TECHNOLOGY · AI INFRASTRUCTURE

After the Compute Glut: The Trillion-Dollar Restructuring of AI Interconnect Systems

Compute is no longer the constraint. The structural restructuring of AI interconnect has begun.

Lead Analyst: Chuck Guo (■■) · Contributing: CJ Muse (Cantor Fitzgerald), Tim Savageaux (Northland Capital), Karl Ackerman (BNP Paribas) · May 2026 · CK-2026-OPT-001

INVESTMENT HIGHLIGHTS

Six key investment findings

• 01 Valuation dislocation: the combined market cap of all listed optical interconnect companies is below that of Micron alone — yet the sector stands at the starting point of a historic restructuring.

• 02 Bottleneck migration: GPU utilisation data shows that once AI clusters scale beyond thousands of GPUs, inter-GPU communication overhead matches compute itself. The marginal constraint on AI has shifted from GPUs to interconnect.

• 03 Copper-optical coexistence: not a simple replacement, but scene differentiation. Copper retains cost and reliability advantages in Scale-up (short-distance GPU-to-GPU) scenarios; optical is irreplaceable in Scale-out (inter-rack, longer-distance) scenarios.

• 04 CPO timeline: the endgame architecture, but 2026 is primarily a validation year. Large-scale commercial deployment is constrained by manufacturing yield and operational complexity; scale-up expected post-2028.

• 05 2026 certainty: the most certain revenue increment lies in 1.6T pluggable module ramp, LPO/NPO transition deployment, and PCB/ABF substrate + silicon photonics foundry capacity expansion.

• 06 Value chain shift: profit pools are migrating from module assembly toward chips and advanced packaging. Cross-stage supply capability is the core driver of long-term alpha.

SECTION I

Compute Is Not the Bottleneck — Interconnect Is

1.1 The Key Data Point: 8,000 GPUs, Actual Utilisation Only 38%

In 2024, Meta disclosed in an internal AI infrastructure study that its large-scale training clusters, once scaled to thousands of GPUs, sustained actual GPU utilisation below 40%. The limiting factor was not compute capability itself, but inter-GPU communication overhead — which had come to match the compute workload.

This is not an isolated observation. Engineering teams at Anthropic, Google DeepMind, and others have described the same pattern at multiple technical forums: the larger the cluster, the higher the GPU idle-time fraction and the lower the fraction of effective compute time.

1.2 How the Communication Bottleneck Forms

Large-model training relies on two core parallelism strategies: Tensor Parallelism and Expert Parallelism (MoE architecture). Both require frequent, large-scale synchronisation of intermediate compute results across GPUs — the AllReduce operation.

The root cause: AllReduce communication volume grows super-linearly with GPU count. As clusters scale from hundreds to tens of thousands of GPUs, the data volume to be synchronised grows far faster than the compute added. Beyond a critical scale, adding GPUs no longer improves training efficiency — it worsens communication congestion.

Three key data points on the interconnect bottleneck

• In AI training clusters, inter-GPU communication overhead accounts for ~31% of total energy — comparable to compute itself (38%).

• A single 1.6T optical transceiver consumes ~30W, of which the DSP chip accounts for ~18W.

• NVIDIA test data: CPO (Co-Packaged Optics) switches save ~70% of communication power versus traditional optical modules.

1.3 Three Physical Constraints of Copper Cable

- Bandwidth ceiling: copper's maximum data rate is constrained by material physics and cannot keep pace with AI's per-generation bandwidth doubling requirement (400G → 800G → 1.6T → 3.2T).
- Distance attenuation: copper's effective transmission distance does not exceed ~2 metres; beyond this, signal degradation is severe. AI datacentre inter-rack distances typically run to tens or hundreds of metres.
- Power explosion: GPU power per card has risen from H100's 700W to GB300's 1,000W+. At this power density, tens of thousands of copper cables in a cluster already account for over 30% of total datacentre energy consumption.

Optical fibre's physical properties precisely address all three: single-fibre bandwidth is orders of magnitude higher than copper (WDM further multiplies it), effective transmission distance reaches kilometres, and transmission power consumption is negligible.

> ***"Within five years, all AI datacentre interconnect will move fully to optical. This is not an industry prediction — it is the inevitable conclusion of the laws of physics."***

SECTION II

A Systematically Undervalued Value Chain

2.1 From Pluggable to Co-Packaged: Technology Evolution and Value Migration

Optical interconnect is not a single technology substitution event, but a structural migration in which optical components continuously move closer to the GPU. Each evolutionary stage involves a clear redistribution of value.

Stage 1 — Pluggable Optical Transceivers (current mainstream; accelerating in 2025–2026)

The pluggable optical transceiver is the most prevalent form of optical interconnect in today's datacentres. GPU output electrical signals travel 30–50cm along PCB copper traces to the backplane transceiver, where E/O conversion occurs before high-speed optical transmission.

Five core components: laser (InP material; silicon cannot lase due to indirect bandgap), modulator (EML integrates laser and modulator on a single chip), photodetector (PD), DSP (~35% of module BOM cost, fabricated by TSMC on silicon process), and micron-precision fibre coupling assemblies (FAU).

