# Optical Blast Radius: When a $10 Laser Bricks a $40,000 GPU

Second Order Labs™ · June 2026

ABSTRACT

The industry crowned silicon photonics as the answer to AI's copper wall. But optics doesn't remove the scaling constraint, it relocates it onto TSMC's packaging lines and a handful of Indium Phosphide fabs.

*Keywords:* Silicon Photonics, Co-Packaged Optics, Indium Phosphide, TSMC COUPE, AI Interconnect, Advanced Packaging

> *"The key challenges for CPO are optical fiber alignment, packaging yield and whether systems can continue operating without interruption when components need replacement."*
>
> — Kim Jung-ho, The Elec, AI Data Center Optical Communications & Interconnect Tech Conference

A ten-dollar laser can now destroy a forty-thousand-dollar GPU. That single fact exposes the flaw in the industry's favorite story about photonics. Copper interconnects have hit a physical breaking point that trillion-parameter models cannot engineer around [1]. Silicon photonics is the crowned solution. But the coronation is premature.

The dominant narrative treats optical I/O as a frictionless silver bullet: swap copper for light to escape the power crisis and scale GPUs. That framing is wrong in a specific and expensive way. Photonics does not eliminate the AI scaling constraint. It relocates it.

> *Optics doesn't break the ceiling on AI scaling. It picks the ceiling up and moves it from the data-center power grid onto TSMC's packaging lines and a handful of Indium Phosphide fabs.*

Solving a visible constraint like interconnect power silently manufactures three new ones: packaging capacity, laser supply, and asymmetric failure risk. I call

this the second-order bottleneck. Every player racing toward co-packaged optics is walking into a set of chokepoints that mainstream coverage ignores. I map all three below, along with the strategic inversion they set up.

## I. Why Copper Hit a Wall and Optics Only Moved It

The Copper Wall is the point where signal loss and power draw of electrical wiring become the dominant constraint on the entire system[2]. Next-generation XPUs demand bandwidths climbing toward 28.8 Tbps. At that figure, traditional copper interconnects and electrical SerDes architectures become mathematically unviable on power alone[3].

Optics answers this cleanly on the physics. Ayar Labs' TeraPHY chiplet delivers up to 5x to 10x higher bandwidth. It also provides 10x lower latency and 4x to 8x better power efficiency than pluggable optics plus electrical SerDes[4]. Lightmatter's Nick Harris frames the shift bluntly. Photonics' moment arrived not because it suddenly got better, but because copper reached its limit[5]. On that narrow question, the industry is right.

The error is treating a bandwidth solution as a scaling solution. Compute stopped being the primary bottleneck a while ago. The constraint moved to I/O. Optics addresses I/O. What it ignores is that building optical interconnect at scale demands the exact same scarce industrial inputs as building the GPU it connects. You do not escape a bottleneck by shifting load onto an input that is already constrained. You just make the constraint harder to see.

