




PHOTONIC INTEGRATED CIRCUITS

ISSUE II 2026

 AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

PICMAGAZINE.NET

WHY SCALE-UP AI NETWORKS DEMAND SCALABLE **OPTICAL TEST**





www.EVGroup.com

INDUSTRY LEADING WAFER PROCESSING EQUIPMENT FOR PHOTONIC INTEGRATED CIRCUITS

■ Optical and Nanoimprint Lithography

■ Temporary and Permanent Wafer Bonding

■ NILPhotonics® Competence Center



EVG®7300
Multifunctional UV Nanoimprint Lithography System

GET IN TOUCH to discuss your process development and manufacturing needs

www.EVGroup.com


Reuniting materials for scalable photonics

➤ The long-standing ambition to integrate thin-film lithium niobate with ultra-low-loss photonic platforms has often been constrained less by physics than by fabrication reality. What is now becoming clear is that these constraints are not insurmountable, but they demand a change in approach. Rather than forcing incompatible materials into a single process flow, a more practical route lies in separating their fabrication and then reuniting them with precision.

This reflects a simple but important truth. The conditions required for optimal material performance are rarely aligned. Stoichiometric silicon nitride relies on high-temperature processing to achieve propagation losses below 0.2 dB/m, while lithium niobate degrades well below those temperatures. Combined with concerns around lithium contamination in CMOS environments, it is no surprise that monolithic integration has remained limited.

Wafer bonding offers a compelling way forward. By allowing each material to be processed under ideal conditions before combining them at low temperature, the platform avoids the compromises that have historically defined hybrid approaches. What stands out is not only the preservation of material quality, but the indication that this can be achieved at the wafer scale with yields approaching those required for manufacturing. This begins to shift the discussion from feasibility to deployability.

The decision to leave the lithium niobate non-patterned is equally significant. Etching has long introduced defects that degrade electro-optic performance. By instead confining the optical mode within the silicon nitride and coupling it into the lithium niobate only where needed, the design preserves



material integrity while maintaining efficient modulation. This hybrid-mode approach highlights a growing maturity in photonic design, where materials are used according to their strengths rather than forced into compromise.

These developments arrive at a time when photonic integrated circuits are becoming central to system performance. As data-intensive applications push interconnects to their limits, the challenge is no longer just speed, but scalability within existing manufacturing ecosystems. Compatibility with standard silicon platforms is therefore critical.

Challenges remain, particularly around long-term reliability and scaling to larger wafers. Yet the broader direction is clear. The future of photonics is unlikely to rely on a single material system. Instead, success will depend on how effectively different materials can be combined without sacrificing their advantages. In that context, modular integration may prove to be the most practical and powerful path forward.

Editor

Sarab Chopra sarab.chopra@angelbc.com

Contributing Technical Editor

Richard Stevenson richard.stevenson@angelbc.com +44 (0)1923 690215

Sales & Marketing Manager

Shehzad Munshi shehzad.munshi@angelbc.com +44 (0)1923 690215

Design & Production Manager

Mitch Gaynor mitch.gaynor@angelbc.com +44 (0)1923 690214

Graphic Design and Digital Media Assistant

Harvey Watkins harvey.watkins@angelbc.com

Publisher

Jackie Cannon jackie.cannon@angelbc.com +44 (0)1923 690205

Sales and Product Manager

James Cheriton james.cheriton@angelbc.com +44 (0)2476 718970

CEO Sukhi Bhadal sukhi.bhadal@angelbc.com +44 (0)2476 718970

CTO Scott Adams scott.adams@angelbc.com +44 (0)2476 718970

Published by

Angel Business Communications Ltd, 6 Bow Court, Fletchworth Gate,

Burnsall Road, Coventry CV5 6SP, UK. T: +44 (0)2476 718 970

E: info@angelbc.com W: picmagazine.net



PIC Magazine is published four times a year on a controlled circulation basis. Non-qualifying individuals can subscribe at: £105.00/e158 pa (UK & Europe), £138.00 pa (air mail), \$198 pa (USA). Cover price £4.50. All information herein is believed to be correct at time of going to press. The publisher does not accept responsibility for any errors and omissions. The views expressed in this publication are not necessarily those of the publisher. Every effort has been made to obtain copyright permission for the material contained in this publication. Angel Business Communications Ltd will be happy to acknowledge any copyright oversights in a subsequent issue of the publication. Angel Business Communications Ltd © Copyright 2026. All rights reserved. Contents may not be reproduced in whole or part without the written consent of the publishers. The paper used within this magazine is produced by chain of custody certified manufacturers, guaranteeing sustainable sourcing. US mailing information: Pic Magazine, ISSN 1096-598X, is published 4 times a year, March, May, August and December by Angel Business Communications Ltd, Unit 6, Bow Court, Fletchworth Gate, Burnsall Rd, Coventry CV5 6SP, UK. The 2026 US annual subscription price is \$198. Airfreight and mailing in the USA by agent named Air Business Ltd, c/o Worldnet Shipping Inc., 156-15, 146th Avenue, 2nd Floor, Jamaica, NY 11434, USA. Periodicals postage paid at Jamaica NY 11431. US Postmaster: Send address changes to Pic Magazine, Air Business Ltd, c/o Worldnet Shipping Inc., 156-15, 146th Avenue, 2nd Floor, Jamaica, NY 11434, USA. Printed by: The Manson Group. © Copyright 2026. We strive for accuracy in all we publish; readers and contributors are encouraged to contact us if they recognise an error or omission. Once a magazine edition is published [online, in print or both], we do not update previously published articles to align old company names, branding, marketing efforts, taglines, mission statements or other promotional verbiage, images, or logos to newly created or updated names, images, typographic renderings, logos (or similar) when such references/images were accurately stated, rendered or displayed at the time of the original publication. When companies change their names or the images/text used to represent the company, we invite organizations to provide Angel Business Communications with a news release detailing their new business objectives and/or other changes that could impact how customers/prospects might recognise the company, contact the organisation, or engage with them for future commercial enterprise.

Scale with Confidence Infrastructure Test for the AI Revolution

VIAVI MAP-300 is optimized for compact, high-throughput silicon photonics testing - from development through manufacturing - from wafer through packaging devices.

- Superior testing speed
- Lowest cost-of-ownership
- High-density hardware with multithreaded test execution
- Global footprint with expert local support
- Triggering with deep buffer capture



Learn More:


<https://www.viavisolutions.com/en-us/products/map-300-multiple-application-platform>

Contents

COVER FEATURE

Why scale-up AI networks need scalable optical test

Scale-up AI networks are pushing optics directly onto switches and XPU ASICs. To avoid a manufacturing bottleneck, the industry is shifting optical tests from low-parallelism, custom configurations to scalable, automated, IC-style methods.



20 How PI enables PICs: Interconnects from package - to planet-scale

As photonic interconnects move from network scale to chiplets and co-packaged optics, precision alignment is emerging as the critical bottleneck.



26 Shifting supply chains in the era of photonics

As AI drives the demand for interconnect bandwidth beyond the limits of copper, photonics is emerging as a critical part of the AI hardware stack.

32 Integrating thin film lithium niobate on low loss PIC platform

Ligentec and UCSD have demonstrated 110 GHz modulators by integrating thin-film lithium niobate with low-loss silicon nitride waveguides on standard silicon wafers

36 Electro-optic polymers move from promise to production

As AI data centres push networking to its limits, electro-optic polymers are moving beyond laboratory demonstrations.

