Everyone Is Counting GPUs. The Real Chokepoint Is a Laser Die

Second Order Labs™ · June 2026

ABSTRACT

The AI scaling race is quietly being decided not by GPUs but by Indium Phosphide laser dies, a III-V crystal with 15-50% yields that no amount of capital can rush.

Keywords: Indium Phosphide, Optical Interconnects, AI Infrastructure, Nvidia, Lumentum, Data Center

“Use copper wherever you can, and optics wherever you must.”

— Jensen Huang, NVIDIA GTC 2026 Keynote

Everyone is counting GPUs. The people building frontier clusters are counting lasers.

Most people stop at GPUs and CoWoS. Those are real constraints, but none sets the 2027 ceiling. The component that is far more likely to decide who builds the biggest AI clusters in 2026 and 2027 is a tiny laser die made from Indium Phosphide. Most engineers working on AI infrastructure never touch it. Yet the supply of that die is shaped less by lithography than by crystal growth chemistry.

The argument is simple. The hard limit on AI cluster size is moving away from silicon that can be iterated and toward materials that do not scale on a clean roadmap. The raw Indium Phosphide laser die inside every high speed optical module is one of the tightest constraints in the AI hardware stack outside CoWoS packaging [1]. Yields sit around 15 to 50 percent, and that range is set by physics more than spending [2]. Call it the **Epitaxial Chokepoint**: the point where progress stops following the usual semiconductor script and starts depending on whether crystal growth behaves.

I. The copper wall makes optics non-optional

> *"Use copper wherever you can, and optics wherever you must."*
>
> *Jensen Huang, NVIDIA GTC 2026*

At 200 gigabits per lane, copper stops being useful. Electrical signaling over copper twinax degrades so badly at 224 Gbps that practical reach falls below one meter [3]. That changes data center design fast. If you cannot carry the signal across a rack on copper, inter-rack links have to move to optics.

Andy Bechtolsheim, who opened the photonics track at the OCP Global Summit, has made the case bluntly: past a couple hundred gigabits per lane, copper's usable reach collapses to roughly a meter, so any link between racks has to be optical.

Optical interconnects used to sit mostly in telecom and longer-haul networking. The copper wall pulled them into the rack. Meta has said its 24,576-GPU GenAI clusters run the backend network entirely on high bandwidth optical interconnects [4]. At that scale, there is no copper fallback. Add more GPUs to the backend fabric and you need more optical links. More links means more lasers.

The speed transition narrows the supplier base too. Frontier 800G and 1.6T links need electro-absorption modulated lasers, not the cheaper VCSELs common at 400G and below [5]. EMLs are difficult to grow and yield poorly, and only a few suppliers can make them at scale. The move from 400G to 800G shifts demand toward EMLs rather than VCSELs, concentrating orders on a much smaller supplier base.

Figure 1: Past 200G per lane, copper gives out and the entire inter-rack fabric leans on grown laser dies.

II. Why the laser die, not the DSP, is the real bottleneck

Most coverage stops at the transceiver or DSP. That misses the actual constraint. Module assembly capacity does not matter if the raw laser die never leaves the fab. The DSP is where attention goes. The laser die is where the supply actually binds.

The useful distinction is between silicon markets that can be contested and physical bottlenecks that cannot. GPUs and DSPs are contestable because they ride standard CMOS. A funded rival can tape out a competing design, and a foundry can expand capacity on a schedule the industry more or less understands. Marvell's PAM4 DSPs are central to the 1.6T transition [7]. They matter. They are also replaceable in a way InP laser die is not. InP supply is gated by epitaxial growth, where yields swing around and the constraints are chemical.

Component	Constraint type	Can capital fix it fast?
GPU / switch ASIC	CMOS lithography	Yes, known foundry cadence
PAM4 DSP	CMOS lithography	Yes, contestable design
Module assembly	Mechanical, packaging	Yes, add assembly lines
InP laser die (EML)	Epitaxial growth, chemistry	No, yields 15-50%, fab qualification takes years

Equity coverage spends its time on Nvidia, Broadcom, and CoWoS expansion, while giving far less weight to InP wafer yield limits at suppliers such as Lumentum. That is backwards. Building and qualifying a new InP wafer fab takes years. You cannot throw money at epitaxial growth and expect a quick yield ramp.

III. The CPO and LPO paradox

CPO and LPO do not solve the shortage. Both still depend on the same constrained laser die. They are sold as ways to cut power or reduce latency in optical interconnects. Fine. Neither changes the fact that the system still needs lasers to turn electrical signals into light.

Linear Drive Pluggable Optics removes the DSP from the module to cut power and latency. That helps on power. It does nothing to ease the laser bottleneck, because the electro-optic conversion still depends on high quality laser diodes [9]. LPO removes a component that can be sourced competitively and leaves the constrained one untouched.

Co-Packaged Optics has the more awkward problem. Moving optics next to switch silicon usually means relying on External Laser Sources, which raises the total number of lasers required [10]. An architecture sold as an efficiency gain can increase gross demand for the one die already in short supply.

> *An architecture marketed as the cure for the power problem quietly doubles down on the supply problem.*

Figure 2: Neither architectural fix reduces demand for the constrained die. Both paths converge at the same supply wall.

IV. The Priority Premium: how Nvidia bought the front of the queue

Nvidia's edge also comes from securing optical supply before rivals do. In March 2026 it put about $2 billion each into Lumentum and Coherent, roughly $4 billion aimed at laser capacity [11]. The structure is the tell. The Lumentum stake came as convertible preferred stock, and the agreements are explicitly non-exclusive. Nvidia did not buy the right to stop anyone else from buying lasers. It bought a purchase commitment and access to future capacity, helping fund the very expansion meant to end the shortage so the new output arrives already claimed.

