

LIGHTWAVE[®]

AI data center providers seek power and bandwidth promise in co-packaged optics

Co-packaged optics (CPO), a high-speed networking technology that integrates optical components (lasers, photodetectors) directly with switch/compute chips (ASICs) in the same package, continues to show promise. Advocates of CPO maintain that it reduces power consumption by over 80% and increases bandwidth density by shortening electrical traces to millimeters. Used primarily in data centers and AI/ML, CPO enables faster, more efficient data transfer with the necessary speed and efficiency for “scale-up” (within servers/racks) and “scale-out” (between racks) networks. Despite CPO’s promise, the technology is still in its early stages. A recent Lightwave survey revealed that while CPO (17%) and next-generation pluggable optics beyond ZR/ZR+ (14%) are drawing interest, adoption appears cautious. In this new eBook, we’ll examine CPO’s potential to scale AI clusters and racks beyond what copper can support, adoption realities, and how new double-sided wafer probe test solutions can address CPO validation that traditional methods can’t meet.

ID 11289813 | Data Center © Maximilian Pogonii | Dreamstime.com



CONTENTS

ARTICLES

Scaling Photonics Test for the Next Decade: Lessons from Wafer to Co-Packaged Optics3

Parallel Photonics Test – Unlocking Scalable PIC and CPO Manufacturing7

Co-Packaged Optics and the AI data center: From skepticism to strategic adoption 12

AI scale-up raises the stakes for transceivers and CPO adoption..... 18

ADVERTISEMENTS

Teradyne and Quantifi Photonics..... 21

RESOURCES

Teradyne..... 22

Quantifi Photonics..... 22

Scaling Photonics Test for the Next Decade: Lessons from Wafer to Co-Packaged Optics

Translating laboratory precision to semiconductor manufacturing workflows

ID 1098390 © Peter Rust | Dreamstime.com

BY MATT GRIFFIN, SENIOR PRODUCT MANAGER, OPTOELECTRONICS TEST, TERADYNE

COP MARKET - 2024-2035 REVENUE OF OES FO SCALE-OUT AND SCALE-UP

Source: Co-packaged Optics for Data Centers 2025 report, Yole Group




Figure 1. Co-Packaged Optics for Data Centers 2025
Yole Group

Silicon photonics (SiPh) and co-packaged optics (CPO) are moving beyond the domain of specialized research labs and prototype demonstrations. No longer niche technologies, they are now scaling toward high-volume

manufacturing as data center operators seek ways to move more data at lower power.

According to Yole Group, annual shipments of optical interconnects and co-packaged devices will grow to *hundreds of millions* over the next decade, driven by the power demands of AI infrastructure and high-performance computing. These demands reflect

how today's AI systems scale "up" within GPU nodes as well as "out" across larger multi-rack clusters, expansion that places unprecedented strain on electrical interconnects and accelerates the shift toward photonics.

The power savings promised by CPO are striking. NVIDIA and Broadcom cite energy savings of 3.5x when replacing copper with optical interconnects, and IBM has reported reductions as high as 5x for certain configurations. These are exciting improvements and explain why optical engines and CPO modules are moving quickly toward mainstream deployment.

But as device volumes climb and power envelopes rise, test is the real challenge.

Traditional optical validation methods were designed for low-volume, precision-oriented laboratory environments. Production-scale photonics demand something fundamentally different.

The industry is now grappling with how to translate the precision of optical measurement into the efficiency, repeatability, and throughput required for semiconductor manufacturing. Once a “contained” step at the end of a process, test has become a central engineering function that determines how quickly photonics technologies can be industrialized. Nowhere is this more evident than at wafer sort, a structural bottleneck for SiPh manufacturing. Scalable test at this insertion is a prerequisite for achieving the known-good-die economics of semiconductor manufacturing.

Integrating test processes for automation, speed, and consistency

The growth in demand is matched by growth in complexity. Modern photonic devices combine dense electrical and optical interfaces, and their performance depends on the seamless interaction of those domains. Channel counts have risen from tens to hundreds or even thousands per device, and electrical signaling continues to accelerate.




Figure 2: Co-packaged optics (CPO). CPO is a packaging innovation that integrates silicon photonics chips with data center switches or GPU computing devices onto a single substrate

Many co-packaged modules integrate four to 32 optical engines, each requiring precise optical alignment and electrical characterization, and housed in thermal environments that regularly exceed a kilowatt of dissipation. All of this must be tested as an integrated system, not in isolation.

Traditional optical validation, designed for small-batch laboratory work, is insufficient. Photonics test built around slow, expert-driven procedures using manual alignment and single-channel instrumentation is a model that cannot scale into a production-level environment.

For photonics to achieve semiconductor-class manufacturing, photonics test must deliver semiconductor-class automation, speed, and consistency. Test must integrate electrical and optical domains at every insertion point from wafer to optical engine to fully assembled CPO module. The process must also produce data that supports yield learning across high-volume production.