40 Why optical scale-up must be multiplexed

Single-wavelength CPO works for scale-out, but future AI system scale-up demands a fundamentally different architecture built around multi-wavelength slow and wide networks.

NEWS

06 3D-printed plug simplifies fibre coupling

07 PICs drive next-gen imaging

08 Eindhoven starts 6-in InP waferfacility

09 External light sources emerge as key to scaling CPO

10 CPO advances AI infrastructure scaling

11 Photonics drives €370B spinout boom

12 Quantum boost for PIC market



3D-printed plug simplifies fibre coupling to photonic chips

Researchers at Heidelberg University have developed a 3D-printed plug interface that enables low-loss, scalable connections between optical fibres and photonic integrated circuits, supporting automated and cost-effective PIC manufacturing.

RESEARCHERS at Heidelberg University have developed a 3D-printed coupling approach that connects optical fibres to PICs using a plug-like interface, enabling low-loss and scalable optical connections for next-generation computing and communications systems.

Led by Prof. Wolfram Pernice at the Kirchhoff Institute for Physics, the research demonstrates a new method for linking fibre arrays to photonic chips without the complex active alignment processes traditionally required. The results were published in *Science Advances*.

Photonic integrated circuits use light instead of electrons to transmit

information, offering higher bandwidth and improved energy efficiency compared with conventional electronic systems.

However, coupling light efficiently between optical fibres and chips typically requires positioning accuracy within a few micrometres, making large-scale manufacturing challenging and costly.

The Heidelberg team addressed this issue by fabricating a 3D-printed optical interface directly onto the chip surface.


Acting as a plug, the structure aligns with standardised fibre arrays and redirects light using total internal reflection couplers designed for

telecommunications wavelengths between 1,500 and 1,600 nm.

The approach enables broadband transmission with minimal optical loss while relaxing alignment tolerances.

Using the concept, researchers successfully addressed a neuromorphic photonic processor featuring 17 optical ports, demonstrating the potential for automated and reproducible assembly of photonic systems.

The plug-and-play coupling method could support scalable manufacturing of PIC-based technologies and hybrid electronic-photonic systems, with potential applications spanning AI hardware, optical communications, and sensing platforms.



PICs drive next-gen imaging

PICs are enabling compact, energy-efficient imaging systems for applications ranging from LiDAR and augmented reality to medical diagnostics.

PHOTONIC INTEGRATED CIRCUITS are emerging as a key technology for next-generation imaging systems, offering new possibilities for miniaturisation, performance, and cost reduction.

By integrating multiple optical components onto a single chip, PICs enable precise control of light while significantly reducing system size, power consumption, and manufacturing complexity.

The technology is already influencing applications such as light detection and ranging (LiDAR), augmented reality (AR),

and optical coherence tomography (OCT).

In automotive systems, PIC-based LiDAR replaces bulky mechanical scanning components with on-chip optical phased arrays, enabling more compact and scalable solutions for advanced driver assistance systems.

PICs are also supporting advances in AR devices through integrated RGB laser modules capable of generating compact holographic projection systems.

In healthcare, integrated photonic chips are helping to reduce the size and cost of OCT diagnostic devices, opening the possibility for portable and point-of-care imaging solutions.

Despite these advances, engineering challenges remain, including optical losses, fabrication precision, and system integration.

Ongoing research aims to address these issues and further expand the role of PICs in future imaging technologies.


Quantum dot lasers tolerate extreme feedback

RESEARCHERS at King Abdullah University of Science and Technology (KAUST) and the University of California, Santa Barbara have demonstrated that quantum dot (QD) lasers can maintain telecom-grade performance even under extreme optical feedback, eliminating the need for isolators in PICs.

Using a custom experimental platform that delivers continuous feedback up to 0 dB, the team directly measured the coherence collapse threshold of QD Fabry–Perot lasers at -6.7 dB (21.4% return).

Despite operating near this limit, the lasers supported 10 Gbps external modulation, stable operation between 15 – 45 °C, and over 100 hours of continuous feedback operation with excellent device-to-device reproducibility.

The findings, published in *Light: Science & Applications*, were reinforced by theoretical modelling based on the Lang–Kobayashi analysis, which



confirmed that scale PIC cavities further enhance QD feedback tolerance.

Compared to quantum well, quantum wire, and VCSEL platforms, the QD Fabry–Perot devices showed superior stability under feedback.

By removing the need for optical isolators, this work simplifies PIC packaging, improves manufacturability, and reduces system costs.

It also provides clear design rules for integrating QD lasers into next-generation PIC systems for communications, sensing, LiDAR, and large-scale photonic platforms.

This research underscores the potential of QD lasers to enable energy-efficient, reliable, and cost-effective PICs, accelerating their adoption in high-performance optical technologies.

Eindhoven starts 6-in InP wafer facility

New pilot line aims to scale indium phosphide photonic chip manufacturing for AI, 6G and datacenter applications.

CONSTRUCTION has begun on a new industrial facility in Eindhoven to produce photonic chips on 6-inch indium phosphide wafers, marking a major step toward scaling European photonic integrated circuit manufacturing.

The project is led by the Netherlands Organisation for Applied Scientific Research (TNO) in collaboration with High Tech Campus Eindhoven and forms part of the broader PIXEurope consortium. The initiative is supported by a €150 million investment under the European Chips Act.

Partners in the project include Eindhoven University of Technology, PhotonDelta and SMART Photonics, intending to link research and development with scalable industrial production of photonic integrated circuits.

The new facility is expected to be the world's first dedicated 6-inch indium phosphide wafer production line, enabling higher-volume manufacturing of photonic chips used in applications such as AI data centres, 6G communications, supercomputing and medical technologies.

Plans for the pilot line were first announced in 2024, with detailed development plans revealed in 2025.

The project aims to accelerate the transition from research prototypes to commercially deployable photonic technologies in Europe.

The move also reflects growing global competition in InP manufacturing. Coherent has also announced efforts to establish 6-inch indium phosphide wafer production at its facilities in Texas and Sweden.

NVIDIA signals silicon photonics shift at GTC

AT ITS GTC keynote this week, NVIDIA CEO Jensen Huang underscored the growing importance of silicon photonics in scaling next-generation AI infrastructure, marking what many see as a pivotal shift away from traditional copper interconnects.

The company's latest updates to its Spectrum-X platform, including the introduction of co-packaged optics (CPO), signal a deeper integration of photonics and electronics to address mounting bandwidth, power and latency challenges in large-scale AI systems.

According to Graham Reed, Director of CORNERSTONE and a leading figure in silicon photonics, the announcement represents a significant milestone for the industry.

"NVIDIA's update on the Spectrum-X switch with co-packaged optics is an important moment, confirming that silicon photonics is central to next-generation AI infrastructure," he said.

Reed highlighted that while the advantages of photonics for scale-out architectures are well established, NVIDIA's commitment to the COUPE process technology, developed in collaboration with TSMC, is particularly notable.

The approach integrates photonic and electronic components directly on-chip, helping to overcome the physical limitations of conventional interconnect technologies.

He also pointed to a broader strategic shift within NVIDIA. Despite a long-standing reliance on copper-based interconnects for scale-up systems, the company is now placing photonics at the core of its future platforms, including Vera Rubin Ultra.