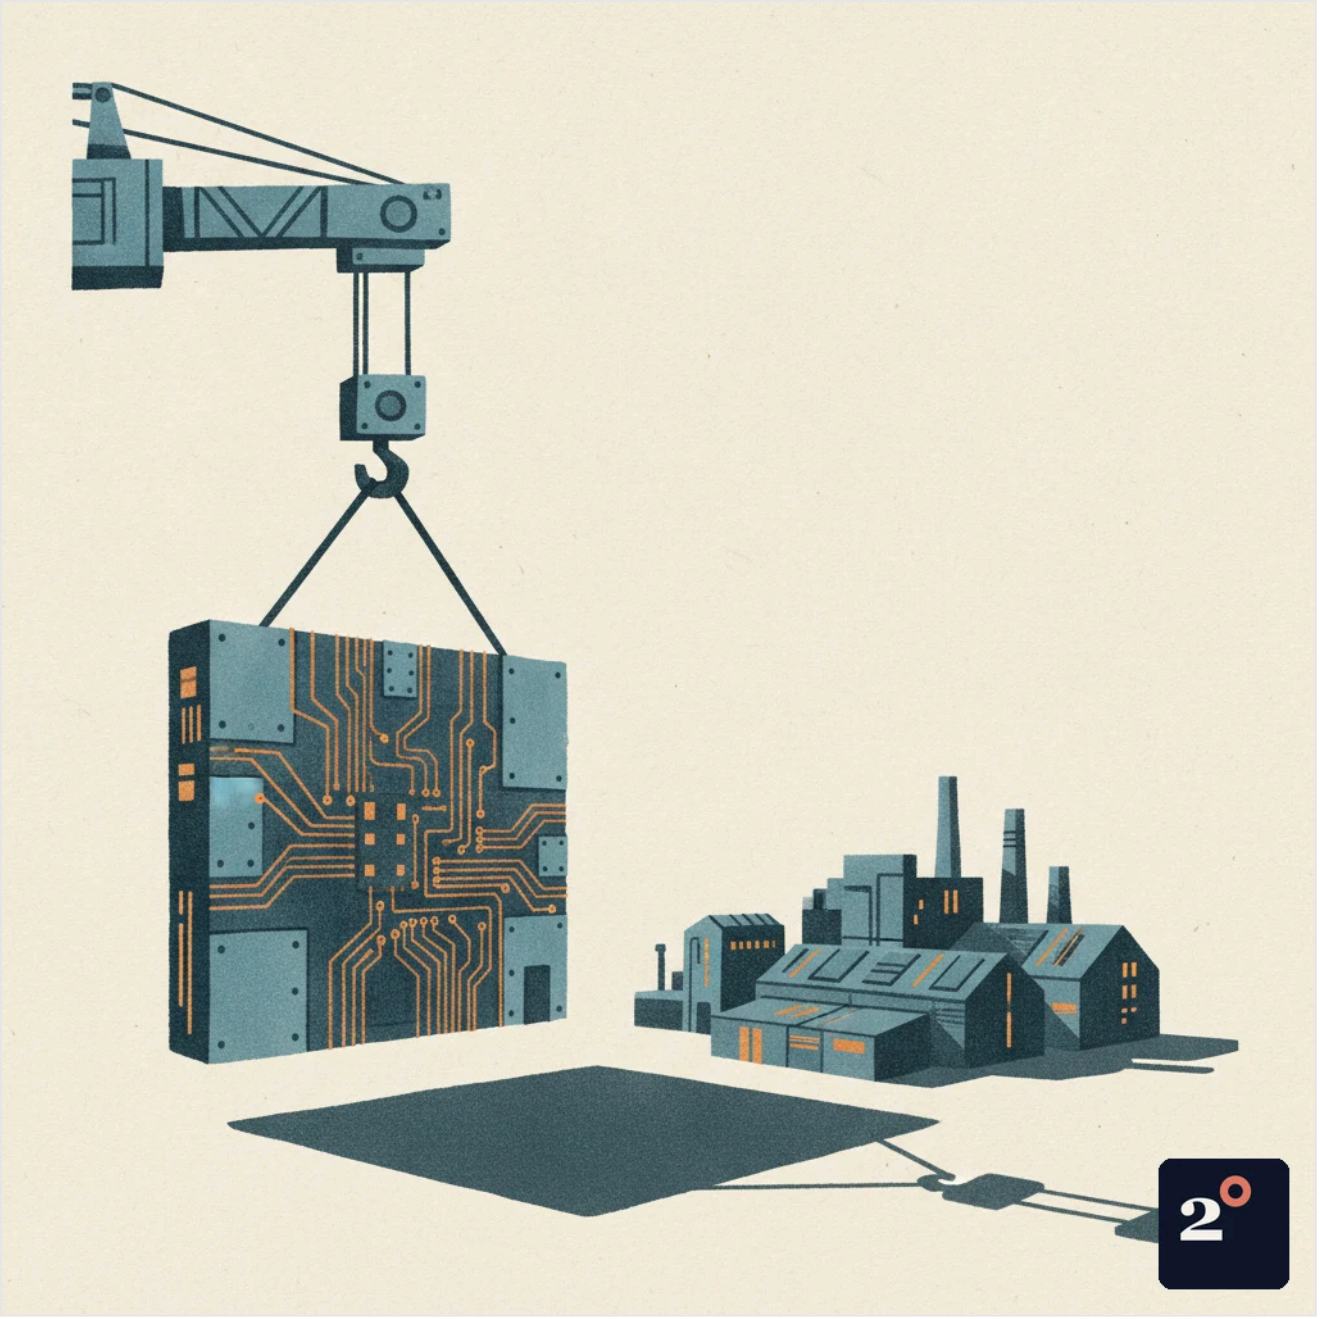

**Figure 1:** Optics relocates the AI scaling wall rather than removing it.

## II. Packaging Cannibalization: Optical Engines Fight GPUs for the Same Foundry Lines

Co-packaged optics requires advanced packaging, specifically the CoWoS and SoIC processes that already sit at the center of AI compute. TSMC's COUPE platform uses SoIC-X chip stacking to place an electrical die directly on top of a photonic die, achieving ultra-low impedance and superior energy efficiency versus conventional stacking[6]. That is genuinely a better mousetrap. It also consumes the exact packaging capacity that HBM stacks and AI accelerators are fighting over right now.

I call this Packaging Cannibalization. Every wafer of advanced-packaging capacity spent building the optical interconnect is a wafer not spent building the accelerator it serves. CoWoS-S already integrates high-bandwidth memory to logic in a 2.5D silicon interposer[7]. When optical engines join that queue,

they do not get a separate factory. They enter the existing queue, fighting AI accelerators for TSMC's limited tool time. The second-order effect is a supply squeeze that compounds. More AI demand drives more optical interconnect, which consumes more of the capacity that AI accelerators need, tightening the very bottleneck optics was supposed to relieve.

TSMC understands its pricing power here perfectly. COUPE is positioned to occupy the same role in silicon photonics that CoWoS occupies in AI packaging: a foundry-and-packaging platform, not a merchant switch product[8]. The roadmap extends to panel-level CoPoS by 2029, framed explicitly as the successor to CoWoS[9]. Samsung, sensing a second monopoly forming, released its own silicon photonics foundry roadmap at OFC 2026 to apply competitive pressure[10].