Current mainstream: 800G, rapidly transitioning to 1.6T in 2026. 1.6T target shipments ~2 million units in 2026, ~150% YoY increase.

Stage 2 — LPO / NPO Transition Solutions (high 2026 near-term visibility)

LPO (Linear Pluggable Optics) removes the highest-power DSP chip, replacing it with linear drive — cutting power to roughly one-third of traditional modules while preserving hot-swappability. NPO (Near-Packaged Optics) moves the optical engine adjacent to the switch chip on the PCB, maintaining removability but reducing distance.

Bernstein research projects LPO shipments could exceed CPO by 2030. LPO/NPO are not simply a transition to be superseded by CPO — they are an independently large market.

Stage 3 — CPO (Co-Packaged Optics) (endgame; gradual volume from 2027)

CPO integrates optical components directly inside the chip package, shrinking E/O conversion distance from 30–50cm to millimetres, reducing communication power ~70% versus traditional pluggable solutions.

Two core CPO components: Silicon Photonics PIC (integrates modulator, waveguides, detectors on SOI wafer; silicon cannot lase — light source must be supplied externally) and ELS (External Laser Source; independent InP laser module; externally replaceable because InP lasers generate heat and have limited lifespan).

CPO creates ELS as an entirely new component category. Existing laser vendor production lines are locked to EML supply via long-term agreements (LTAs); CPO requires CW (Continuous Wave) lasers using a different process — lines cannot be directly converted. This 'process mismatch' is the root of ELS supply scarcity.

2.2 Six-Layer Value Chain Structure

Tier	Function	Key Companies	Critical Constraint
① Platform / Architecture	Defines AI datacentre architecture standards; controls technology roadmap selection	NVIDIA, Broadcom, Cisco	Architecture definition power
② Optical components & modules	Lasers, modulators, detectors, module assembly	Coherent (COHR), Lumentum (LITE), AAOI	InP capacity extremely scarce

Tier	Function	Key Companies	Critical Constraint
③ Interconnect chips	DSP, CPO controllers, retimers	Broadcom (AVGO), Marvell (MRVL), Credo	Duopoly structure entrenched
④ Foundry / Manufacturing	Silicon photonics PIC foundry (SOI), compound semiconductor foundry	Tower Semi (TSEM), TSMC, Fabrinet	TSEM 70% capacity locked to 2028
⑤ Advanced packaging & test	CPO hybrid integration, reliability verification	TSMC COUPE, POET, AEHR Test	CPO mass production yield — key gate
⑥ Upstream materials	SOI substrates, InP substrates, optical fibre, epitaxial layers	Soitec, AXT, IQE, Corning (GLW)	Soitec monopoly on SOI (~95% share)

SECTION III

Capital Has Already Voted

3.1 NVIDIA's $7.2bn: Locking Down the Next Constraint

NVIDIA rarely makes large-scale strategic investments in upstream suppliers. Yet in the first half of 2026, NVIDIA completed three targeted investments in optical interconnect upstream:

Date	Target	Amount	Strategic Intent
2 Mar 2026	Lumentum (LITE)	$2.0bn	Lock in preferential supply of InP EML/CW lasers through 2027–2028
2 Mar 2026	Coherent (COHR)	$2.0bn	Lock in full-stack optical component and module capacity, plus multi-billion-dollar purchase commitments
6 May 2026	Corning (GLW)	$3.2bn	Build three dedicated optical factories in NC and TX; scale interconnect capacity 10x

Historical parallel: NVIDIA's last comparable action was its strategic positioning in memory upstream during the HBM capacity tightening cycle. This round of $7.2bn in targeted optical interconnect investment signals that NVIDIA has designated optical interconnect as the next critical constraint that could limit AI cluster expansion.

> ***"NVIDIA's $7.2bn targeted investment in optical interconnect upstream is the equivalent of declaring: optical interconnect is the next structural constraint of AI infrastructure."***

3.2 Five Hyperscalers: >$300bn Capex Pointing in the Same Direction

The five major hyperscalers (Google, Microsoft, AWS, Meta, Oracle) combined capex in the first nine months of 2025 exceeded $300bn — surpassing any prior full-year peak on record. The vast majority is flowing to AI infrastructure, and interconnect's share of that total is rising.

3.3 Key Stock Performance (2024 Jan – Apr 2026, base = 100)

Ticker	Company	Performance	Key Driver
AAOI	Applied Optoelectronics	+3,360% ($5 → $173)	1.6T order ramp + US-made narrative dual catalyst
LITE	Lumentum	+1,518% ($55 → $890)	NVIDIA $2bn strategic investment; +20% on announcement day
COHR	Coherent	+1,204% ($8 → $365)	Full-stack optical integration; backlog >$3.9bn
TSEM	Tower Semiconductor	+609% ($35 → $248)	Silicon photonics PIC foundry; 70% capacity locked to 2028
GLW	Corning	+469% ($29 → $165)	Global optical fibre leader; NVIDIA $3.2bn dedicated factory investment
NVDA	NVIDIA (reference)	+280% ($50 → $242)	AI core compute benchmark

SECTION IV

CPO: Structural Reorganisation of the Profit Pool

4.1 Value Migration: Who Benefits, Who is Pressured

• DSP chips (pressure): CPO architecture shrinks E/O conversion distance to millimetres, greatly reducing the need for complex error-correction coding. DSP will be simplified or removed. Structural pressure on DSP design companies over the medium-to-long term.