This transition is expected to support increasingly complex configurations, such as NVL576 and future architectures like Kyber NVL1152.

The move reflects a growing consensus across the industry that copper

interconnects alone are no longer sufficient to meet the performance and efficiency demands of advanced AI workloads.

Beyond its technical implications, the announcement also carries strategic significance for regional ecosystems.

Reed noted that the rapid pace of development in photonics technologies underscores the need for countries such as the UK to act quickly to maintain competitiveness.

"The UK must act immediately to leverage its world-leading expertise and secure a role in this emerging photonics ecosystem," he said.

As hyperscale AI infrastructure continues to evolve, NVIDIA's endorsement of silicon photonics is likely to accelerate adoption across the sector, reinforcing its role as a foundational technology for future high-performance computing systems.

External light sources emerge as key to scaling CPO

New architectures that separate lasers from high-heat processors are gaining traction as AI datacentres push the limits of co-packaged optics.

EXTERNAL LIGHT SOURCES (ELS) are emerging as a critical enabler for scaling co-packaged optics (CPO) in next-generation AI datacentres, as the industry looks to overcome mounting thermal, reliability and integration challenges.

As AI workloads continue to expand, traditional electrical interconnects are increasingly unable to meet the bandwidth and energy efficiency demands of modern data centre architectures.

CPO has been widely identified as a solution, integrating optical components directly with GPUs and switch packages to deliver high-speed, low-power connectivity. However, the approach introduces new challenges, particularly around thermal management and laser stability.

ELS architectures address this issue by relocating temperature-sensitive lasers away from high-power processing units.

This separation improves wavelength stability, enhances system reliability and simplifies serviceability, making ELS an increasingly attractive approach for both scale-out and scale-up optical systems.

Recent industry activity highlights growing momentum behind this shift. Siverts Semiconductors, O-Net Technologies and Enablence Technologies have announced a collaboration to develop an advanced ELS module designed to support CPO deployment in AI and high-performance computing environments.

The solution combines laser arrays with photonic distribution technologies to

enable scalable, high-density optical interconnects.

The move reflects a broader industry trend, as data centre operators and technology providers seek practical pathways to scale optical connectivity without compromising performance or efficiency.

With market forecasts pointing to rapid growth in CPO adoption over the coming decade, ELS solutions are expected to play a central role in enabling reliable and manufacturable architectures.

As AI infrastructure continues to evolve, the shift towards ELS-based designs signals a growing recognition that optical innovation at the system level will be essential to sustaining performance gains in increasingly complex computing environments.



CPO advances AI infrastructure scaling

Integration of photonics with computing is set to overcome bandwidth and energy limits in next-generation data centres.

THE RAPID growth of artificial intelligence is accelerating the shift toward large-scale, distributed compute architectures, driving demand for new interconnect technologies such as co-packaged optics (CPO).

CPO integrates optical engines closer to switch and processor chips, replacing traditional front-panel transceivers to reduce signal loss, improve bandwidth density, and lower power consumption. This approach is emerging as a key enabler for scaling AI infrastructure, particularly as copper interconnects reach their physical limits.

Recent advances in electronic and photonic integration have already delivered significant gains, including a 64-fold increase in bandwidth density and a fivefold improvement in energy efficiency over the past five years. However, challenges remain in bringing

CPO to high-volume manufacturing.

A key issue is the mismatch between the bandwidth density of AI chipllets and current optical interconnects, highlighting the need for more compact and scalable photonic integrated circuits. Fiber array design, wavelength stability, and thermal management also present ongoing technical hurdles.

Manufacturing complexity is another barrier, with CPO systems requiring the integration of lasers, photonic circuits, and fiber arrays across multiple process nodes.

Active alignment techniques, while necessary for performance, can limit throughput and increase costs. In addition, conventional permanent bonding methods reduce repairability, prompting industry efforts to develop detachable fiber array solutions.

Testing and yield optimisation are also critical concerns. The adoption of “known-good” optical engines, verified through wafer-level testing prior to final assembly, is seen as essential for improving manufacturing efficiency.

However, not all integration approaches are compatible with current testing methods, particularly in advanced 3D configurations.

Despite these challenges, co-packaged optics is widely viewed as a cornerstone technology for future AI and hyperscale data centre systems.

Continued progress will depend on close collaboration across the semiconductor and photonics ecosystem, spanning design, packaging, and system integration.



Photonics drives €370B spinout boom


Photonics is emerging as a core driver of Europe's rapidly growing deep-tech spinout ecosystem.

A MAJOR new study has revealed that Europe's university spinouts have grown into a technology ecosystem worth €370 billion, with photonics and quantum science playing a central enabling role across key industries.

The "European Spinouts Report 2025", produced by Dealroom.co, analysed more than 17,000 university spinouts created since 1990, including 7,300 deep-tech and life sciences companies.

These firms now employ over 160,000 people and collectively represent one of Europe's most dynamic sources of science-driven innovation.

Photonics is highlighted as a foundational technology underpinning advances in sectors ranging from



semiconductors and robotics to climate technologies and advanced sensing systems. The science of harnessing light enables critical applications in communications, imaging, manufacturing and diagnostics. Dr Lutz Aschke, President of Photonics21, said Europe's strength in

optical science is translating into globally competitive companies, adding that ensuring these spinouts scale within Europe will be key to long-term industrial success.

The report also points to rapid acceleration in recent years, with nearly 40% of total spinout value created by companies founded since 2015.

Since 2019, spinouts have accounted for around 40% of all new deep-tech and life sciences startups, reflecting growing investor interest in science-led innovation.

As global competition intensifies in areas such as semiconductors and quantum technologies, the findings highlight the strategic importance of converting Europe's research excellence into scalable industrial capability.

NIST advances robust photonic packaging

RESEARCHERS at the National Institute of Standards and Technology (NIST) have developed a novel packaging method for PICs that allows them to withstand extreme environments, including high radiation, ultrahigh vacuum and cryogenic temperatures.

The breakthrough centres on improving one of the most critical aspects of photonic packaging, the connection between optical fibres and chips.

Conventional polymer-based adhesives used in fibre-to-chip bonding tend to degrade under harsh conditions, leading to misalignment and device failure.

To overcome this, the NIST team adopted hydroxide catalysis bonding (HCB), a technique originally developed by NASA for precision optical systems.

The method creates an inorganic, glass-like bond at the molecular level, enabling highly stable and precise alignment without relying on traditional adhesives.

Testing showed that PICs packaged using HCB maintained performance after exposure to rapid temperature cycling, intense ionising radiation and high vacuum conditions.

The bonding approach also demonstrated strong mechanical

stability at elevated temperatures, suggesting significantly improved resilience compared to conventional packaging methods.

The advance could expand the use of photonic technologies in demanding applications such as quantum computing, space systems, nuclear instrumentation and industrial sensing, areas where conventional electronic and photonic packaging has struggled to perform reliably.

While the current bonding process requires several days to complete, researchers say further engineering development could make it viable for large-scale manufacturing.

Quantum boost for PIC market

Quantum technologies are accelerating demand for next-generation photonic integrated circuits and new material platforms.

THE PHOTONIC INTEGRATED CIRCUIT (PIC) industry is entering a new phase of innovation as quantum technologies accelerate demand for advanced optical chip platforms, according to insights shared by IDTechEx.

