator
MACsec layer-2 encryption	Not viable across routed domains	Supported	Supported
Capacity scaling method	New service order per increment	Add wavelengths to existing pair	Add wavelengths, then strands, then routes
Optics and topology choice	Provider-controlled	Partially controlled	Fully operator-controlled

Figure 7: Interconnect options compared across the capabilities that AI data center interconnect depends on. Path diversity, MTU control, and MACsec availability are all properties inherited from the physical layer rather than configured above it.

Read alongside Figure 4, this is the same argument arriving from a second direction. Strand count determines how much traffic a campus can carry; the physical layer choice determines what the operator is permitted to do with it.

8. FIBER BANKING: DESIGNING FOR FUTURE GROWTH

Fiber banking means installing more strands, and more duct capacity for future strands, than current lit demand requires — and treating the unlit portion as a deliberate reserve rather than idle inventory. For AI campuses whose power roadmap already anticipates several future expansion phases, banked fiber is the mechanism that keeps each of those phases from requiring new civil construction.

Two market conditions make banking more attractive today than it was even a few years ago, and they pull in opposite directions on cost. The first lowers the barrier: commercial terms around dark fiber leasing have loosened, and indefeasible right-of-use (IRU) agreements that once ran 20 to 30 years are now available in terms as short as five years in many markets, reducing the financial commitment required to secure capacity ahead of need. The second raises the stakes: as Section 3.4 and Section 10 describe, the long-haul market is consolidated and supply-constrained, with lead times measured in quarters and hyperscale buyers actively competing for the same routes. Banking is attractive precisely because capacity secured in advance is cheaper and faster than capacity sought in a tight market — not because the market has become easy.

A well-designed Neofiber route — BAAC's standard structure of four primary conduits, each holding two 7-way microduct bundles sized for high fiber-count cable — allows an operator to light a fraction of total capacity now, while the remainder sits banked and ready. The marginal cost of that reserve, purchased once during initial construction, is a small fraction of what a second trenching project would cost later. Banked capacity is, in effect, an option on future growth purchased at construction-cost pricing rather than emergency-project pricing.

8.1 Banking is not only about glass

Three assets beyond the fiber itself belong in the same reserve, and each is routinely underestimated.

Spare microduct is the cheapest form of technological insurance available. Fiber installed underground today has a usable life beyond twenty years and plausibly beyond thirty. Over that horizon, hollow-core fiber or other glass types may reach price points and cable sizes that make them worth deploying. An operator holding a vacant microduct can install that cable without removing anything; an operator with fully occupied conduit faces a rebuild. Keeping at least one microduct empty is a deliberate decision to leave room for a technology that does not exist yet.

Amplifier land is the second. Rack requirements per fiber pair are genuinely uncertain, because two technology trends pull in opposite directions: amplification advances increase the number of fiber pairs a rack can serve, while C+L band DWDM increases the equipment required per pair. The defensive planning assumption is roughly eight fiber pairs per rack. On a route built to 864-pair capacity, that implies land for modular amplifier structures supporting on the order of 108 racks — a figure that has to be secured as land, at construction, whether or not it is built out.

Power service capacity is the third, and it follows directly. Sizing the electrical service to a full light-up scenario while energizing incrementally avoids the twelve-month augment cycle described in Section 5 at precisely the moment an operator is trying to turn up new capacity quickly.

Asset	Why It Must Be Sized on Day One	Practical Standard
Conduit and duct	Civil cost dominates; reopening the trench repeats nearly the full expense	Four primary conduits per route
Microduct	Preserves a pathway for future glass types without removing installed cable	Two 7-way bundles per conduit; keep at least one spare
Fiber strands	Cable lead times run up to a year; bulk purchase lowers per-foot cost	Size to the power roadmap's final phase, not phase one
ILA land	Amplifier huts are land- and permission-bound; adding sites later takes 12 months or more	Allocate for roughly 8 fiber pairs per rack across full route capacity
ILA power	Modern amplification draws roughly 5 kW (100 A DC) per rack; augments are slow in remote locations	Plan service capacity for full light-up, energize incrementally

Figure 8: What a fiber banking strategy actually reserves. Each asset shares the same economics — cheap to secure during initial construction, slow and expensive to add once the route is in service.