• Module assembly (pressure): pluggable optical modules as an independent product form factor will be displaced by CPO. Core value of traditional module assembly houses dilutes.

• Silicon photonics PIC foundry (benefit): CPO forces adoption of SOI-process silicon photonics PIC. TSEM's PH18 platform, with process leadership and >50 deeply bound design clients, has 70% of capacity locked to 2028.

• ELS External Laser Source (benefit; new category): entirely new component class created by CPO. Incumbent production lines locked to EML LTAs; overflow demand flows to independent suppliers such as Sivers (SIVE).

• Switch ASIC and platform layer (ultimate beneficiary): NVIDIA and Broadcom hold architecture definition rights. CPO proliferation only deepens moats.

4.2 The Realistic CPO Commercialisation Timeline

The 'CPO mass market in 2026' narrative circulating in the market requires careful correction. Three barriers:

Three barriers to CPO mass commercialisation

• Reliability threshold: a failed pluggable module is hot-swapped in minutes on-site. A failed CPO optical engine — integrated inside the switch — requires full chassis return. Downtime costs are unacceptable for cloud providers. Zhongji Innolight has explicitly stated: 'no hyperscaler customer is planning large-scale CPO deployment in 2026–2027.'

• Transition solution effectiveness: LPO reduces power to one-third of traditional modules; NPO further shortens the optical engine-to-chip distance. Before CPO reliability is fully validated, hyperscalers have ample reason to prefer transition solutions.

• Scale-up scenario copper advantage: in ultra-short-distance GPU-to-GPU interconnect scenarios, copper's cost and reliability advantage remains unmatched. LightCounting projects copper will still hold ~50% of the 1.6T interconnect market in 2029.

Our timeline: 2026 = small-volume validation in switch-side Scale-out scenarios; 2027 = CPO moves from samples to initial customer deployment; post-2028 = scale ramp. Goldman Sachs' $910bn CPO forecast for 2028 represents an optimistic ceiling scenario, not a 2026 baseline.

SECTION V

Timing Matters as Much as Direction

Optical interconnect is the next major battleground of AI infrastructure; CPO is the endgame form — direction certainty is beyond question. But on this track, timing precision matters as much as directional conviction: entering too early risks years of waiting; too late risks missing the highest-elasticity window.

5.1 2026: Position in Certainty, Not in Long-Dated Options

The most certain revenue increment in 2026 comes not from the purest CPO concept names, but from the infrastructure segments that must be upgraded before CPO can scale:

- 1.6T pluggable module ramp: ~2 million units targeted in 2026 (+~150% YoY); ASP markedly higher. Direct beneficiaries: LITE, COHR, AAOI.
- LPO/NPO transition deployment: 800G LPO in mass production; 1.6T LPO in development. One-third power reduction makes these the preferred power-efficiency tool for hyperscalers in 2026–2027.
- PCB/ABF substrate upgrade: Rubin platform uses 44-layer midplane and M8-grade low-loss copper-clad. AI server PCB value per unit rises from ~$100–150 to ~$300. ABF substrate ASP rising ~5–7% quarterly due to T-glass shortage.
- Silicon photonics foundry reservation: TSEM's $1.3bn silicon photonics contract (est. 2027 delivery) is itself a direct signal of capacity scarcity.

5.2 2027: CPO Enters Initial Customer Deployment

Key metrics to track: NVIDIA Quantum-X / Spectrum-X Photonics real deployment scale and customer feedback; Broadcom Tomahawk 5-Bailly CPO switch customer roster expansion; CPO procurement decisions at CoreWeave, Lambda Labs, Meta, and Google; whether ELS supply chain has entered stable volume production.

5.3 The Three Structural Bottlenecks That Cannot Be Bypassed

Segment	Source of Scarcity	Key Company	Current Status
SOI substrates	Soitec holds ~95% global share; sole certified supplier for three major silicon photonics foundries	Soitec (SOI.PA)	Every silicon photonics PIC requires Soitec SOI wafers; no substitute
Silicon photonics PIC foundry	PH18 SOI platform has bound >50 design clients; process migration cost prohibitive	Tower Semi (TSEM)	70% capacity contracts locked to 2028; $1.3bn new contract announced
InP laser capacity	EML lines already locked to LTAs; CPO requires CW lasers, different process — lines not directly convertible	LITE, COHR, SIVE	LITE CEO publicly stated company itself must purchase CPO lasers from open market

Lead Analyst: Chuck Guo (■■) | Charles & Kwok Multi-Asset Research | Report No.: CK-2026-OPT-001 | May 2026