PICs, miniaturised optical systems fabricated on semiconductor wafers, are already critical to high-speed data communications, particularly in AI data centres.

Their role is now expanding as quantum computing, sensing, and communications move from laboratory research to commercial deployment.

Quantum technologies rely heavily on precise control and manipulation of light. Traditional experimental setups, often involving bulky optical tables and complex laser arrangements, are impractical for real-world applications.

PICs offer a scalable alternative, enabling these optical systems to

be integrated into compact, robust, and manufacturable chips. This shift is especially significant for quantum computing architectures.

Companies such as IonQ and Quantinuum use photonics to control trapped ions, while PsiQuantum and Quandela are developing systems where photons act as qubits. Scaling these technologies will depend heavily on advances in PIC design and manufacturing.

A key challenge identified by IDTechEx is the limitation of traditional silicon-based photonics.

While silicon and silica dominate current PIC production due to mature fabrication ecosystems, they are not always suitable for quantum applications, particularly those requiring operation in visible wavelengths or ultra-low noise environments. As a result, the industry is exploring alternative materials.

Platforms such as silicon nitride, thin film lithium niobate, and barium titanate are gaining attention for their optical performance and electro-optic properties. However, these materials face barriers including higher costs, smaller wafer sizes, and limited foundry availability.

Despite these challenges, the commercial outlook is strong. IDTechEx forecasts that PICs tailored for quantum technologies could represent a US\$12.6 billion market opportunity by 2046, driven by applications across quantum computing, secure communications, and advanced sensing.

Looking ahead, the convergence of photonics and quantum technologies is expected to reshape the PIC landscape, pushing innovation beyond silicon and opening new opportunities for materials, fabrication, and system integration.



Luceda Photonics Design Platform

Luceda IPKISS

First-time-right
photonic IC design software

Automate and integrate all aspects of
your photonic design flow in one tool,
using one standard language.



Luceda Design Kits

Gain access to a wide range of PDKs, ADKs, and TDKs.
Design and tape out to your foundry of choice.

OFC 2026

Visit us at our booth
to learn more!


Booth
#965




New Product Luceda DRC

Identify and resolve design issues
instantly within your design environment
with a single click, and proceed to
successful tape-outs with confidence.


CONTACT US TO
TRY IT TODAY!



Component / Circuit
Layout



Design Rule Checking
in Luceda DRC



DRC-Clean
Layout

Why scale-up AI networks demand scalable optical test

Scale-up AI networks are pushing optics directly onto switches and AI accelerators. To avoid a manufacturing bottleneck, the industry is shifting optical test from low-parallelism, custom configurations to scalable, automated, IC-style methods.

BY ANDREW YICK, SENIOR DIRECTOR, PRODUCT & TEST ENGINEERING, MARVELL

THE DOMINANT CHALLENGE in modern AI infrastructure is not just the performance of a single accelerator but scaling up to thousands of accelerators (XPUs) in a cluster. Training and inference workloads now depend on an interconnect that can stitch these accelerators into a high-bandwidth, low-latency system, where performance is governed as much by the network as by the compute itself.


As these systems scale, physics asserts itself. Electrical links over copper hit a practical ceiling as routing density and channel loss collide, turning the loss bandwidth product into an impassable constraint. The choice is binary: either move electrical-to-optical conversion closer to the Application-Specific Integrated Circuit (ASIC) or surrender the link budget. Thus, to bypass this electrical wall, optics must migrate from the board edge and onto the ASIC package.

This progression is not an architectural preference; it is a physical necessity. Every serious scale-up roadmap eventually converges on the same outcome: optical interfaces migrating from front-panel modules (including optical-engine-based pluggables) to Near-Packaged Optics (NPO), and ultimately to Co-Packaged Optics (CPO) integrated alongside the compute die. CPO is not the goal in itself, but the inevitable consequence of interconnect bandwidth scaling.

The bottleneck moves to test

Silicon photonics platforms today deliver extraordinary functionality and bandwidth density, but the industry is hitting a manufacturing wall. Historically, optical testing has evolved as a specialized, low-volume practice. Test cells are often custom-built, alignment-heavy, and single-site. Optical instruments typically reside outside the Automated Test Equipment (ATE) ecosystem in rack-and-stack configurations, and throughput is measured in minutes rather than seconds. While these approaches suffice for labs and pilot lines, they break down at high channel counts and production volumes.

The issue is not the measurement rigor, it's manufacturing efficiency at scale. As optics move inward toward the ASIC, we still need the same level of test coverage, but delivered with IC-style parallelism, repeatability, and automation.



The practical advantage of a mainstream ATE test head is its high-density instrumentation and native multi-site architecture, allowing one test program to run across multiple devices in a single touchdown. This efficiency stems from the ATE's ability to synchronize complex, high-parallelism tasks such as coordinating laser sweeps and power meter captures across all sites simultaneously. To better realize this system-level coordination at optical wafer test, next-generation opto-electrical probe cards integrate the optical probe head and on-card fine alignment directly into the probe card. This integrated architecture allows optical wafer testing to inherit the semiconductor industry's high-volume manufacturing model.

Figure 1 illustrates a fully integrated optical test cell from Teradyne, built around the UltraFLEXplus platform with Photon 100 and a TEL Precio XL production wafer prober. An integrated opto-electrical probe card assembly, auto-loaded by the prober, brings the optical probe head to the wafer while maintaining a standard ATE test flow. Electrical instrumentation resides natively inside the tester, while the optical sources and measurement resources are consolidated in an external rack that is fully interfaced and safety interlocked to the test head. Optical connections are then blind-mated through a standard load board at the prober docking interface.

Figure 2 highlights a Technoprobe integrated opto-electrical probe card architecture. First-generation integration using static optics with expanded-mode interfaces is already deployed in volume production today, with solutions available from vendors such as Jenoptik. The figure shows the next-generation evolution: cards with integrated piezos for sub-micron fine alignment.

The solution also integrates a native on-card controller and a local high-voltage step-up to drive the piezos, avoiding the need for high-voltage sources from the tester. By combining the prober's coarse positioning with local high-speed alignment at the device interface, this Outsourced Semiconductor Assembly and Test (OSAT)-friendly architecture eliminates bulky optomechanics, effectively closing the gap between lab-grade precision and factory-friendly volume.

Shift left or absorb the cost

As optical functions are integrated into multi-chip modules (MCM), the full-rate optical behavior may only be uncovered once the CPO-Integrated XPU package is fully assembled – after substantially more silicon value has already been added and when rework may be limited or unavailable.


Optical devices are no longer pluggable modules; they are built into engines and chiplets, co-packaged alongside ASICs, memory, and substrates. At this level of integration, a single failing optical channel can force the scrap of an entire high-value assembly.