9. ENERGY AND COMPUTE OPTIMIZATION ENABLED BY DARK FIBER

Fiber strategy is not only a connectivity question — it is increasingly an energy and compute efficiency question. As data center power availability becomes the binding constraint on AI growth in many markets, operators are increasingly distributing training and inference workloads across multiple sites rather than concentrating them behind a single utility interconnection. That strategy is only viable if the network between sites can move data at the volume and latency the workload requires; otherwise, distributing compute across sites simply relocates the bottleneck from the power meter to the network port.

9.1 Multi-site training is already operational practice

This is no longer a forward-looking scenario. Google has stated that its Gemini models were trained on TPU pods distributed across multiple data centers, using model parallelism within pods and data parallelism across them, with network latency and bandwidth sufficient to sustain synchronous training. Google has publicly described distributing training partly for resilience — so that a power problem at one site does not halt the run — and partly because space and power constraints at any single site cap how large a cluster can be. In April 2026, Google announced network architecture intended to connect more than one million TPUs across multiple data center sites into a single training cluster, alongside support for up to 960,000 GPUs across multiple sites.

Independent analysis reaches the same conclusion from first principles. Epoch AI's October 2025 assessment of decentralized training concluded that it is technically feasible to distribute a multi-gigawatt-scale training run across dozens of sites separated by thousands of kilometers, and that distributing training beyond a single data center is a viable route around local power constraints — noting that training clusters powered by a single local grid interconnection are unlikely to exceed roughly 10 GW. Academic work on distributed training governance makes the driver explicit: power demands for the largest training runs already exceed 100 MW and are projected toward multiple gigawatts by 2030, and power constraints are expected to continue pushing labs toward multi-site training.

Every one of those architectures is a statement about fiber. Synchronous training across sites converts inter-site bandwidth and latency into a hard determinant of how much of the purchased GPU fleet is actually usable. The power roadmap sets how many accelerators a campus can host; the fiber plan sets whether accelerators in two campuses can work as one.

9.2 Removing network-induced idle time

Dark fiber's dedicated, unshared paths also remove a subtler inefficiency. Shared, packet-based transport introduces variable latency and jitter as it queues traffic from many customers across common infrastructure; AI training jobs that rely on tight synchronization between GPUs can lose meaningful throughput waiting on a congested or high-jitter link, effectively wasting compute cycles that were paid

for but not usefully used. A dedicated dark fiber path, engineered for the workload rather than shared across tenants, keeps those GPU cycles doing model work instead of waiting on the network.

Put together, dark fiber planned in step with the power roadmap turns connectivity into a lever that supports — rather than constrains — energy and compute optimization. It allows power-constrained operators to spread load across sites without a bandwidth penalty, and it removes network-induced idle time from GPU clusters, both of which translate into more delivered AI compute per megawatt of power procured.

10. MARKET VALIDATION: HOW THE LARGEST BUYERS ARE ACTING

The arguments in this paper are engineering arguments, but they are being confirmed commercially, in public, by the parties with the most complete view of AI network demand. Three data sets are worth putting in front of any operator still treating fiber as a late-stage procurement item.

10.1 Long-haul dark fiber demand has stepped, not drifted

Zayo's 2026 Bandwidth Report — drawn from purchasing behaviour across roughly 6,000 customers — found that demand for long-haul dark fiber doubled between 2024 and 2025 as its largest customers connected AI data centers, with long-haul fiber counts sold rising from 8,353 to 16,894 strands. Hyperscalers and carriers still account for about 95 percent of long-haul purchases, but Zayo identifies AI-native companies and neoclouds as an emerging class of buyer now purchasing dense fiber counts in both metro and long-haul. In April 2026, Zayo announced an anchor-customer agreement covering 8,000 route miles of new builds and overbuilds, explicitly framed as addressing a potential bandwidth gap, and describes AI-focused infrastructure investment spanning more than 15,000 route miles across North America.

Two details in that report matter for design rather than market sizing. First, two separate 2025 purchases by Zayo's largest hyperscale customers were at 864-fiber count — the same order of magnitude this paper recommends banking toward. Second, Zayo's own product leadership notes that customers increasingly want pre-built, ready-to-activate capacity with guaranteed intervals rather than bespoke engineered solutions, which is precisely the outcome that day-one over-build produces.

10.2 Carriers are rebuilding around pre-positioned capacity

Lumen reported roughly $13 billion in Private Connectivity Fabric agreements with hyperscalers, neoclouds, social platforms, and AI companies — including a purpose-built, high-capacity network for Anthropic — and is scaling from 17 million intercity fiber miles at year-end 2025 to a planned 47 million by 2028 and approximately 58 million by 2031. Its stated approach is instructive: pre-deploying high-capacity long-haul paths ahead of demand so customers can turn up 100G and 400G faster, securing

long-term IRU rights to underground pathways across dozens of additional routes, and deploying next-generation cable designed to fit two to four times more fiber into existing conduit.

