on 16-Oct-2026

BERNSTEIN
SOCIETE GENERALE GROUP

15 October 2026

U.S. Multi Industry & Electrical Equipment

AI Data Center Interconnects: Copper Inside, Optical Outside, CPO in Between (webinar recap)



Varun Govindaraj
+1 917 344 8543
varun.govindaraj@bernsteinsg.com



Daniel Zhu
+1 917 344 8309
daniel.zhu@bernsteinsg.com

Specialist Sales



Steve Song
+1 917 344 8401
steve.song@bernsteinsg.com

Last month we initiated coverage of US interconnect manufacturers and networking equipment, with the evolving copper-versus-optics debate across AI data center interconnects a key theme. We also recently hosted an expert webinar with Brad Smith, a former NVIDIA leader with extensive experience across the copper and optical interconnect ecosystem. Summarizing the key takeaways from the discussion:

Rising AI workloads fundamentally alter the interconnect problem: Traditional data centers historically were ~90% copper, ~8% multi-mode optics and ~2% single-mode optics, but AI fundamentally changes the architecture: a single AI application can run across a million+ GPUs. This creates substantially greater bandwidth, reach and port-density requirements. *The debate is therefore about where the boundary between the 'copper and optical' ultimately sits.*

Copper inside, optics outside, CPO in between: Short-reach scale-up connectivity inside the rack remains predominantly copper *(including through Vera Rubin / Vera Rubin Ultra)*, while longer-reach scale-out / scale-across connections increasingly require single-mode optics. ***The framework is therefore relatively simple: use copper inside the rack for as long as possible, CPO increasingly around network switches, and optics for longer-reach connections outside the rack.***

Physics ultimately defines the copper-to-optics transition: Copper reach compresses rapidly as signaling speeds increase, from ~2.5m at ~100Gbps to ~18 inches at ~200Gbps and ~12 inches at ~400Gbps; multi-mode optics also becomes reach-constrained at higher speeds. DSPs can extend electrical reach but introduce additional power, cost and latency problems, while single-mode optics becomes increasingly necessary across the hundreds-of-meters distances found in large AI data centers. Within the rack, the gating factor for greater optical adoption is ultimately whether rising GPU/HBM density pushes package-level I/O beyond what short-reach copper cables can practically support.

CPO addresses an emerging switch-level constraint: As switch bandwidth and port density increase, the electrical connection between the switch ASIC and front-panel optical modules itself becomes challenging. CPO moves electrical-to-optical conversion adjacent to the switch silicon *(thereby shortening that electrical path from ~18 inches to ~0.5 inches).* The expert contrasted ~128 × 800Gbps ports on a current NVIDIA switch with ~512 × 800Gbps using CPO *(a 4x increase in capacity). The resulting increase in switch radix can reduce the physical space associated with millions of pluggable transceivers.*

CPO adoption will be gradual: Silicon photonics (SiPh) is the underlying chip technology enabling increasingly dense optical connectivity around the switch, while CPO is the system-level architecture integrating those connections with the switch silicon.

Overall, the transition toward CPO appears to be later than we had initially expected, with meaningful adoption only from 2030+; until then, scale-up is expected to remain predominantly copper.

on 16-Oct-2026

Varun Govindaraj +1 917 344 8543 varun.govindaraj@bernsteinsg.com 15 October 2026

BERNSTEIN TICKER TABLE

			14 Oct 2026		TTM	Reported EPS				Reported P/E (x)		
			Closing	Price	Rel.							
Ticker	Rating	Cur	Price	Target	Perf.	Cur	2025A	2026E	2027E	2025A	2026E	2027E
APH (Amphenol)	O	USD	82.69	99.00	12.6%	USD	1.67	2.65	3.31	49.5	31.2	25.0
TEL (TE Connectivity)	O	USD	210.25	243.00	(22.1)%	USD	9.26	11.46	12.98	22.7	18.4	16.2
LITE (Lumentum)	O	USD	896.33	1,220.00	427.1%	USD	2.05	8.37	22.40	438.2	107.0	40.0
COHR (Coherent)	O	USD	263.32	350.00	105.5%	USD	3.53	5.58	9.51	74.5	47.2	27.7
GLW (Corning)	M	USD	140.89	140.00	39.5%	USD	2.54	3.22	4.16	55.5	43.8	33.9
SPX			7,718.94									