These are the challenges that have guided Teradyne’s work with Quantifi Photonics and the larger photonics ecosystem. Together, these companies have collaborated to bring several test innovations out of the lab and into production

environments, most notably Teradyne Photon 100, the industry's first high-volume, double-sided wafer probe test solution.

High-volume, double-sided wafer test

The platform probes the wafer electrically from the top and couples light into it from the bottom. A straightforward approach, but in practice, its impact is substantial.

Wafer sort is where SiPh economics succeed or fail, and Teradyne Photon 100, integrated with Teradyne UltraFLEXplus, makes wafer test truly production-capable. The system enables full electro-optical characterization in a single automated step, without moving wafers across multiple setups. The entire process takes place within one mechanical environment that eliminates drift, reduces handling, and improves data consistency. By enabling truly parallel electro-optical characterization across many wafer sites, Teradyne Photon 100 brings wafer-level throughput much closer to semiconductor norms.

What elevates this architectural improvement to a production-ready solution is its ability to perform multi-site opto-electrical alignment, always one of the most difficult parts of photonics test. A lab engineer might spend minutes adjusting a single fiber to find the right coupling point. In manufacturing, this is not a sustainable approach.

The new method automates alignment using robotic motion, machine vision, fiber-array inspection and cleaning, and real-time optical feedback. This system simultaneously aligns and tests multiple sites on a wafer while maintaining sub-micrometer precision.

A previously time-consuming manual step is transformed into a predictable and repeatable

process that can be executed across shifts, operators, and facilities. This level of automation mirrors semiconductor workflows in their ability to manage probe stability and site-to-site repeatability.

The most far-reaching outcome of this integration may be its impact on data quality and how that affects reliability and yield learning. Capturing optical and electrical measurements under identical alignment, temperature, and mechanical conditions delivers a tightly correlated dataset that reveals how these domains interact and influence each other. When electrical and optical results share timestamps, metadata, and environmental context, they form a diagnostic signature that is far more informative than isolated measurements.

For example, some deployments have demonstrated small variations in wafer-level optical signatures, such as wavelength drift or power uniformity; these in turn have been linked to downstream performance issues in assembled CPO modules.

This type of correlation was difficult to uncover when optical and electrical domains were tested separately. The data produced by Teradyne Photon 100 provided an early indicator of reliability, allowing design teams to make evidence-based decisions about device structures, materials, and process windows.

It also moves test into a more proactive role, where it informs design and process decisions rather than reacting to them. The same test architecture extends from wafer to engine to module; this keeps data structures consistent across insertions, strengthening correlation and accelerating yield learning.

Reshaping expectations for product and process design

Deploying such systems has taught us a great deal about the realities of scaling photonics test. A key lesson is that automation is about consistency as much as throughput. Manual optical alignment is an impressive skill, but it introduces variation. Automated multi-site alignment and handling significantly reduce variation, produce cleaner data, and prevent risk, especially as device volumes scale.

The most important takeaway is that photonics test is inherently multidisciplinary. Electrical performance, optical coupling, thermal behavior, and mechanical tolerances all interact, and scalable photonics test requires these engineering disciplines to operate as a unified function. Historically, each domain optimized its own component of the problem, but today that kind of separation only slows progress.

Teradyne Photon 100 is emblematic of a new design philosophy in which test is not an afterthought but an active influence that shapes device architecture. Companies can smooth the path to high-volume yields by designing with test in mind; for example, positioning optical I/O for automated coupling, structuring metallization for robust probing, and planning thermal pathways for stable measurement. This alignment is becoming increasingly important as fabs expect optical test insertions to operate with the same repeatability and automation used in electronic wafer test.

These new insights inform how the industry must navigate the next decade. CPO modules and photonic engines will become even more complex, integrating more optical interfaces, higher electrical bandwidths, and significantly higher thermal loads.

With devices approaching or exceeding 1,000 watts per package, optical performance must be evaluated under real operating conditions rather than ambient or lightly loaded scenarios. Optical alignment must hold steady under thermal expansion; electrical characterization must accommodate dynamic shifts in power levels; and measurement frameworks must combine optical, electrical, and thermal data in real time. Here, test will reveal behavior as much as it verifies performance.

Why the next decade? Architectures being built now enable the precision of the laboratory to coexist with the efficiency and resilience of the factory. These innovations are designed to support multiple generations of photonic devices, well into the 2030s.

In the end, the speed at which the photonics industry can commercialize next-generation architectures will depend on both the creativity of its designs and the maturity of its test systems. Integration of optical, electrical, and thermal awareness into a continuous, automated test process will define what is possible in high-volume silicon photonics and co-packaged optics.

The organizations that unify these domains into a single automated, coherent workflow will set the pace for how quickly SiPh and CPO technologies can reach true commercial scale. This is the foundation of the next decade, and the path that will carry photonics into its industrial future.