That last point deserves emphasis, because it is the same insight from the opposite side of the transaction. A national carrier with existing conduit is spending capital specifically to raise the fiber density inside pathways it already owns — because the pathway is the expensive, slow, permission-bound asset, and the glass inside it is not. An operator building a new campus has one opportunity to make that decision cheaply, and it occurs before the trench is closed.

10.3 Even the compute vendors are buying the route

Analyst reporting from Needham and Wolfe Research in mid-2026 indicated that NVIDIA has been acquiring long-haul dark fiber across the United States, on routes with fiber counts reported at up to 100 pairs, in a program estimated at $5–10 billion over three years. NVIDIA has not publicly confirmed the scope, and these figures should be treated as analyst estimates rather than disclosed fact. The signal, however, is hard to misread: a company whose business is selling accelerators has concluded that owning physical transport is worth billions in capital, because control of the route determines how quickly compute can be delivered.

Taken together with the hyperscalers' privately owned subsea systems described in Section 3, a consistent pattern emerges across every tier of the AI infrastructure stack. The most sophisticated buyers are not purchasing bandwidth. They are purchasing routes, at high strand counts, ahead of demand, and holding the unlit remainder. That is the definition of a Neofiber strategy, and it is available to operators well below hyperscaler scale through a carrier-neutral partner willing to build to it.

CONCLUSION: TREAT FIBER LIKE POWER

AI and neocloud data center developers already treat power capacity, cooling, and land as first-order capital planning inputs, forecast years ahead of first operation. Dark fiber deserves the same treatment. The bandwidth curve for AI clusters is steep, non-linear, and directly tied to the same GPU density decisions that drive the power roadmap — which means the fiber plan can, and should, be derived from the same numbers that drive the power plan, rather than requested afterward as a service order.

BAAC Innovation designs and builds the outside-plant Neofiber infrastructure — strand count, duct and microduct structure, route diversity, and splice architecture — that AI and neocloud operators need to support GPU clusters from first power through their fourth or fifth expansion phase. Our standard route design of four primary conduits, each carrying two 7-way microduct bundles sized for high fiber-count cable, is built specifically to let operators light what they need today while banking capacity for the scale they are planning toward.

For operators evaluating a new AI campus, co-location footprint, or multi-site training cluster, we welcome the opportunity to translate a power and compute roadmap into a fiber design built to match it from day one.

About BAAC Innovation Inc.

BAAC Innovation Inc. is a Calgary-based utility and telecom infrastructure contractor specializing in long-haul fiber, horizontal directional drilling, trenchless construction, outside-plant (OSP) design, and civil works, with operations across Western Canada and expanding into the United States. BAAC designs and builds carrier-neutral dark fiber infrastructure — the Neofiber class described in this paper — via pipeline right-of-way repurposing and greenfield construction, supporting hyperscale, neocloud, and enterprise AI infrastructure clients.

APPENDIX A: GB300 NVL72 INTER-SITE BANDWIDTH REFERENCE

The table below is the full data set behind Figure 4, provided here in reusable table form. Figures assume AllReduce-style gradient synchronization for large-model training, using a conservative estimate of approximately 200 Gbps of inter-site bandwidth per scalable-unit (SU) pair, and 400GbE DWDM optics as the baseline strand-count reference. Coherent DWDM systems supporting 80 or more wavelengths per fiber pair can extend the usable life of a given strand count before new fiber construction becomes necessary.

SUs / Site	GPUs / Site	Est. Inter-Site Bandwidth	Min. Strands (400GbE DWDM)	Practical Fiber Bundle
1	72	~200 Gbps	2 strands	12F minimum
2–4	144–288	~400–800 Gbps	4–8 strands	12–24F
5–8	360–576	~1–1.6 Tbps	10–16 strands	24F (saturates)
9–16	648–1,152	~1.8–3.2 Tbps	18–32 strands	48F / DWDM upsell
32–64	2,304–4,608	~6–13 Tbps	64–128 strands	96F or 144F multi-tube
128	9,216	~26 Tbps	~200+ strands	288F multi-tube
500+	36,000+	~100+ Tbps	1,000+ strands	864F+ / multi-route

These figures are engineering planning estimates, not a substitute for a site-specific traffic study. Actual inter-site bandwidth depends on the training topology in use, the ratio of data-parallel to model-parallel communication, checkpointing frequency, and whether inference and training traffic share the same route.