O - Outperform, M - Market-Perform, U - Underperform, NR - Not Rated, CS - Coverage Suspended

APH, TEL estimate is Adjusted EPS; APH, TEL valuation is Adjusted P/E (x);

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

We rate APH (PT $99) and TEL (PT $243) Outperform.

We rate LITE (PT $1220), COHR (PT $350) Outperform and GLW (PT $140) Market-Perform.

RESEARCH INDEX

US Multi Industrials: Can we (inter)connect? Initiating coverage with Outperform on Amphenol and TE Connectivity

US Networking: Scaling the bandwidth wall. Initiating with a positive view

人在草木间

on 16-Oct-2026

DETAILS

We hosted an expert webinar with Brad Smith, a former NVIDIA leader with experience across both copper and optical interconnect technologies. Smith has spent roughly half of his career across the copper and optics ecosystem, including serving as CEO of an optical solutions company, working in business development at a silicon-photonics foundry funded by Intel's silicon-photonics business, running transceiver product lines, and subsequently joining Mellanox (later acquired by NVIDIA). Most recently, he worked across product marketing and management for the copper and optics product portfolio at NVIDIA.

The discussion focused on the ongoing copper-versus-optics debate across AI data center interconnects, including the underlying physics driving the transition, where copper versus optics fit within current and future AI architectures, the evolution of CPO and competing architectures, and what ultimately determines how far optics moves inside the rack.

What is the debate all about, and why does it exist in the first place?

The starting point for the copper-versus-optics debate is that ***AI fundamentally changes the way compute infrastructure needs to be interconnected***. Historically, the expert stated that ~90% of data center interconnects were copper, ~8% multi-mode optics and ~2% single-mode optics. This was broadly how the market operated prior to the acceleration of AI beginning August 2022. In traditional search / social media data centers, a single CPU could run ~10-12 applications simultaneously through virtualization. AI workloads turns that model to an extreme: one application can run across as many as a million GPUs.

That architectural change places very different requirements on the network. According to the expert, single-mode optics (relevant to AI workloads) support upto ~500m-2km reaches required across increasingly large AI data centers. In contrast, both copper and multi-mode optics become progressively less viable as the combination of bandwidth and distance increases. This is ultimately what sits behind the current debate: vastly larger numbers of GPUs / CPUs increasingly need to operate together as one system, creating significantly greater requirements for bandwidth, port density and hence interconnects - both in terms of numbers and reach. **The question is therefore where the boundary between copper and optics ultimately sits as AI architectures continue to scale**.

Framing the problem and key terminology

It is helpful first to understand the different layers of AI connectivity. The expert described "**scale-up"** as *connectivity within the rack*, where GPUs are connected together to behave as a larger virtual GPU; "**scale-out"** *connects multiple racks together*; and "**scale-across"** *connects entire data centers together*. The discussion also referenced NVIDIA Fusion / NVLink connectivity as a mechanism to connect an external AI-chip subsystem into the broader NVIDIA GPU, networking and software stack, referred to in the discussion as "**scale-in'**. We could eventually have **"scale-to-space"** as orbital data centers eventually emerge.

The distinction matters because *inside and outside the rack remain fundamentally different connectivity problems*. Inside the rack, distances are short and copper remains the predominant medium. Vera Rubin illustrates the current approach: the expert described ~18 inches as the maximum relevant copper wire length (for up to ~200Gbps interconnection speeds), with NVLink switching positioned centrally in the rack and compute trays above and below, connected through ~5,000 short copper cables. Even NVIDIA's Vera Rubin generation was described as retaining copper PCBs (with layers of copper traces carrying high-speed signals).