About the Author

[Matthew Griffin](#) is the Senior Product Manager for Optoelectronics Test at Teradyne. Over the past eight years, he has led advancements in test technologies for silicon photonics, co-packaged optics, microLED, ADAS camera interfaces, and machine vision for production inspection.

Parallel Photonics Test – Unlocking Scalable PIC and CPO Manufacturing

ID 11079615 | Data © Fernando Cortés | Dreamstime.com

BY KEES PROPSTRA, VP OF PRODUCT MARKETING, QUANTIFI PHOTONICS

With contributions from Teradyne and ecosystem partners

Photonics is entering a new phase of industrialization. What began as a domain of precision laboratory experiments has matured into the foundation of large-scale data movement in high-performance computing, AI training clusters, and hyperscale data centers. As system designers push architectures to scale both “up” within GPU nodes and “out” across multi-rack fabrics, copper interconnects are hitting hard limits of power, reach, and density all at once. Photonic integrated circuits (PICs) and co-packaged optics (CPO) have emerged as the path forward, bringing higher beachfront density, multi-wavelength engines, and $\geq 100\text{G}$ high-speed I/O to the center of system design.

Yet these advances highlight a difficult truth: the industry cannot deploy PICs or CPO at scale until it can test them at scale.

While electronics has a long head start, with decades of automation and standard

workflows shaping a test ecosystem that naturally supports high volume, photonics developed differently. Optical test grew out of environments where precision and flexibility mattered more than throughput, an approach that just doesn’t work once devices leave the lab. Bridging that gap is now one of the industry’s most urgent challenges.


Reframing Traditional Test for Industrial-Scale Photonics

As PICs and CPO devices become denser, the limits of legacy test approaches become obvious. Wafer-based silicon photonics, in particular, raises the bar. Every die needs to be tested. Optical coupling must be repeatable across many sites. Electrical and optical measurements must remain stable as temperature and mechanical conditions change. And all of this must happen fast enough to make economic sense, keeping pace with semiconductor-scale throughput.

Many teams have learned (sometimes the hard way) that problems are much easier to solve when they are discovered earlier rather than later. That's driving a broader shift toward "shift-left" thinking. Instead of treating test as something to optimize once designs are finished, engineers are accounting for it during architecture and layout. Features that support alignment, predictable optical access, and test reduction are increasingly designed in from the beginning.

Parallel Test as a Practical Necessity

In early development, engineers still need broad visibility. Passive, active, and high-speed characterization is required to reveal and understand device behavior. As designs stabilize in design verification testing (DVT), the emphasis shifts to automated, volume testing over process, voltage and temperature. By the time products reach high-volume manufacturing (HVM), only a small subset of measurements is typically required, providing good



The "Shift-Left" approach – test insertions from wafer level screening through die test, assembly validation, and module characterization to final product testing - illustrating how shift-left strategies catch issues early and reduce costly failures downstream.

This is in addition to more traditional 'shift-left' thinking. More extensive testing in development and validation allows for a reduced test regimen in high-volume manufacturing, which we refer to as the testing funnel. In the manufacturing process testing as early as possible to screen potential failures while avoiding costly failures downstream is more generally referred to as shift-left.

This isn't a philosophical shift so much as a practical one. Planning your testing strategy early in the development cycle and in general testing more extensively early in the development and manufacturing process are a recipe for success in the pursuit of cost-effective, ultra-high-volume manufacturing of photonic ICs and co-packaged optics.

performance and high yield. Quantifi Photonics' test-funnel approach reflects this reality by maintaining continuity from early development through high volume manufacturing (HVM), while progressively tightening and automating the test set.

Once devices include more than a handful of optical paths, serial testing is no longer viable. Scaling photonics test is fundamentally a parallelization problem. Parallel test architectures and modular instruments make this possible. Multi-channel laser sources, high-density optical power meters, and parallel high-speed I/O testers allow multiple channels to be stimulated and measured at the same time. For PICs and CPO engines with 8, 16, or 32 optical paths, this capability isn't a nice-to-have - it's essential.

Just as important, the same measurement hardware can now be reused across multiple stages. A power meter used on a benchtop to evaluate an early device can later appear

in a design-verification rack and eventually in an automated wafer-sort or module-test cell. This continuity keeps measurements correlated across the product lifecycle and eliminates the need to reinvent test infrastructure with every transition from the lab to the factory.


What Parallel Test Looks Like on the Floor

In development labs, modular PXI-based optical instrumentation gives teams the flexibility to experiment as devices and test requirements evolve. As programs move forward, those same modules support multi-site design-verification setups, where stability and repeatability matter just as much as measurement coverage.




In addition to a modular PXI-based approach, high-channel-count parallel testing is enabled with scalable test platforms such as the QCA/QCR electrical and optical sampling oscilloscope platform (left), and the Power 1410 high density power meter (center). The ELSFP-compliant external laser source (right) is the latest addition to the PXI portfolio.