➤ Figure 1: Teradyne UltraFLEXplus with Photon 100 single-sided optical wafer probe cell (Courtesy of Teradyne)

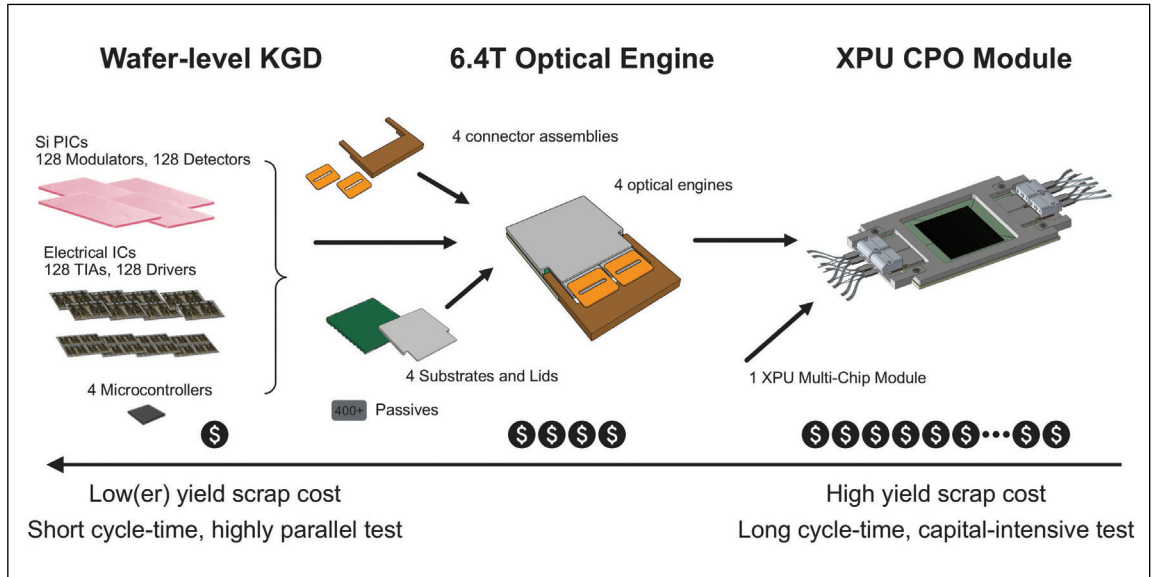
Late discovery is therefore economically untenable. The only viable strategy is to shift meaningful optical test earlier in the manufacturing flow, specifically to the wafer level, but only if those early measurements are predictive and correlated to full-rate performance in the final assembly. This logic is familiar to semiconductor manufacturing, but its implications are more severe for integrated optics because the financial penalty for a late-stage failure is so high. Establishing Known-Good-X (KGX) is no longer optional: where “X” is the unit being certified as known good – die, optical engine, or chiplet – before it is integrated into a higher-value assembly.

Figure 3 illustrates the direct impact of test timing on manufacturing economics. As devices move through the assembly process, defect discovery becomes disproportionately expensive, making “shift-left” testing a fundamental requirement for scalable manufacturing economics. This strategy prioritizes short cycle-time highly parallel tests at the wafer level, in contrast to the long cycle-time and capital-intensive tests required once optics are integrated into a full CPO-Integrated XPU module.



➤ Figure 2: Technoprobe integrated opto-electrical probe card for high-volume production (Courtesy of Technoprobe)

➤ Figure 3: “Shift-left” optical test improves manufacturing economics by reducing scrap exposure and maximizing throughput



A unified optical test insertion framework

Scaling optical test requires a structured framework with explicit test insertions, providing a common language across design, test, and manufacturing. Figure 4 illustrates this hierarchy:


- Insertion 1 (Photonic IC Wafer):** The PIC wafer is screened using electrical and optical parametric measurements to establish Known Good Die (KGD) at the earliest possible stage. This initial gate employs high-parallelism optical scans to generate a parametric map of the wafer, which may include critical spectral characterization, modulator efficiency, and polarization-dependent loss (PDL) measurements. By identifying marginal components early, the exponential cost of integrating defective silicon into high-value downstream assemblies is avoided.
- Insertion 2 (PIC + EIC):** This occurs once the photonic and electronic ICs are bonded, representing the first stage of multi-die integration. Electro-optic transfer characteristics may be measured at speed and correlated back to initial wafer-level results (e.g., modulation response, eye quality, and other high-speed

figures of merit). This correlation ensures that the bonding or die-attach process has not degraded the PIC’s performance, while also assessing the health of the inter-die connectivity.

- Insertion 3 (Optical Engine):** Engines or chiplets undergo final calibration and functional testing after fiber attach or connectorization. These tests establish KGx status for the assembly, serving as the final gate to prevent defective units from propagating into multi-chip modules. Validation at this stage typically expands to include link-budget verification and thermal characterization to ensure that fiber-to-chip coupling remains stable across the full operating temperature range.

While the final test insertions, including the CPO XPU and system-level assemblies, are likewise critical, their primary role shifts to validating overall functionality. The final test insertions are not where component-level yield problems should be discovered. By shifting the burden of discovery to early PIC-centric insertions, this structured approach ensures that yield is managed at the most cost-

➤ Figure 4: Optical test insertion hierarchy for scaling photonics test from wafer screening to system-level



effective stages of the production flow.

Industrializing the optical test methodology

The transition from treating optics as a specialty test to an IC-style methodology is already underway. This philosophy now extends beyond the wafer to the chiplet level (Insertion 3), where the challenge shifts to verifying integrated assemblies.

Figure 5 illustrates a chiplet-level optical test cell built around an Advantest V93000 EXA and an MPI DTS650-DI chip-level prober. In this configuration, the optical alignment function is integrated into the prober itself, enabling bottom-side or edge coupling of the optical chiplet within a standard OSAT-style manufacturing footprint. This removes the external opto-mechanics that typically obstruct direct-dock operation and complicate automation. Crucially, the tester remains on a standard ATE platform with a conventional load board and docking interface. Electrical resources reside natively in the test head, while optical resources reside in a dedicated test cabinet and are routed into the cell under full tester control. The result is a manufacturing-ready system



in which optical chiplets can be handled, aligned, and verified alongside digital content without resorting to custom bench setups.

These proof points change the conversation. The question is no longer whether a scalable optical test is possible, but how quickly the ecosystem can standardize around it. To reach a true global scale, several pieces must still fall into place:

- **Native ATE Instrumentation:** We require native optical instrumentation that resides directly inside the ATE test head, reducing the reliance on external optical racks either adjacent to or bolted onto the test cell.
- **Defined Standards:** The ecosystem requires industry-wide agreement on standards for scale-up optical networking—including wavelengths, fiber counts, modulation formats, and target power levels as well as unified calibration and correlation methods.
- **Physical Interfaces:** Standardized physical interfaces, such as common PIC port definition, detachable connector solutions, and blind-mate docking, are essential for manufacturing adoption.

➤ Figure 5: Chiplet-level optical test cell combining an Advantest V93000 EXA direct-dock test head with an MPI DTS650-DI prober with integrated optical alignment (Courtesy of Advantest and MPI)



BOOK YOUR REPRINT TODAY!

A reprint from **PIC Magazine** amplifies your editorial exposure and provides authoritative third-party validation. It's a powerful tool for **sales and marketing**, helping you showcase innovation with independent credibility. Use it across **events, pitches, newsletters, and social channels** to reinforce trust, extend visibility, and strengthen your position in the PIC industry.



Contact: Shehzad Munshi
shehzad.munshi@angelbc.com



Accelerating the adoption of these standards will determine how quickly “optics as an IC” moves from a custom, low-volume specialty test to a mainstream manufacturing reality.