Future architectures can remove much of these copper cables without necessarily removing copper in its entirety. The expert described a Feynman-era architecture in which compute blades connect into a midplane PCB, replacing many of the individual cables with copper traces on the board itself, with NVLink switching positioned at the back. The architecture discussed included 4 such assemblies within a rack and was described as approaching ~1MW per rack. The expert also referenced commentary that implementing an equivalent architecture with optics would add *~60kW of power and $2M of cost*, reinforcing why manufacturers are unlikely to use optics for an ~18-inch connection while copper remains technically viable.

This leads to the expert's broader framework for how the interconnect architecture evolves: **copper inside the rack for as long as possible; CPO increasingly around network switches; and optics on the outside for the longer-reach connections between racks and across facilities** *(ie scale out / scale across)*. CPO is therefore central to the current discussion, but should not be interpreted as implying an imminent full transition from copper to optics throughout the AI data center, but a more gradual one.

What eventually pushes optics deeper into the scale-up domain is the evolution of the compute package itself. The expert used the progression from Hopper to Blackwell to illustrate this. Hopper was described as having one GPU die with HBM

on 16-Oct-2026

around it. TSMC's CoWoS, *('chip-on-wafer-on-substrate')* allows the memory to be positioned adjacent to the compute silicon. Blackwell moves to two dies with eight stacks of 12 HBM chips each, while future architectures are moving toward four dies and substantially more HBM on a package described as roughly the size of a mobile phone.

As more compute and memory are concentrated within a package, I/O requirements rise accordingly and networking hence becomes the next constraint. The expert therefore identified **the amount of silicon and memory that can be placed on a package as a key gating factor for when scale-up connectivity needs to become optical**. The broader message remains evolutionary: the industry will continue using PCBs (with copper traces) and high-speed copper connections until package-level I/O requirements exceed what short-reach electrical connections can support within practical limits.

Underlying physics: why the need for single-mode optics at all?

The primary limitation of copper is physical: **as signaling speed increases, achievable cable length declines**. The expert illustrated this through the historical progression of copper cable links. A conventional Cat 5e/6-style connection can support ~1Gbps over ~100m. At ~25Gbps, the reach falls to ~7m; at 50Gbps, ~3m; and at 100Gbps, ~2.5m. At the higher speeds relevant for AI workloads, passive copper reach compresses further, with the discussion citing ~18 inches for ~200Gbps and ~1 ft / 12 inches at ~400Gbps. In practical terms, the expert suggested that copper effectively starts to disappear at ~400Gbps as its usable reach becomes increasingly limited.

The expert explained the physics relatively simply: sending an electron through a wire creates a radio wave that radiates energy away from the conductor, leaving less energy available to maintain the signal farther along the wire. As signaling speeds rise, maintaining signal integrity therefore becomes progressively more difficult. Put simply, the faster the signal, the shorter the copper connection becomes.

One way to extend electrical reach is to add a digital signal processor (DSP) that compensates for signal degradation and noise, creating an active electrical cable (AEC). The expert cited Credo as an supplier and referenced ~$1,200 cost (per unit) and ~100-200 nanoseconds of latency for an AEC. With DSPs, ~200Gbps links can be pushed to ~2 meters, although the expert indicated that the approach becomes challenged at higher speeds and that ~400Gbps copper effectively runs out of reach.

Importantly, multi-mode optics faces its own version of the reach problem. Historically, multi-mode optics provided a relatively inexpensive bridge between copper and the more expensive single-mode optics. However, multi-mode optics gets constrained with ~100m of reach at 50Gbps and ~50m at 100Gbps, while ~200Gbps multi-mode was described as not yet having emerged as a viable technology and remaining in R&D stage. For large AI data centers where connections can extend hundreds of meters, this increasingly pushes the architecture toward single-mode optics.