Distance-Based Fiber Architecture Reference

How a dark fiber route is engineered depends heavily on the distance between sites, independent of the scalable-unit math above:

- **Under 120 km:** coherent optics can typically be integrated directly into routers or switches at each end, with no in-line amplification required — the simplest and lowest-cost architecture.
- **120 km to 500 km:** routes generally require an optical line system with in-line amplifiers spaced roughly every 100 km to maintain signal strength, while still avoiding full DWDM transponder chassis.
- **Beyond 500 km:** long-haul routes introduce additional signal regeneration and amplifier planning, and fiber availability and construction cost both become larger factors in the route's overall economics.

Matching the optical architecture to the distance between sites — rather than defaulting to the same design at every link — is part of what keeps dark fiber's cost advantage intact as campuses grow beyond a single metro.

APPENDIX B: GLOSSARY

Term	Definition
Dark Fiber	Unlit optical fiber strands leased or owned by an operator and lit with the operator's own equipment, rather than purchased as a managed, carrier-lit service.
Neofiber	BAAC's term for the class of dark fiber engineered specifically for AI infrastructure: high-strand-count, low-latency, route-diverse long-haul built into power-first sites, as distinct from the lower-strand, density-following long-haul of the cloud era.
DWDM	Dense Wavelength Division Multiplexing — a technology that transmits many independent wavelengths of light down a single fiber, multiplying the effective capacity of a strand pair without new construction.
Scalable Unit (SU)	A pod of tightly coupled GPUs, such as an NVIDIA GB300 NVL72 system, that must communicate with its peers at extremely high bandwidth to synchronize gradients during training or to serve large models during inference.
IRU	Indefeasible Right of Use — a typically prepaid lease of specific dark fiber strands, treated as a capital lease by the customer and as deferred revenue by the fiber provider. Terms historically ran 20 to 30 years; shorter terms, in some markets down to five years, are now available.
OSP	Outside Plant — the physical fiber infrastructure (conduit, duct, cable, splice enclosures) outside of a building, connecting sites across a campus, metro, or long-haul route.
Fiber Banking	Installing more strands, and more duct capacity for future strands, than current lit demand requires, and holding the unlit portion as a deliberate reserve against future growth.
Physical Route Diversity	Two or more paths between the same endpoints that share no common conduit, right-of-way, or structure — as distinguished from logical or contractual diversity, where separate circuits may still traverse the same physical plant.
Triversity / Quadversity	Requirements for three or four physically separate paths between endpoints. Under triversity, one production fiber pair is protected by two failover pairs, so every three pairs of raw supply carry one pair of production traffic.
ILA	In-Line Amplifier — a site along a long-haul route housing optical amplification equipment. ILA sites require land, structures, and electrical service, and are frequently the binding constraint on lighting additional fiber.
East-West / North-South Traffic	East-west traffic moves between data centers (machine to machine); north-south traffic moves between a data center and end users. East-west flows are typically one to two orders of magnitude larger and are kept on private backbone.
Hollow-Core Fiber	An emerging optical fiber design guiding light through an air core rather than solid glass, offering lower latency. Retaining spare microduct allows future cable types to be deployed without removing installed fiber.
DCI	Data Center Interconnect — the network connection between two or more data center sites, and the architecture and equipment used to establish it.
EVPN / VXLAN	Overlay networking technologies that encapsulate traffic for transport across a data center fabric. VXLAN encapsulation adds 50 bytes of overhead to each frame, which is why end-to-end jumbo frame support matters on interconnect links.

Term	Definition
MACsec	A layer-2 encryption standard. MACsec-encrypted traffic cannot cross a layer 2 boundary or be routed across an IP network, so it is available on dark fiber and CWDM/DWDM interconnects but not across routed third-party services. NVIDIA identifies it as its recommended DCI security implementation.
ELINE / ELAN	Third-party Ethernet virtual circuit services — point-to-point (ELINE) and multipoint-to-multipoint (ELAN) — delivered over a provider's network, for example as an MPLS pseudowire or a virtual private LAN service.
Colored / Gray Optics	Colored optics transmit on a specific DWDM wavelength and can feed a passive multiplexer directly; gray (colorless) optics transmit on a standard wavelength and require an active DWDM system to convert them. The choice depends on the xWDM equipment in use, cost, and available lead time.
Transponder	A device historically used in optical networks to convert and boost signals for long-distance transmission, increasingly being replaced by coherent pluggable optics installed directly in routers or switches.
Wavelength Services	A carrier-managed connectivity service providing a dedicated optical connection between two points at a fixed bandwidth increment, such as 100G or 400G.