| Chokepoint | What it constrains | Who controls it |
|---|---|---|
| Advanced packaging (CoWoS/SoIC) | Both GPUs and optical engines | TSMC (COUPE), Samsung challenging |
| Indium Phosphide lasers | Light generation for all optics | InP fabs, Veeco tooling |
| ELSFP / modular design | Failure isolation and uptime | System architects (NVIDIA et al.) |

**Figure 2:** Optical engines and AI accelerators share the same advanced-packaging line, competing for the same tools.

## III. Silicon Cannot Lase, So Indium Phosphide Becomes the Hidden Chokepoint

Silicon is superb at computation and terrible at generating light. Every high-performance optical transceiver, no matter how advanced its packaging, still depends on Indium Phosphide lasers and photodetectors[11]. That material dependency sits entirely outside the traditional silicon supply chain. It is nowhere near ready for the volume AI is about to demand.

The evidence is in the capital. Veeco announced more than $250 million in equipment orders specifically to manufacture InP lasers, with deliveries starting in 2026 and accelerating sharply in 2027[12]. Orders of that size signal that current InP capacity is inadequate for what is coming. Aeluma put it plainly. The AI data-center buildout is outpacing the photonics supply chain, and customers are evaluating quantum-dot laser technology purely to bridge near-term supply gaps[13].

**Figure 3:** More than $250M in tooling orders, with deliveries ramping through 2026-2027, marks the scramble to build InP capacity that does not yet exist.

Read the timeline against the roadmaps and the mismatch jumps out. COUPE targets volume production in 2026. The InP laser capacity to feed it does not significantly accelerate until 2027. The packaging platform arrives before its critical input does.