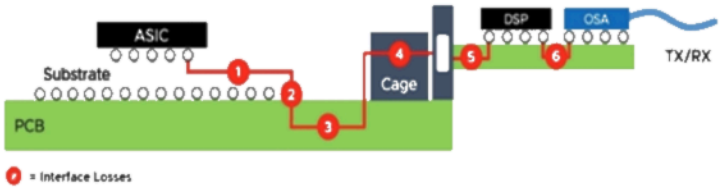
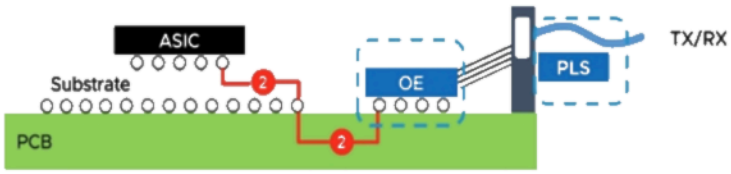
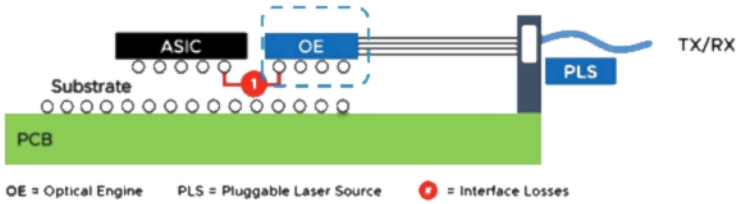
However, single-mode optics introduces its own cost and complexity. The expert highlighted the challenge of precisely aligning fiber with lasers and photodetectors, is ~30-35% of transceiver cost at an increasingly sensitive 9-micron core for single-mode fiber *(roughly the size of a red blood cell, think about aligning two red blood cells together)*. This reinforces the broader economic logic of the architecture: single-mode optics becomes increasingly important as reach increases, but there is little reason to absorb its additional cost and complexity for an 18-inch connection that copper can still serve.

The optical transceiver architecture itself is also evolving to reduce power and cost. A conventional transceiver combines the transmitter and receiver: on the transmit side a laser is modulated to encode the signal, while the receiver uses a photodetector and amplification to recover it. Optical links also experience signal-integrity issues, including reflections across interfaces, and conventional transceivers therefore contain DSPs to maintain signal integrity. The expert cited approximately $150 of cost (per unit) and ~15-20W of power associated with the DSP.

This has led to three broad approaches: Full-retimer optics (FRO) retains a DSP within the transceiver; low-power optics (LPO) eliminates the DSP entirely; and a hybrid approach retains signal processing on only one side of the link. The trade-off is straightforward: removing DSP functionality can reduce power and cost, but requires a more carefully engineered (and increasingly sensitive) link to maintain signal integrity.

on 16-Oct-2026

EXHIBIT 1: **How will the packaged optics evolution work?**

Method	Illustration	Description
Traditional switch with pluggable optical transceivers *(current technology)*		• **Conversion at the front panel**: Electrical signals (carrying data) travel from switch ASIC across PCB traces and interfaces to a pluggable transceiver, *where the signal is converted from copper to optical* • **Digital signal processor (DSP) compensates for signal loss**: DSP equalizes and retimes the degraded high-speed electrical signal before the optical subassembly transmits it over fiber • **Serviceable but power-intensive:** Transceivers are readily replaceable, but long electrical paths, multiple interfaces and DSP processing increase power requirements and constrain bandwidth
Near-packaged optics (NPO)		• **Optical conversion moves beside ASIC**: OE is removed from the front-panel transceiver and placed on the PCB; *close to, but on a separate package from, the switch ASIC* • **Shorter electrical path**: Signal traverses only a short board-level electrical path before being converted to optical; *fiber then carries the signal to the front panel* • **Integration–serviceability optimization**: NPO reduces interface losses and power versus traditional pluggable transceivers while potentially retaining a "pluggable" OE, improving serviceability
Co-packaged optics (CPO)		• **Optical conversion moves *onto* the ASIC package**: OE sits on the same package substrate as the switch ASIC, *bringing electrical-to-optical conversion within millimeters of silicon* • **Electrical path is minimized**: A very short package-level interface connects ASIC and OE; optical fiber carries the signal from the package to the front panel, *largely eliminating thick board-level high-speed copper cables* • **Maximum efficiency & tighter integration**: CPO offers the lowest path loss and highest bandwidth density, while laser remains externally pluggable for thermal management and replacement