REFERENCES

[1] Ciena Corporation, "Dark Fiber, AI, and the Future of Connectivity," 2025.

[2] Fetterolf, Peter, ACG Research, "Comparing Total Cost of Ownership for Dark Fiber, Carrier Ethernet, and Wavelength Services for Data Center Interconnect at 100G and Beyond," 2025.

[3] IEEE 802.3 / Ethernet Alliance bandwidth roadmap, referenced for per-link bandwidth targets.

[4] Corning and OFS optical cable specifications, referenced for strand-count and cable-structure comparisons.

[5] Grand View Research, "Dark Fiber Networks Market" industry analysis, as cited in Ciena Corporation (2025).

[6] Cisco, "Cisco AI Readiness Index," as cited in ACG Research (2025).

[7] McKinsey & Company, "Telco Operators' Growth Lies in AI Infrastructure," as cited in ACG Research (2025).

[8] Linkewire, "Data Center Interconnect Market," February 2025, as cited in ACG Research (2025).

[9] ACG Research, "The Economic Benefits of Routed Optical Networks," as cited in ACG Research (2025).

[10] Federal Aviation Administration statements on the Dallas TRACON telecommunications outage, September 20, 2025, as reported by CNN and Aviation Week Network.

[11] American Airlines internal communication on the September 19, 2025 fiber cuts, as reported by Aviation Week Network and NBC 5 Dallas–Fort Worth, September 2025.

[12] Associated Press, "AT&T attributes mass 911 outages in 3 states to fiber cuts made by third parties," September 26, 2025.

[13] Reuters, reporting on Baltic Sea submarine cable and Estlink 2 incidents, November–December 2024 and February 2025.

[14] Al Jazeera, "As undersea cables break off Europe and Taiwan, proving sabotage is tough," March 10, 2025.

[15] Taiwan Ministry of Digital Affairs, "2025 Report on the Analysis of Causes of Damage to Taiwan's Submarine Communication Cables and Improvement Measures," 2026.

[16] International Telecommunication Union and International Cable Protection Committee figures on annual cable faults and causes, as reported by Agence France-Presse, February 2025.

[17] Meta Platforms, Project Waterworth announcement, February 2025; and Lightreading, "Meta unveils Project Waterworth subsea cable initiative," February 17, 2025.

[18] Oxford Internet Institute, expert commentary on privately owned subsea cable systems, University of Oxford, March 10, 2025.

[19] TeleGeography subsea cable investment projections, as reported by CNBC, November 8, 2025.

[20] Zayo Group, "2026 Bandwidth Report," 2026.

[21] Lumen Technologies, 2026 Investor Day materials and Form 8-K, February 2026; and "Lumen Accelerates Multi-Billion-Dollar Network Expansion to Meet Soaring AI Demand," September 22, 2025.

[22] Needham & Company and Wolfe Research analyst commentary on NVIDIA long-haul dark fiber acquisition, June–July 2026, as reported in trade and financial press. Figures are analyst estimates; NVIDIA has not publicly confirmed program scope.

[23] Gemini Team, Google, "Gemini: A Family of Highly Capable Multimodal Models," arXiv, 2023; and "Gemini 1.5," arXiv, 2024.

[24] Google Cloud, "AI infrastructure at Next '26," April 2026.

[25] DatacenterDynamics, "Training Google's Gemini: TPUs, multiple data centers, and risks of cosmic rays," 2026.

[26] Epoch AI, "Could decentralized training solve AI's power problem?", October 28, 2025; and "Can AI scaling continue through 2030?"

[27] "Distributed and Decentralised Training: Technical Governance Challenges in a Shifting AI Landscape," arXiv:2507.07765, 2025.

[28] NVIDIA Corporation, "DCI Topologies," NVIDIA Networking Documentation, 2026.

[29] Dark fiber market advisory analysis on regional demand modeling, fiber pair forecasting, and right-of-way diversity, 2026.

[30] BAAC Innovation Inc., internal engineering analysis: NVIDIA GB300 NVL72 inter-site bandwidth and fiber strand modeling, 2026.

References 5 through 9 above were cited within the ACG Research and Ciena Corporation source materials referenced in this paper; they are listed here for transparency regarding the original data lineage and have not been independently re-verified by BAAC Innovation Inc.