Production environments raise the stakes even further. Modern wafer-sort cells may combine



The "Photonics Test Bench" forms the core building block for validating next generation 1.6T and 3.2T optical engines and CPO architectures. This configuration supports active alignment, insertion loss, ring resonator characterization, eye diagrams/TDECQ, polarization optimization, and receiver sensitivity testing – with a compact PXI and benchtop test instrument setup - pictured right.



dozens of optical channels, multiple laser sources, and parallel power meters in a single automated system. Electrical probing and optical coupling happen together, which dramatically improves throughput for silicon photonics testing. Without that level of parallelism, wafer-level test quickly becomes the bottleneck.

Module-level test introduces different constraints. Optical engines and subassemblies must be validated as systems, often under thermal load and at full data rate. Parallel high-speed testing makes it possible to exercise multiple lanes at once, evaluating eye quality, bit-error rate, and receiver sensitivity without extending test time beyond practical limits. For 3.2T and 6.4T-class engines, this approach is a hard requirement for cost-effective, high-throughput testing and manufacturing.

From Characterization to Screening

As architectures densify, the value of the test funnel becomes clearer: its ability to convert deep early characterization into a minimal, production-ready screening strategy. R&D may require dozens of measurements, while statistical correlation allows high-volume manufacturing to rely on a far smaller set of predictive tests.

Quantifi's modular architecture supports this exploration-to-production continuity. When instruments, APIs, and data structures remain aligned throughout test processes, correlation becomes significantly easier. Manufacturers gain both speed and clarity in understanding which measurements truly matter for yield and reliability.

Where Parallelism and Economics Come into Play

Wafer-level test has become one of the most critical points in silicon photonics manufacturing. Occurring at the top of the manufacturing process, characterization must be a fast and efficient screening process to prevent faulty wafers from progressing downstream. PICs require electro-optical validation at 100% coverage, and the economics of high-volume SiPh hinge on efficient known-good-die workflows. At wafer sort, parallel test becomes economically mandatory.

Module-level test sits between wafer processing and full CPO assembly, and introduces its own complexity. Optical engines and modules must be

validated as assembled systems, combining optical connectivity, thermal management, and high-speed digital analysis. Parallel testing allows these factors to be evaluated together, ensuring that only modules meeting full performance requirements move forward, where failures are far more costly.

Why Ecosystem Cooperation Matters

Scaling photonics test is not a single-vendor problem. It requires an ecosystem in which optical measurement, ATE-level automation, device handling and precision alignment technologies operate as one cohesive system. Quantifi's acquisition by Teradyne reflects this convergence, pairing modular optical instrumentation with system-level automation and data management.

This cooperation is critical as PIC and CPO devices enter volume production, creating the semiconductor-like discipline required to achieve scale. When optical instrumentation is integrated with ATE-class automation and shared data infrastructure, manufacturers can



No single company can solve scalable photonics test on its own. Teradyne's UltraFLEXplus platform and recently announced Photon 100 integrates Quantifi Photonics PXI-based optical instruments and connects to the customer's semiconductor probe of choice to enable automated optical-electric testing for SiPh and CPO manufacturing.

maintain consistent test strategies from wafer to module to final package. That consistency enables correlation, accelerates yield learning, and ultimately makes scaling more predictable.

Looking Ahead

The next decade will shape how far silicon photonics and co-packaged optics can go. Devices will continue to add channels, operate at higher speeds, and dissipate more power. Meeting those demands requires test architectures that are parallel, modular, and grounded in manufacturing reality.

Test is no longer just a checkpoint at the end of the line. It's becoming a constraint designers must work with, and increasingly, a source of competitive advantage. As photonics manufacturing matures, the ability to validate complex devices efficiently and at scale will play a decisive role in which architectures succeed.

About the Author

[Kees](#) is VP of Product Marketing at Quantifi Photonics and has over two decades of experience across photonics R&D, product management, and technology marketing. He previously held roles at JDS Uniphase, Tektronix, and MultiLane, and holds an MSc in Applied Physics and a postgraduate degree in Modern Optics from the University of Twente.

Co-Packaged Optics and the AI data center: From skepticism to strategic adoption

CPO's implications aren't just about technology; they also reflect a range of issues, including user psychology, risk tolerance, data center culture, and the changing relationship between infrastructure buyers and their suppliers.

BY LAWRENCE GASMAN / CIR

Co-packaged optics (CPO) has become one of the most talked-about technologies in the AI data center world. Vendors and standards bodies aggressively position CPO as the answer to AI's bandwidth, latency and power crises, yet many users remain in a quandary, wondering whether they need CPO; whether it is mature enough, and whether adopting CPO will create more operational risk than performance benefit. CIR has recently completed [a study of CPO's prospects given the current mix of technical prospects and user skepticism](#).