Design-for-test: the shared responsibility

As optics matures into an IC-style manufacturing problem, design and test must converge earlier in the product cycle. Scalable optical test cannot be retrofitted as a final stage of development; it must be explicitly enabled by design. Optical Design-for-Test (DfT) is therefore no longer a luxury, but a requirement.

Specific architectural features determine whether a device can be economically tested at scale:

- **Access Points:** Specific points for parametric measurement must be integrated into the design. These access points may include dedicated optical DfT paths and integrated monitor diodes that allow for in situ structural checks.
- **Correlation Support:** Built-in support is needed to correlate data across various test insertions. By burning unique IDs and calibration trims during wafer-level screening, a digital imprint is created that ensures parametric traceability in downstream test insertions.
- **Parallel Test Strategies:** Parallel testing is essential for high-volume throughput—maximizing tester-resource utilization and enabling multi-site operation. Supporting simultaneous tuning and calibration of all transmitter/receiver channels, without compromising measurement integrity, is key to meeting AI-scale takt times. Built-In Self-Test (BIST) further enables this parallelism by generating at-speed stimuli and leveraging loopbacks to screen channels in situ, reducing dependence on external optical instrumentation.

Without these front-end considerations, even the most advanced ATE infrastructure will struggle

For the test community, the implication is clear: the industry must move away from the legacy optical module era practices in favor of integrated, automated solutions that leverage existing semiconductor infrastructure.

to maintain yield, test coverage, and cycle-time targets.

For the test community, the implication is clear: the industry must move away from the legacy optical module era practices in favor of integrated, automated solutions that leverage existing semiconductor infrastructure.


Optics as a first-class manufacturing domain

Scale-up AI networks depend on optics, and optics, in turn, now depends on test. What began as a specialized, lab-centric practice must now mature into a first-class manufacturing capability, aligned with the rigorous methods that have sustained semiconductor scaling for decades.

This transition requires deep collaboration across design, test, instrumentation, and manufacturing, and it is already well underway across the ecosystem. As scale-up AI clusters advance, optics must follow the same DfT, shift-left, and ATE-enabled manufacturability playbook that has long governed high-volume electronics. The mantra for the upcoming years is simple: If it cannot be tested like an IC, it will not scale like one.

FURTHER READING

- Enabling Efficient PIC Testing for Mass Production – Chip Scale Review, Vol. 29, No. 3 (June 2025). This cover feature reports validation data for integrated probe-card approaches for scalable, high-volume testing of photonic integrated circuits.
- Silicon Photonics Comes of Age: Scaling Test for the AI Era – Semicon Taiwan 2025 Keynote, Taipei, TW (September 11, 2025). A keynote detailing the evolution of silicon photonics testing and specific test requirements for high-volume AI infrastructure.
- Marvell Unveils AI-Focused Data Center Connectivity Innovations Ahead of OFC 2026 – Marvell Press Release (March 12, 2026). Details the Marvell Photonic Fabric™ technology platform designed to meet the reach, bandwidth, and latency demands of next-generation AI clusters.
- Building Scalable Testing & Custom Compute for the AI Era – Counterpoint Research YouTube Interview (November 27, 2025). A discussion on the evolving semiconductor testing landscape, specifically regarding challenges in Co-Packaged Optics (CPO) and innovations in optical testing.
- IEEE Heterogeneous Integration Roadmap (HIR) 2024 Edition, Chapter 17: Test Technology. A widely referenced industry roadmap outlining collaborative requirements for testing complex heterogeneous systems, including photonic chiplets
- Custom Silicon: A Sea Change for Semiconductors – Marvell Blog (December 16, 2025). An exploration of how custom accelerators and Co-Packaged Optics (CPO) are becoming essential for scaling compute infrastructure efficiently.




PROFESSIONAL SIMULATION TOOLS & EXPERT DESIGN SERVICES



Integrated Devices, Waveguides & Fibers

VPIdeviceDesigner™




Photonic Integrated Circuits

VPIcomponentMaker™ Photonic Circuits



Fiber Amplifiers & Lasers

VPIcomponentMaker™ Fiber Optics



Transmission Systems

VPItransmissionMaker™ Optical Systems



Lab-proven DSP Algorithms

VPItoolkit™ DSP Library



Quantum Key Distribution

VPItoolkit™ QKD

We empower you to
define the cutting edge.



Request an
evaluation



How PI enables PICs: Interconnects from package - to planet-scale

As photonic interconnects move from network scale to chiplets and co-packaged optics, precision alignment is emerging as the critical bottleneck. Physik Instrumente (PI) is addressing this challenge with high-speed, multichannel alignment systems that enable scalable, high-yield PIC manufacturing.

BY DAVID FORER – DIRECTOR MARKET STRATEGIES – SEMICONDUCTOR, PI PHYSIK INSTRUMENTE, SCOTT C. JORDAN – PRINCIPAL, AIKIDOVATION LLC. (FORMER HEAD OF PHOTONICS; SR. DIRECTOR, NANOAUTOMATION; PI FELLOW AT PI USA / PI (PHYSIK INSTRUMENTE))

PHOTONIC INTEGRATED CIRCUITS (PICs) integrate optical functions (generation, modulation, routing,



➤ PI's F-712. MA2 Fully Automated Alignment System

and, increasingly, computation) onto chips. As bandwidth demand grows and power budgets shrink, the industry is steadily pushing photonic interconnects closer to the silicon: from the global network backbone all the way down to the package, to the dies themselves, and between chiplets in a package. That shift is not just a bullet point on a roadmap. It changes, fundamentally, how semiconductors get tested and built. In most PIC workflows today, the bottleneck is no longer the optical design; it's the packaging. Physik Instrumente (PI) addresses this head-on with precision motion subsystems and alignment firmware that turn what used to be slow, serial active alignment steps into fast, repeatable, factory-grade processes.

From transoceanic links to chip-to-chip optics: a compressed history

Photonics started at the largest scales: transoceanic and transcontinental fiber networks that made long-distance communication viable. From there,

the frontier moved steadily inward through regional and metro networks, then fiber-to-the-home, each step shrinking the span while multiplying the endpoints. The next jump brought photonics inside the data center, replacing copper in building-to-building links and rack-to-rack connections. Now the industry is pushing further still: co-packaged optics (CPO) and related architectures aim to place optical components on the same substrate as switching or compute silicon, cutting electrical trace lengths and the power and signal losses that come with them. The appeal is real: better power-per-bit, higher capacity and lower latency by moving optical conversion right up next to the ASIC. [R5] [R6] [R7]

AI is accelerating this. Training and inference clusters generate massive east-west traffic, while power and cooling have become first-order design constraints rather than afterthoughts. Latency, once a secondary concern, is now a competitive differentiator, especially in disaggregated architectures and edge AI deployments. Meanwhile, the ambitions for optical interconnects keep expanding outward too: low-Earth-orbit networking and ‘data centers in space’ concepts are gaining traction as routes to unique power and cooling environments. Whether those particular visions pan out or not, the broader signal is unmistakable: PIC-enabled interconnects are evolving into a cross-scale platform, from chiplets in a package all the way to global links. The common thread running through all of it is manufacturable, high-yield optical coupling.