Optical conversion progressively moves closer to the switch ASIC; shortening the electrical signal path and thereby ***reducing interface losses, power consumption and signal-conditioning requirements***

Source: Broadcom, Bernstein Asia Tech Hardware Team

Varun Govindaraj +1 917 344 8543 varun.govindaraj@bernsteinsg.com

on 16-Oct-2026

Deep dive on co-package optics (CPO)

Co-packaged optics (CPO) addresses a different problem: ***as switch bandwidth and port density increase, the electrical connection between the switch silicon and conventional front-panel optical modules itself becomes increasingly difficult***. Traditional switch architectures place pluggable transceivers at the front panel, physically separated from the switch ASIC. CPO brings this electrical-to-optical conversion (through optical engines) adjacent to the switch chip, shrinking the distance between the electrical switch silicon and the optical interface from ~18 inches to ~0.5 inches (according to the expert). Hence, CPO is not synonymous with replacing every copper connection inside the rack, nor does it imply that every conventional optical transceiver disappears. Rather, CPO addresses the challenge of fitting increasing amounts of optical bandwidth around increasingly large and powerful switch silicon chips.

CPO itself is not a new technology. Earlier efforts to manufacture CPO at scale struggled because there was insufficient market demand. The current cycle is different because AI workloads require substantially more bandwidth and connectivity, while investment in the underlying manufacturing ecosystem has accelerated. NVIDIA was characterized as an important driver of the current CPO push, alongside increasing activity from Broadcom, Cisco and Marvell. The expert noted that NVIDIA CPO switches began shipping in summer 2026, but characterized this as the beginning of the technology cycle, with more meaningful practical adoption expected from 2030 onwards and CPO initially deployed at the highest levels of the network before moving progressively downward, with technological maturity probably stretching into 2035 and beyond.

The maturity of silicon photonics (SiPh) and the associated manufacturing ecosystem are central to enabling this transition. *Please note that SiPh is the underlying chip technology, while CPO is the system-level packaging / architectural framework that integrates optical connectivity around the switch silicon.* SiPh experienced an earlier investment cycle from ~2015-18 but did not develop as anticipated and remained more expensive than competing technologies, leading many fabs to step back production. The expert argued that the technology has since caught up materially with market requirements, citing SiPh solutions enabling upto ~400Gbps links.

One technology highlighted was the ring-resonator modulator, which can materially reduce the size of the optical modulation element. The expert contrasted a traditional SiPh modulator, described as closer to postage-stamp size (in terms of scale), with a ring resonator cited at ~20 microns in diameter. This becomes increasingly important as the number of connections around the switch grows and individual optical elements need to become smaller. The expert discussed NVIDIA's investment in R&D and TSMC SiPh in this context, while GlobalFoundries was also cited among companies investing into the broader manufacturing ecosystem.

Ultimately, the attraction of CPO is best understood through switch radix *(or the total number of external ports / connections that a switch can support)*. The expert described large AI data centers as requiring ~15,000-25,000 network switches stacked across 4-5 layers because each individual switch has a limited number of ports. He contrasted a current NVIDIA switch configuration with ~128 800Gbps ports against a CPO configuration supporting ~512 800Gbps ports *(a 4x increase in port capacity)*. At that level of density, conventional pluggable transceivers become increasingly difficult simply because of the physical space (and amount of power) they consume.