CPO demos at OFC and ECOC, and perhaps these days GTC, are impressive. Yet, if you walk into most data centers today—outside of a small number of hyperscale facilities—you will find few CPO deployments. What you find if you talk to data center managers is a mix of cautious

curiosity, engineering-led skepticism, and perhaps an industry quietly preparing for a technology transition that may take a decade to fully play out.

However, CPO is not just a technology story. It is a story about user psychology, risk tolerance, data center culture, and the changing relationship between infrastructure buyers and their suppliers.

Why CPO is back in the spotlight

CPO is not new. The concept – packaging photonics and electronics close together – goes back years, to early work by IBM on supercomputer interconnects and to “flyover” interconnect concepts. What has changed is the emergence of AI as the defining workload of modern data centers. Before the AI boom, CPO was pitched as a

broadly applicable innovation—useful for sensors, telecom, high-performance computing, and data center interconnect. This “shotgun” positioning generated early excitement but ultimately failed to sustain demand.

AI-driven demand is more focused. CPO is framed as a solution to specific and real problems in AI data centers: electrical interconnects are becoming too lossy and too power-hungry as the industry moves toward 112G and 224G SerDes, and as switch ASICs approach 51.2T and beyond. In that context, CPO’s value proposition becomes easy to explain. By shortening electrical paths and moving optical interfaces closer to the ASIC, CPO promises better energy efficiency, improved bandwidth density, and a path toward future ultra-high-radix switches (51.2T to 204.8T). In the table below we chronicle the consensus view on where CPO can best show off its apparent advantages.

CPO equivalents and easier to service. Network managers from a couple of decades ago would be shocked that pluggability could be so easily disposed of. Some level of confidence might be re-established by pointing out that the disappearance of pluggability may also lead to fewer field failures.

And for all the talk of CAPEX savings, early generation CPO switches may cost more than conventional pluggable switches. Caveat emptor!

How potential CPO users really feel:
The hyperscalers tale

Among data center operators, there is a divide between fascination and distrust about CPO. Operators acknowledge that CPO appears technically compelling but also see CPO as something that could turn out as an operational headache. CIR’s research suggests that even within the traditional data center

Advantages Currently Claimed for Co-Packaged Optics (CPO)	
Power efficiency	Shorter traces among ASICs and optics, reduction in DSP power and lower I/O drive requirements. These may result in a 20–40% reduction in interconnect power at 800G to 1.6T and ~5–15W saved per 800G port plus a 200–500W reduction per switch. These savings aggregate to multi-MW levels at hyperscale, with thousands of switches. These factors also reduce cooling needs.
CAPEX	Fewer power distribution upgrades and less need for cooling plant expansion. Also, CPO may enable higher rack density and fewer switches producing the same aggregate bandwidth.
OPEX	Lower costs for electricity and cooling are possible.

CIR

This seems too good to be true, and this may be the case. The mistrust of potential CPO users is linked to CPO being initially more complex than pluggables and the front panel design of conventional data center boxes being simpler than

market (i.e., non-hyperscale), awareness of CPO remains limited. Outside hyperscale environments, the average data center manager may know little about CPO. The same can be said of established pluggable transceiver suppliers as well.

From an engineering perspective, CPO seems to matter most when the longer-term goal is scaling to 102.4T ASIC. Then power becomes the limiting factor, with hints that “extreme” port densities will be required. This tends to mark CPO as a hyperscaler technology. Indeed, Microsoft, Meta, Google, and Amazon, are already running internal trials of CPO. It’s not experimentation for its own sake. They are looking for anything that can save them from higher power budgets in the future. The hyperscalers see CPO as part of a broader architectural shift: photonic fabrics, denser racks, and the possibility of scaling AI clusters and racks beyond what copper can support.

In this sense, hyperscalers view CPO not as an isolated technology upgrade, but as a key enabling element of the next AI infrastructure generation. Also, unlike the owners/managers of enterprise data centers and more modest cloud and edge data centers, hyperscalers’ willingness to adopt is reinforced by their already accepting nontraditional supply chains. Unlike enterprises and smaller data center operators, hyperscalers are comfortable with vendor lock-in if the performance gains justify it. They have procurement leverage, engineering teams capable of designing around supplier weaknesses, and, in many cases the ability to demand custom solutions. For hyperscalers, the question is not “Should we deploy CPO?” but “How quickly can we industrialize it?”

Enterprise and colocation operators: “Show me reliability”

CIR notes that outside hyperscale environments there is little evidence of CPO deployments today, despite references to smaller users in the trade press. Even if some small-scale CPO deployments do exist, they are not widely visible

or influential enough to make much difference. Enterprise and colocation operators and other smaller operators have a fundamentally different culture to the hyperscalers.