PIC interconnects inside the package: chiplets, CPO, and the war on latency

Inside advanced packages, PICs appear in two main configurations. First, optical engines sit adjacent to compute or switch on silicon (the CPO model). Second, as photonic bridges between chiplets (think memory ↔ accelerator links) where electrical channel limits and power density make optical alternatives attractive. Both approaches rely on dense arrays of optical I/O, often tens of channels at once, which immediately sets up the assembly challenge: you need to align arrays of fibers, lenslets, interposers, or waveguide couplers with sub-micron accuracy without sacrificing production throughput. Every review of photonic-chip packaging says the same thing: optical I/O coupling and interconnection are the central scaling challenges. [R8]

From a manufacturing perspective, the uncomfortable truth is that packaging and testing often dominate both time and cost, especially when legacy active alignment is used as a serial, single-channel process. As channel counts climb, the penalty compounds: more elements to align, more degrees of freedom to optimize, and more repetitions across wafer test, module assembly, and final package verification. A small inefficiency in alignment becomes a hard capacity and yield ceiling

at volume. That’s why ‘alignment economics’, not just alignment accuracy, has become a strategic lever for anyone running a serious PIC program.


From a technical perspective, our use of the word “accuracy” regarding alignment performance spotlights the additional need for an evolved vocabulary for this field. Ordinarily, accuracy regards the performance of a positioner versus a perfect ruler. This goes to manufacturing quality, encoder principle, and other mechanical subtleties, and can be seriously impacted by design choices like fixturing offsets and environmental variables such as temperature. In the context of alignment, however, accuracy reflects the ability of the system to determine the mutual position and orientation that achieves best coupling according to the figure of merit (usually, but not always, optical power). This goes more to the alignment system’s algorithms than its physical details. In fact, highly accurate alignment can be achieved with open-loop mechanisms lacking any motion metrology at all, as was the case for the earliest active alignment engines in the late 1980s.

The recurring challenges - and how PI addresses them

Yield: coupling accuracy, drift control, and device safety


Yield in the PIC assembly is tightly coupled to positioning accuracy. Coupling losses can escalate significantly with even small offsets, particularly for edge coupling and other tight mode-field interfaces where the tolerances leave little room for error. PI’s active alignment approach treats coupling as a closed-loop optimization problem: measure optical power (or another figure of merit) in parallel with multi-DOF motion, converge on the optimum, then hold it. This is spelled out clearly in PI’s own active alignment publications and documentation, which identify silicon photonics testing and packaging as key drivers for both nanoscale accuracy and economic efficiency. [R1]

PI’s Fast Multichannel Photonics Alignment (FMPA) systems bring precision mechanics together with intelligent, firmware-level alignment commands running in dedicated digital controllers (the E-713 platform being the flagship example) to enable fast, simultaneous alignment and tracking of multichannel couplings across multiple degrees of freedom (DOFs) and multiple inputs and outputs. [R2] [R9]



➤ F-713.HAX
Compact,
High-Speed
6-DoF
Photonics
Alignment
System

➤ PILightning®
Firmware
feature for
alignment
systems -
Optimization
of first light
detection



There's also a practical device safety benefit to this architecture. By integrating motion, sensing, and alignment logic inside the controller itself, systems can execute scans efficiently, without the stop-start latency of PC-mediated control loops, while operating within well-defined constraints. When you're dealing with expensive wafers or fragile fiber arrays, that predictability matters as much as raw speed.

Productivity: throughput, reproducibility, and new approaches



Throughput matters enormously in PIC manufacturing. Alignment steps repeat constantly, and as channel counts keep rising, any slowness in that step compounds. PI reports that FMPA can cut fiber-array alignment time in silicon photonics packaging by roughly a factor of 100, down to about one second, compared to conventional approaches. [R3]

But this isn't purely a speed story; it's a reproducibility story. When alignment is automated and embedded in the motion control stack, the result depends less on operator skill and more on traceable control parameters. That creates a pathway to process capability metrics, recipe management, and scalable QA: exactly the kind of rigor the semiconductor industry already expects in front-end manufacturing.

Worked examples: FMPA, PILightning, and PINovAlign

Example A: Multichannel optimization: Once multichannel PICs became the norm, aligning fibers became an obvious bottleneck; you could see it in cycle times. There are several reasons: foremost is simple geometric dependencies between DOFs, where adjustment of one axis impacts others.

Formerly, a recursive, looping, and serial approach was required to bring everything into consensus alignment. This was time-consuming. PI's FMPA was designed specifically to address this, enabling parallel optimization across channels and degrees of freedom simultaneously, all implemented at the firmware level in PI's controllers. [R2] [R10]

Example B: Time to 'first light': Getting to that initial coupling point is often the most painful part of active alignment, particularly when tolerances are tight and the search space is large. PI's PILightning feature extends the FMPA algorithm package specifically to accelerate this initial acquisition phase and improve overall alignment efficiency on compatible controller platforms. [R4]

Example C: Compact, cost-optimized alignment: Not every station needs a full hexapod, and PI recognizes that. The PINovAlign (F-141) is aimed at applications that need fast, accurate multi-axis motion in a compact footprint: a cost-optimized system with fast alignment algorithms baked in, purpose-built for PIC test and assembly environments [R11] [R12]

Trends shaping PIC assembly and why general solutions still work

PIC device diversity is relentless: grating couplers, edge couplers, integrated lasers, externally coupled sources, photodiodes, phased arrays, interposers, new fiber types, and the list keeps growing. It's the fundamental reason for the field's lack of standards. The formats evolve, but the underlying physics stays consistent. Most coupling problems come down to finding either a clear optimum (a peak in received

power) or a centroid computable from sensor data. That's why a well-designed general alignment engine that can scan, evaluate a merit function, and optimize across multiple axes translates across a wide range of PIC products even as device geometries change year to year.

PI reinforces this generality through modular motion and control building blocks: stacked linear stages for XYZ travel, rotary axes for theta adjustments, and parallel-kinematic platforms (hexapods). On the control side, the E-713 motion controller supports synchronous, high-precision operation across up to 16 logical axes with mixed drive types, giving tool builders the flexibility to scale their architectures as products evolve. [R9]

Passive alignment motivation and why active alignment matters more than ever

The multi-decade interest in passive alignment approaches makes sense as a reaction to legacy active alignment's well-known slowness. Companies like Teramount are building around wafer-level self-aligning optics and wider assembly tolerances to enable higher-volume packaging with more conventional assembly lines. [R13] [R14]

But passive and active alignment are not actually at odds. Many products still need at least one active alignment step, or at minimum an active verification step, to account for real-world variability and squeeze out maximum coupling efficiency. And honestly, much of the push toward passive alignment is a pushback against outdated active alignment workflows specifically.

Modern active alignment, with parallel optimization, fast scans, and embedded control logic, changes the cost equation considerably. PI makes this explicit: the evolution from manual single-fiber methods to fully automated array alignment is a continuum, and modern tooling can combine passive placement with high-speed active optimization wherever that combination makes sense. [R15]



Adjacent semiconductor subprocesses where PI's capabilities translate

PIC packaging and test lines function as interconnected elements of broader advanced-packaging flows, not as discrete, independent processes. Several neighboring subprocesses draw on exactly the same core competencies PI brings high-precision pick-and-place, metrology and inspection positioning, lithography-related alignment. In most of these contexts, tool architects want coarse-fine implementations: long travel approach moves and format changes, then high-resolution nanopositioning for final alignment and drift compensation. PI's photonics alignment systems are built around hybrid mechanisms specifically designed to deliver long travel without trading away stability, resolution, or alignment speed. [R2]

➤ Compact Fast Alignment System for Photonics and Fiber Optics, F-141 PINovAlign®

Tooling architecture: flexibility, modularity, and EtherCAT-based scalability

One underappreciated constraint in PIC manufacturing equipment is architectural longevity. Tool builders and in-house integrator teams need platforms that can scale in axis count, support mixed actuator types, and stay serviceable across multiple product generations, not just for the

Dedicated webinars for the photonic integrated circuits (PIC) industry

Using our 30+ years' experience in B2B vertical technical markets, and as the publisher of PIC Magazine, we offer effective webinars, ZOOM interview and virtual events. We help you get your message, to your desired audience, by marketing to over 30,000 PIC professionals.