Higher radix potentially allows the network to become flatter, reducing the number of switching layers and intermediate connections. Reliability also becomes increasingly relevant at these scales. The expert cited ~6-7 million transceivers in a single data center (today); with that many modules deployed, individual failures become frequent enough to matter operationally, a transceiver failure every few hours and the system potentially entering checkpoint mode once sufficient failures accumulate. *CPO therefore addresses more than headline bandwidth: port density, physical space, network topology and system-level reliability all become part of the rationale*

Alternatives to CPO: co-packaged copper (CPC) and near-packaged optics (NPO)

CPO is not the only technology in response to the electrical-reach and front-panel-density problem. One alternative is co-packaged copper (CPC), where high-speed copper connectors are placed directly around switch silicon (substrate). In CPC, the connector is brought directly onto or very close to the ASIC substrate. Instead of the high-speed signal traveling a relatively long distance through copper traces on the PCB (as in the current technology today), the signal enters a purpose-built copper cable almost immediately outside the silicon. TE Connectivity, Molex and Amphenol were cited as companies working on CPC solutions across switch chips and other semiconductor packages.

on 16-Oct-2026

A second alternative is near-packaged optics (NPO). Instead of fully co-packaging the optics with the switch silicon (as is the case with CPC), NPO moves the transceiver inward from the front panel and positions it closer to the switch, including architectures where the optical modules point outward at a 90-degree orientation and provide additional channels. DSP functionality can also be removed depending on the implementation. The trade-off is physical size. As switch chips themselves become larger, surrounding them with sizable optical modules results in an increasingly larger package. The expert illustrated this by describing a large switch chip as roughly the size of a cell phone and the adjacent optical modules as the size of "small hot dogs"; once assembled, the overall footprint can approach the size of a laptop. NPO therefore solves part of the reach problem without taking the optics all the way into the package, but physical density eventually becomes a constraint.

Multiple architectures are thus developing in parallel, while CPO itself is expected to begin at the upper layers of the network and broaden over time as manufacturing ecosystem matures.

Other points of discussion

- **Inference / custom ASICs:** Given that inference is application-specific with custom ASICs, it is difficult to prescribe one universal networking architecture for inference workloads. Custom ASICs can be optimized for specific workloads rather than providing the generalized GPU architecture offered by NVIDIA / AMD today which is more relevant for training workloads. Google was cited as an example for custom-made ASIC associated with YouTube workloads. These custom ASICs were described as being 'first-generation' ~500W TDP and air cooled. Distributed inference introduces another model, with clusters of different computers contributing to workloads and software bringing the outputs together to create tokens. Despite these differences, the expert's near-term conclusion remained consistent: inference does not inherently require optics inside the rack today, and copper should continue to be used where bandwidth and reach allow it.

- **Optical circuit switching (OCS):** The expert was also asked whether Google has an advantage through optical circuit switching (OCS). He emphasized that OCS itself is not a new technology, having been used in telecom networks for decades. Unlike an electronic switch, which the expert described as providing ~300 nanoseconds port-to-port switching, OCS acts more like a railroad junction node by establishing a direct optical path between specified ports, for example from port 1 to port 6. This can be useful in reconfiguring resources, such as linking storage capacity to different GPU clusters, but the expert did not view OCS as replacing conventional electronic switching. Coherent was cited as a participant using liquid-crystal technology, while a MEMS-based approach by Lumentum was also discussed.

- **Laser architecture evolution:** The discussion also covered the evolution of laser technology. The expert highlighted indium-phosphide lasers, noting that indium is a byproduct of zinc production and that ~70% of supply comes from China. Indium phosphide lasers historically served primarily telecom applications, and the AI optical build-out is placing substantially greater demands on the supply chain; high-power lasers were also described as more challenging from a manufacturing-yield perspective.

- **Fabrinet**: Finally, the expert discussed Fabrinet's positioning as the photonics architecture evolves. NVIDIA has used Fabrinet historically for transceiver manufacturing, while other suppliers also use outsourced manufacturing models. The expert characterized Fabrinet as a mainstream, high-quality manufacturing participant rather than a transceiver technology supplier itself, and noted that the company is scaling its optical capabilities going forward.

The authors would like to thank Sakansh Mittal for his help in preparing this report.