They are not building large proprietary platforms, and they rarely have the engineering staff to run complex optical integration programs. Their attitudes toward CPO are shaped by a different set of priorities: interoperability, multi-vendor supply chains, and field serviceability. This will create different adoption curves. Hyperscalers may jump early, while the rest of the market waits for “proof,” standardized interfaces, and a mature ecosystem. In effect, hyperscalers may act as the industry’s test lab, and enterprises and other smaller operators will be the eventual volume market.

The “Bridge Technology” psychology: LPO and NPO as comfort zones

Meanwhile, one of the most important trends shaping user attitudes is the rise of “stepping stone” solutions. CIR emphasizes that cautious potential CPO users will not leap directly from pluggables to full CPO. Instead, they will adopt intermediate architectures such as NPO and LPO, which provide some benefits of reduced power and improved signal integrity without fully sacrificing modularity.

Operators currently move gradually because they don’t trust early-generation CPO manufacturing yields, thermal behavior, or repair/maintenance models. NPO and LPO allow them to experiment with shorter electrical traces, reduced DSP overhead, and emerging electrical interfaces such as CEI-112G and CEI-224G without rewriting operational playbooks overnight.

LPO appeals strongly to operators focused on power and latency. By removing the DSP, LPO promises lower power consumption and reduced latency, valuable for AI. But it also introduces constraints: shorter reach, stricter host requirements, and tighter signal budgets. NPO, offers proximity benefits without full co-packaging, reducing risk in thermal and manufacturing complexity.

These bridge technologies matter because they will shape the pace of CPO adoption. CPO is “the endgame,” but for enterprise operators, CPO is seen as “next decade technology.” Many believe that even if CPO becomes important, intermediate steps may provide benefit without the risks of full co-packaging.

Thermal reality and the return of pluggability

The most frequently cited technical barrier to CPO adoption is thermal management. Temperature instability causes wavelength drift, accelerated aging, and performance degradation. The very act of bringing optics closer to the ASIC introduces heat-related risks. Optical components—especially lasers and photonic ICs—have strict temperature requirements. The CIR CPO report cites thermal management as one of the biggest factors discouraging CPO adoption today.

However, the most interesting aspect of CPO lasers is not the paradox that they come with their own thermal concerns, but that they might be bringing back pluggability through the back door. As things now stand External Laser Small Form Factor Pluggable (ELSFP), driven by OIF implementors’ agreements, represents a compromise between full CPO integration and traditional modular optics. The logic is simple: lasers fail, lasers degrade, and lasers are best kept in cooler areas. External laser sources

allow replacement without disturbing the switch ASIC package – the return of pluggability in a sense.

From a user attitude standpoint, ELSFP is appealing because it addresses the “repair anxiety” that haunts CPO discussions. Operators may not mind losing the ability to swap out an optical engine if they can at least swap out the laser source. This alone makes the CPO model feel less fragile. However, the external approach also introduces new risks, such as insertion loss and the possibility that one laser failure could affect multiple channels. ELSFP and external laser architecture are likely to play a major role in easing adoption.

Vendor influence: Broadcom and NVIDIA shape perceptions

User attitudes toward CPO are also shaped by the credibility of its champions. Here we note that Broadcom and NVIDIA have emerged as the influential suppliers driving CPO narratives. Broadcom’s early work with its Bailly platform established it as a reference point for switch ASIC integration. NVIDIA, meanwhile, has brought CPO into mainstream AI conversation by integrating it into its Spectrum-X and Quantum-X platforms and then showcasing these systems. publicly.

NVIDIA’s approach is particularly interesting since it reflects an awareness of operator concerns. Its architecture includes detachable optical sub-assemblies, implying a partial modularity model. In effect, NVIDIA appears to be designing CPO systems with manufacturability and replaceability in mind, acknowledging that pure co-packaging without serviceability would be a hard sell.

Vendor strategies matter because users often will not adopt new infrastructure technology until they believe a “safe vendor path” exists. In

networking, trust is often brand driven. Operators will tolerate risk if they believe the supplier can absorb it through engineering support and long-term product stability. Broadcom and NVIDIA are therefore not just suppliers—they are CPO confidence engines.

The supply chain anxiety: “Will this become another lock-in trap?”

That said, CPO changes procurement in ways that make operators uncomfortable.

With pluggables, operators buy optics from multiple suppliers and treat them as interchangeable commodities. CPO threatens this model. If optics are integrated into the switch package, the operator becomes dependent on the switch vendor’s packaging ecosystem and replacement policies. Instead of purchasing interchangeable modules, customers may have to source integrated CPO systems from a single vendor or partnership. Another issue is that with CPO there is the possibility that an optical failure might require replacement of a board, a line card,

or even an entire switch assembly. This is not the kind of truth that data center managers want to hear. CPO violates the instincts of operations teams. Many will interpret CPO as “vendor lock-in disguised as innovation.”

This is why interoperability work at organizations like the OIF and the Advanced Photonics Coalition matter. Users are not only watching for performance—they are watching for ecosystem maturity and multi-vendor credibility.