## IV. Optical Blast Radius: When a Cheap Laser Bricks an Expensive GPU

Integration has a dark side the efficiency numbers hide. When you weld the light source onto the compute package, you fuse their fates. A single laser failure in a poorly designed CPO architecture can render a $40,000 GPU or a $200,000 switch unusable, because you cannot replace the failed component without discarding the whole assembly[14]. I call this the Optical Blast Radius. It measures the total dollar value destroyed when the cheapest component in an integrated optical package fails.

The physics makes this worse. Repeated thermal cycling in densely integrated data-center optics degrades laser lifetime rapidly, precisely the failure mode most likely to trigger in a hot, always-on AI cluster[15]. You are attaching the least durable element to the most expensive one, inside the harshest thermal environment in the building.

> *"The key challenges for CPO are optical fiber alignment, packaging yield and whether systems can continue operating without interruption when components need replacement." Kim Jung-ho, The Elec*

The industry's response reveals the real design constraint. The architecture gaining mainstream traction is ELSFP (External Laser Source Form factor Pluggable), which physically separates the laser and mounts it as a replaceable external module on the front of the switch[14]. NVIDIA's CPO approach does the same, using external light sources specifically to avoid the reliability trap where thermal cycling forces costly network interruptions[16]. The lesson is counterintuitive. The winning optical design is the one that keeps the light source out of the package. Maximum integration is not the goal. Minimizing the blast radius is.

## V. Fabric Value Inversion: How Optics Threatens to Commoditize the GPU

The deepest consequence is strategic, not physical. Perfect optical interconnect makes memory pooling and low-latency fabric across a cluster genuinely viable. That quietly undermines the case for the monolithic, do-everything accelerator. Celestial AI's Photonic Fabric is designed exactly for high-bandwidth, low-latency connectivity across large-scale deployments[17]. It is the kind of fabric that lets many cheaper components behave like one giant machine.

If the fabric becomes good enough, value migrates from the compute node to the optical switching layer that binds nodes together. The GPU risks becoming the commodity. The fabric becomes the moat. The money is already moving to front-run it. Marvell's acquisition of Celestial AI folds Photonic Fabric into a networking incumbent's portfolio[18], a bet that whoever owns the fabric owns the value. The broader consolidation, alongside the industry pivot from component-level pluggable transceivers to platform-level foundry-integrated optical engines[19], tells you where incumbents think the profit pool is heading.

NVIDIA's and TSMC's maneuvers look different in this light. Locking up laser capacity and packaging lines is not only about scaling faster. It is a defensive veto. Control the inputs to the fabric, and you control whether anyone can commoditize your compute with it.

## VI. What Founders and Architects Should Do Before 2027

Treat the second-order bottleneck as your planning baseline, not an edge case. If your roadmap assumes optics arrives frictionless in 2026, it is mispriced. Broadcom is pushing VCSEL-based Near-Package Optics as a bridge to scale AI clusters before full CPO matures [3]. That phasing is a signal to design for a staged transition rather than a clean cutover.

Architect for the Optical Blast Radius from day one. Favor ELSFP-style separation so a cheap failure never strands expensive silicon.

Model your exposure to two chokepoints you do not control: InP laser supply and TSMC/Samsung packaging allocation. Secure position early rather than assuming spot availability.

Watch the Fabric Value Inversion if you build accelerators. The fabric layer is where margin is migrating, and yield hurdles in fiber alignment and packaging are what will gate the timeline.

The companies that win the next cluster generation will not be the ones with the best photonics. Hardware buyers who fail to secure InP laser supply by Q3 2026 will find their $40,000 GPUs sitting idle in warehouses, waiting on a $10 part that simply does not exist.

KEY FINDINGS

- Next-generation XPUs need bandwidths near 28.8 Tbps, the point where copper interconnects and electrical SerDes die on power alone.
- Co-packaged optics runs on the same CoWoS and SoIC packaging lines that HBM stacks and AI accelerators already fight over at TSMC.
- Silicon cannot generate light, so every high-performance optical transceiver leans on Indium Phosphide lasers from outside the silicon supply chain.
- Veeco booked more than $250 million in InP laser tooling orders, with deliveries starting in 2026 and accelerating through 2027.
- One laser failure can strand a $40,000 GPU or $200,000 switch, so the industry is turning to pluggable external laser sources.