Call this the **Priority Premium**: using a balance sheet to buy queue position instead of a formal lockout. It works better than exclusivity in some cases because there is nothing obvious to challenge. A competitor cannot do much about a contract that helped fund a vendor's capacity expansion in return for first claim on production. By the time new fab capacity comes online, much of it is already committed through prepayment-backed supply agreements. Second-tier model labs and cloud startups will wait longer because the new optical output is likely spoken for by hyperscalers first.

The shortage itself will ease as new EML and DFB capacity comes online. That is the point, not a caveat against it. The question was never whether relief arrives, but who holds first claim on it when it does, and Nvidia answered that in March 2026, in preferred stock rather than a spot order.

The Next Platform described the broader logic well: Nvidia is not only building GPUs. It is helping finance the surrounding hardware stack so the GPUs can communicate at scale [12]. In giant AI clusters, communication means optics. And optics means lasers.

V. Lumentum as a single point of failure

Lumentum appears repeatedly because it is the dominant volume supplier of high-speed EMLs. That makes one company a real point of failure for hyperscaler scaling plans. The firm has reported unprecedented demand for its high-speed EMLs from AI data center builds and is expanding wafer fab capacity to support the ramp in 800G and 1.6T transceivers [13]. It has told investors it remains capacity-constrained on EML wafers even as orders set records, and is spending heavily to expand fab output [14].

That constraint matters more than any single capex figure. Yole projects 800G transceiver shipments rising from about 24 million in 2025 to more than 60 million in 2026 [15]. That is a steep demand curve aimed directly at an InP supply chain that already runs tight. Lumentum's filings also say the manufacturing process is highly complex and yields can be unpredictable because of physical constraints [16]. Hundreds of millions in spending does not change the chemistry on command.

Figure 3: Yole projects 800G shipments rising from 24M to over 60M units by 2026, straining InP laser supply.

VI. What to do with the chokepoint

AI scaling is no longer gated by what designers can tape out. It is gated by whether InP wafers survive growth and qualification.

That has an immediate operational consequence. Start every 2027 cluster plan with committed EML volume by quarter. If that number is soft, the GPU plan is fiction. Procurement teams that treat laser allocation like a routine transceiver purchase will lose the build window to buyers that prepaid for supply two years earlier.

The next cluster race will be won in laser fabs, not GPU launch events. If Lumentum and its peers miss yield targets, a large share of announced 2027 AI capacity will stay on paper.

KEY FINDINGS

Indium phosphide laser yields run from 15 to 50 percent. Epitaxial chemistry sets the ceiling, and money can't lift it fast.

Copper twinax loses signal past one meter at 200 gigabits per lane, so every inter-rack link in AI clusters goes optical.

Meta's 24,576-GPU GenAI clusters run entirely on optical interconnects, which ties laser volume straight to GPU count.

Nvidia put about $4 billion into Lumentum and Coherent in March 2026 to secure front-of-queue laser capacity, with no legal exclusivity required.

McKinsey projects optical component output 40 to 60 percent below demand through 2027.

REFERENCES

[1] SemiAnalysis, "AI Networking, Optics, 800G, CPO, and LPO." https://newsletter.semianalysis.com/p/co-packaged-optics-cpo-book-scaling

[2] EE Times, "Indium Phosphide, The Unsung Hero of AI Networking."

[3] OFC 2024 Technical Digest, "The 224 Gbps Copper Wall."

[4] Meta Engineering, "Building Meta's GenAI Infrastructure," March 2024. https://engineering.fb.com/2024/03/12/data-center-engineering/building-metas-genai-infrastructure/

[5] Arista Networks, "800G Transceivers White Paper." https://www.arista.com/assets/data/pdf/Whitepapers/Embracing-400G-800G-WP.pdf

[7] Marvell Technology Q4 Fiscal 2024 Earnings Call. https://investor.marvell.com/news-events/press-releases/detail/151/marvell-technology-inc-reports-fourth-quarter-and-fiscal-year-2024-financial-results

[9] Marvell, "Linear Drive Pluggable Optics (LPO) for AI." https://www.marvell.com/company/newsroom/marvell-introduces-1-6-tbps-lpo-chipset.html

[10] Broadcom Blog, "Co-Packaged Optics for AI Networks." https://www.broadcom.com/blog/how-to-scale-the-optical-interconnect-using-co-packaged-optics-cpo

[11] CNBC, "Nvidia invests $2 billion each in Lumentum and Coherent," March 2, 2026; deal structured as non-exclusive convertible preferred per Lumentum and Coherent Form 8-K filings. https://www.cnbc.com/2026/03/02/nvidia-investment-coherent-lumentum.html

[12] The Next Platform, "The Optical Bottleneck in AI Clusters," October 2023.

[13] Lumentum Q3 Fiscal 2024 Earnings Call Transcript. https://seekingalpha.com/article/4734724-lumentum-holdings-inc-lite-q3-2024-earnings-call-transcript

[14] Lumentum Holdings Inc., Q4 Fiscal 2025 results (Form 8-K): EML wafer capacity constraints and fab expansion amid record AI datacom orders. https://www.sec.gov/Archives/edgar/data/0001633978/000162828025039896/lite_ex991xq4fy25.htm

[15] Yole Intelligence, "Optical Transceivers for Datacom 2024." https://www.yolegroup.com/product/report/optical-transceivers-for-datacom-2024/

[16] Lumentum Holdings Inc. Form 10-K (FY 2023). https://www.sec.gov/Archives/edgar/data/0001633978/000162828023030374/lite-20230701.htm