How adoption will actually evolve: A three-stage pattern


CPO market revenue will grow as operators gain confidence in equipment testing, supply chains, cable management, and cooling. As with all such technologies, the growth of CPO, if it is successful, will be non-linear.

From skepticism to acceptance (2026–2028): CPO will be a plaything for hyperscale AI clusters not a mainstream networking technology. The continued scaling of AI fabrics will force more

serious evaluation of CPO. The limiting factor will be unresolved issues in lasers, packaging yield, thermal design, and testing.

From acceptance to dependency (2029–2032): The next phase will shift from experimentation to reliance. As AI clusters scale toward 100T-class, CPO becomes the only technology to date that

Summary of CPO Markets by Type of Data Center (\$ Millions)



www.cir-inc.com

can do the job. User attitudes may shift from “we are evaluating” to “we cannot scale without it.”

From dependency to optimization (2032–2035): Once CPO is mainstreamed, the conversation changes. It will now focus on which vendor’s CPO architecture is better. This will be the period when non-hyperscale operators begin adopting CPO in meaningful volume,

We are aware our numbers for CPO may seem optimistic to some. The doubters recall networking technologies that have long disappeared without a trace – FDDI springs to mind. One important factor, CIR believes, is the evolution of AI itself. It is possible that smaller language models could reduce the need for high-speed interconnects, leaving CPO as a niche technology. If AI workloads become distributed or less bandwidth-hungry, the urgency of CPO may also decline.

Yet another wildcard is copper in the rack. NVIDIA’s continued use of copper for NVLink reinforces the industry’s long-standing pattern: fiber is adopted only when copper fails. CIR notes that an optical takeover is not in the stars. Fiber will penetrate the rack, but no one knows for sure how fast or how deep.

Coda: CPO as a cultural shift

CPO’s success depends on more than bandwidth density and power-per-bit metrics. It depends on whether operators can trust it. Right now, users are

intrigued but skeptical. Hyperscalers are moving forward because some of them suspect that CPO will be the only long-term strategy for scaling AI fabrics, even if CPO disrupts the service model that has defined optical networking for decades.

Over the next decade, the attitudes of operators will shift from “this looks risky” to “this is how modern AI networks work.” And when that shift happens, CPO will no longer be discussed as optics technology. It will be discussed as an infrastructure destiny.

About the Author

Lawrence Gasman is the President and Founder of CIR. He is a leading authority on deep-technology markets and the founder of CIR. For more than four decades, he has advised multinational corporations, venture investors, and strategic planners on where advanced technologies are heading — and how to act before markets shift. His expertise spans AI infrastructure, hyperscale data centers, optical networking, photonics, advanced processors, quantum technologies, and deep-tech commercialization. Lawrence is known for identifying inflection points early and delivering forecasts grounded in technical reality rather than hype.


Executives rely on his analysis to guide market entry, capital allocation, product strategy, and competitive positioning in complex, fast-moving sectors. A frequent speaker and sought-after commentator, Lawrence has been quoted and interviewed by major media outlets, including TIME and leading industry publications such as Lightwave. His perspectives are valued for their independence, clarity, and long-term view of technology cycles. He has also testified to Congress and advised the White House on telecommunications matters. He is the author of multiple books and publisher of the CIR DeepTech newsletter, read by professionals across technology, finance, and policy.

AI scale-up raises the stakes for transceivers and CPO adoption

LightCounting's Optics for AI report says the top cloud service providers will drive the adoption of CPO for scale-up connectivity.

BY SEAN BUCKLEY

Figure: Sales of optical transceivers and LPO/CPO for AI Clusters



Source: LightCounting

LightCounting's recent optical transceivers and CPO report looks at the correlation with capex.
www.lightcounting.com

As cloud operators' bandwidth needs for AI clusters continue to rise, so do sales of optical transceivers, linear pluggable optics (LPO), and co-packaged optics (CPO) for scale-out and scale-up networks.

LightCounting, in its recent "Optics for AI Clusters" report, estimates that sales of Ethernet optical transceivers (beginning with 100G) and CPO for scale-out and scale-up networks used in AI

clusters reached \$16.5 billion in 2025, and it will reach \$26 billion in 2026.

The research firm said its forecast “corresponds to a 60% growth rate for both 2025 and 2026.”

The data center interconnection (DCI) technology market weighs technology options

While DCI continues to gain momentum, there appears to be no clear winner in the optical technology race yet. A recent Lightwave survey revealed that the market is awaiting a dominant technology. The research found that nearly half of respondents (48%) say no clear technology leader has emerged, illustrating that the DCI industry continues to weigh performance, cost, and interoperability issues. While co-packaged optics (17%) and next-generation pluggable optics beyond ZR/ZR+ (14%) are drawing interest, adoption appears cautious. Even established options like 400ZR and 800ZR/800ZR+ each capture just 10%, signaling that organizations are weighing multiple paths forward rather than committing to a single optical strategy.