REFERENCES


[1] The Era of Light: Photonic Interconnects Shatter the 'Copper Wall' in AI Scaling. http://business.bigspringherald.com/bigspringherald/article/tokenring-2026-1-9-the-era-of-light-photonic-interconnects-shatter-the-copper-wall-in-ai-scaling

[2] Xscaping Limits Blog Series #3: Xscape the Copper Wall.

[3] Beyond the Copper Wall: Scaling AI Clusters with VCSEL-Based Near-Package Optics (NPO), Broadcom Inc. https://www.broadcom.com/blog/beyond-the-copper-wall-scaling-ai-clusters-with-vcsel-based-near-package-optics-npo-

[4] Ayar Labs In-Package Optical I/O Solution, Open Compute Project; TeraPHY Optical I/O Chiplet, Ayar Labs. https://www.opencompute.org/chiplets/46/ayar-labs-in-package-optical-io-solution

[5] Lightmatter CEO Nick Harris: AI's Next Scaling Challenge Is Interconnect, Not Compute. https://www.eetimes.com/podcasts/lightmatter-ceo-ais-next-scaling-challenge-is-interconnect-not-compute/

[6] [News] Silicon Photonics Race Intensifies as TSMC Targets 2026 COUPE Production, TrendForce. https://www.trendforce.com/news/2026/04/01/news-silicon-photonics-race-intensifies-as-tsmc-targets-2026-coupe-production-samsung-eyes-2029-cpo-turnkey/

[7] Electronic Chip Package and Co-Packaged Optics (CPO) Technology for Modern AI Era: A Review, PMC. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12029643/

[8] TSMC COUPE: The Underappreciated Platform Layer for AI Photonic Interconnect; TSMC COUPE: Why the CoWoS Pattern Is Repeating in Silicon Photonics. https://www.atlaspeakresearch.com/report/66ebb8

[9] TSMC targets 2029 for panel-level CoPoS packaging in AI chip push, DigiTimes. https://www.digitimes.com/news/a20260630PD239/tsmc-supply-chain-cowos-packaging-equipment.html

[10] TSMC's COUPE push puts Samsung's silicon photonics ambitions under pressure. https://www.digitimes.com/news/a20260520PD215/samsung-silicon-photonics-tsmc-production.html

[11] Indium Phosphide (InP): The Quiet Bottleneck Behind the AI Infrastructure Boom, YIAZOU. https://yiazou.com/indium-phosphide-inp-the-quiet-bottleneck-behind-the-ai-infrastructure-boom-2/

[12] Veeco Announces $250 Million+ in Equipment Orders for Manufacturing Indium Phosphide Lasers. https://ir.veeco.com/news-and-events/news-details/2026/Veeco-Announces-250-Million-in-Equipment-Orders-for-Manufacturing-Indium-Phosphide-Lasers/default.aspx

[13] Aeluma positive on QD lasers despite impact of US government shutdown, optics.org. https://optics.org/news/aeluma-positive-on-qd-lasers-despite-impact-of-us-government-shutdown

[14] [Supply Chain Anatomy] What is the CPO bottleneck? note. https://note.com/iqoys/n/n0a3a4605bf1a

[15] How Industry Collaboration Fosters NVIDIA Co-Packaged Optics, NVIDIA Technical Blog. https://developer.nvidia.com/blog/how-industry-collaboration-fosters-nvidia-co-packaged-optics/

[16] How Industry Collaboration Fosters NVIDIA Co-Packaged Optics, NVIDIA Technical Blog. https://developer.nvidia.com/blog/how-industry-collaboration-fosters-nvidia-co-packaged-optics/

[17] Marvell Completes Acquisition of Celestial AI.

[18] Marvell to Acquire Celestial AI, Accelerating leadership in data center connectivity.

[19] Three Key Challenges Stand in the Way of CPO Commercialization, The Elec.