CPO to capex correlation

Despite the momentum around CPO and optical transceivers, LightCounting took a sanity check on its forecast to account for capex spending.

While Meta and Oracle plan to double their capex in 2026, other companies have not updated their spending plans.

As a result, LightCounting said its “current assumptions may be conservative” and that it is “very likely that growth in spending will moderate in 2027-2031.”

The research firm noted that optics sales to the top five cloud providers will account for 3.1% of their

capex in 2026 (up from 2.7% in 2025) and increase to 4.1% by 2031.

Cloud provider’s spending will focus on scaling up networking and on rapid growth in GPU bandwidth for both scale-out and scale-up connectivity.

“Adoption of CPO for scale-up connectivity can exceed our forecast and lead to a stronger growth for the market in 2028-2031,” LightCounting said.

Supply chain sanity

While [Meta](#) and [Microsoft](#) also reported their latest quarterly results and increased their 2026 capex projections, supply chain issues could be a factor.

Meta set its 2026 capex guidance at \$115–\$135 billion, a large increase from the \$72.2 billion spent in 2025.

Likewise, Microsoft saw significant increases in capex. Driven by AI infrastructure demand, Microsoft reported an 89% year-over-year increase in capex to \$29.88 billion for Q2 FY2026.

“The sky seems to be the limit for their spending, but supply chain shortages restore some sanity,” LightCounting said.

Overall sales of AI infrastructure products were constrained by supply chain shortages between 2023 and 2025.

Shortages of VCSELs and InP lasers limited shipments of optical transceivers. However, the VCSEL segment bounced back in mid-2024 as NVIDIA moved away from SR8 and adopted DR8 800G transceivers. Meanwhile, InP laser chip suppliers increased capacity in 2025.

"We expect the shortages to ease by mid-2026, which is great news for customers, but it is a double-edged sword for suppliers," LightCounting said. "They can ship more, but the demand tends to drop when shortages go away."

Looking forward, LightCounting does not foresee a downturn, but the market might see a flat quarter or two late this year. As a result, the research firm has "forecast moderation in the market growth for 2027-2031, assuming an eventual return to sanity among the top cloud companies."

CPO M&A accelerates

As momentum for CPO grew, vendors pursued key M&A deals to shore up their positions in the market segment over the past year. Marvell moved to [acquire Celestial AI](#), a next-generation CPO vendor, for \$3.25 billion, and later XConn to broaden its data center switching portfolio with PCIe and CXL products. However, Ciena [shelled out only \\$270 million for Nubis Communications](#), another co-packaged optics start-up, in September 2025. Celestial AI's technology is based on photonic interposers, which are similar to Lightmatter's approach. Both Celestial AI and Lightmatter raised several million dollars in funding in early 2025.

About the Author

Sean is responsible for establishing and executing the editorial strategy of Lightwave across its website, email newsletters, events, and other information products.

TERADYNE + QUANTIFI PHOTONICS

Powering the Next Wave of Optical Advancements for AI Data Centers

Enabling high-volume testing that drives efficiency in silicon photonics innovation

To learn more about photonics testing, visit teradyne.com/silicon-photonics-testing

TERADYNE

**QUANTIFI
PHOTONICS®**
A Teradyne Company

TERADYNE

Teradyne designs, develops, and manufactures automated test equipment and advanced robotics systems. Its test solutions for semiconductors and electronics products enable Teradyne's customers to consistently deliver on their quality standards. Its advanced robotics business includes collaborative robots and mobile robots that support manufacturing and warehouse operations for companies of all sizes. For more information, visit [teradyne.com](https://www.teradyne.com).

Website

[teradyne.com](https://www.teradyne.com)

Email

semitestcommunication@teradyne.com

Contact

<https://www.teradyne.com/company/contact-teradyne/>

QUANTIFI PHOTONICS

A Teradyne Company

Quantifi Photonics, a Teradyne company, provides advanced optical and electronic test solutions designed to enable scalable manufacturing of photonic integrated circuits (PICs), co packaged optics, and next generation optical interconnects. Their flexible, high density architecture supports all stages of the product life cycle - from early development and validation to high throughput characterization, manufacturing, and field deployment. Learn more at [quantifiphotonics.com](https://www.quantifiphotonics.com).

Website

<https://www.quantifiphotonics.com/>

Email

sales@quantifiphotonics.com

Contact

sales@quantifiphotonics.com

Resources

Teradyne Silicon Photonic Testing Applications

<https://www.teradyne.com/silicon-photonics-testing/>

Teradyne Photon 100

<https://www.teradyne.com/products/photon-100/>

Quantifi Photonics product brochure

<https://www.quantifiphotonics.com/download/5678/>

Quantifi Photonics White paper - PXI: The Way Forward for High-Volume Photonic IC Testing

<https://www.quantifiphotonics.com/download/7441/>