In addition to organising and managing your webinar, we can also market your webinar to our specialist databases.

Reach Educate Influence

- Brand Awareness
- Lead Generation
- Thought Leadership



PIC PHOTONIC
INTEGRATED CIRCUITS

Contact: Jackie Cannon
jackie.cannon@angelbc.com
+44 (0)1923 690205

device they're shipping today. Industrial fieldbuses like EtherCAT are a proven way to achieve deterministic multi-axis control while keeping the architecture modular and expandable.

PI highlights EtherCAT-capable controllers in photonics assembly contexts and positions its software stack to support fast application development across common operating systems and programming languages. [R16]

Looking forward: parallelization, embeddability, and partnership

Two requirements will be defined in the next phase. The first is parallelization: running more alignment processes simultaneously, whether in wafer-level testers or high-channel-count probe and coupling stations. The second is embeddability: compact form factors that fit inside dense tools and automation cells without giving up dynamic performance. PI's product and system demonstrators point in exactly

this direction, combining proven fiber alignment systems with multi-axis gantry automation to push toward more fully automated fiber array assembly workflows. [R17]

In the end, PI's value to PIC manufacturing is not any single product; it's an enabling platform: precision motion, firmware-level alignment intelligence, and scalable tooling architecture working together.

As PIC applications stretch from package-scale integration toward network and even 'planet-scale' visions, the manufacturing challenge converges on one thing: fast, repeatable, high-yield optical coupling at ever-higher channel counts. The companies that industrialize alignment, that make it fast, predictable, and scalable, will define what PIC manufacturing looks like at volume. PI is building the subsystems and control engines to make that happen.

REFERENCES


- [R1] PI (Physik Instrumente), "Active Alignment" (physikinstrumente.nl). Accessed 2026-03-12. <https://www.physikinstrumente.nl/en/expertise/technology/controllers-software/active-alignment>
- [R2] PI (Physik Instrumente), "Fast Multi-Channel Photonics Alignment Systems" (pi-usa.us). Accessed 2026-03-12. <https://www.pi-usa.us/en/products/more-products/fast-multi-channel-photonics-alignment-systems>
- [R3] PI (Physik Instrumente), News: "PI Halves Lead Time for F-712.HAx Fiber Alignment Systems" (includes FMPA performance claim). Accessed 2026-03-12. <https://www.physikinstrumente.com/en/about/news-press/news-detailpage/lead-time-f-712>
- [R4] PI (Physik Instrumente), "PiLightning Firmware feature for alignment systems". Accessed 2026-03-12. <https://www.physikinstrumente.com/en/products/fast-multi-channel-photonics-alignment-systems/pilightning-firmware-feature-for-alignment-systems>
- [R5] Nokia blog, "Evolving pluggable optics to reduce power consumption" (discussion of CPO power/latency). Accessed 2026-03-12. <https://www.nokia.com/blog/evolving-pluggable-optics-to-reduce-power-consumption/>
- [R6] Ansys blog, "The Future of Co-Packaged Optics" (benefits: latency, bandwidth, power). Accessed 2026-03-12. <https://www.ansys.com/blog/the-future-of-co-packaged-optics>
- [R7] NVIDIA Developer Blog, "Scaling AI Factories with Co-Packaged Optics..." (overview of CPO direction). Accessed 2026-03-12. <https://developer.nvidia.com/blog/scaling-ai-factories-with-co-packaged-optics-for-better-power-efficiency/>
- [R8] Tian et al., "Photonic Integrated Circuits: Research Advances and ..." (MDPI Photonics, 2025) – packaging/interconnect review. Accessed 2026-03-12. <https://www.mdpi.com/2304-6732/12/8/821>
- [R9] PI (Physik Instrumente), "E-713 Motion Controller" (physikinstrumente.com). Accessed 2026-03-12. <https://www.physikinstrumente.com/en/products/controllers-and-drivers/controller-systems-for-multiple-axes-mixed-drive-types/e-713-motion-controller>
- [R10] PI-USA, "Active Alignment" (FMPA described as firmware-level commands). Accessed 2026-03-12. <https://www.pi-usa.us/en/expertise/technology/controllers-software/active-alignment>
- [R11] PI-USA, "F-141 PINovAlign multi-axis fiber optics / photonics alignment system". Accessed 2026-03-12. <https://www.pi-usa.us/en/products/photonics-alignment-solutions/f-141-multi-axis-fiber-optics-photonics-alignment-system>
- [R12] PI-USA, Datasheet: "F-141 ... Fast Photonics Alignment System" (Aug 2024). Accessed 2026-03-12. https://www.pi-usa.us/fileadmin/user_upload/pi_us/files/product_datasheets/F-141_Fast-4DOF-6-DOF-Photonics-Alignment-System-Datasheet.pdf
- [R13] Teramount, "Technology" (self-aligning optics; passive tolerance concept). Accessed 2026-03-12. <https://teramount.com/technology/>
- [R14] Teramount, "About" (wafer-level self-aligning optics; passive alignment framing). Accessed 2026-03-12. <https://teramount.com/about/>
- [R15] PI-USA Tech Blog, "Photonics Alignment Techniques: Active vs Passive" (Aug 2024). Accessed 2026-03-12. <https://www.pi-usa.us/en/tech-blog/photonics-alignment-techniques-active-passive>
- [R16] PI Japan, "Assembly & Packaging of Photonic Devices" (mentions EtherCAT-capable controller + software stack). Accessed 2026-03-12. <https://www.japan-pi.com/en/knowledge-center/product-and-system-demonstrators/photonics-assembly-and-packaging>
- [R17] PI (Physik Instrumente), "Product & System Demonstrators" (multi-channel fiber assembly concept with gantry). Accessed 2026-03-12. <https://www.physikinstrumente.com/en/knowledge-center/product-and-system-demonstrators>




POWERING AI CONNECTIVITY

Key features of our Optical Interposer-based technology:

- Lower cost, silicon-based hybrid-integration platform
- Enabling next-generation interconnect solutions for 1.6T and above data rates
- Wafer-level chip-scale packaging (no wire bonds and no active alignments)




Our revolutionary products include:



POET Blazar

A high-power, multi-channel External Light Source. Designed to power both AI Scale-up and Scale-out Networks, including co-packaged optics.



POET Teralight

Tx and Rx optical engines for 1.6T and beyond pluggable transceivers.

Our Partners



FOXCONN

LITEON[®]

Adtran



poet-technologies.com

sales@poet.tech

[#POETpowersAI](https://twitter.com/POETpowersAI)



@POETtech



@POETTechnologies



@POETtechnologies



Shifting supply chains in the era of photonics

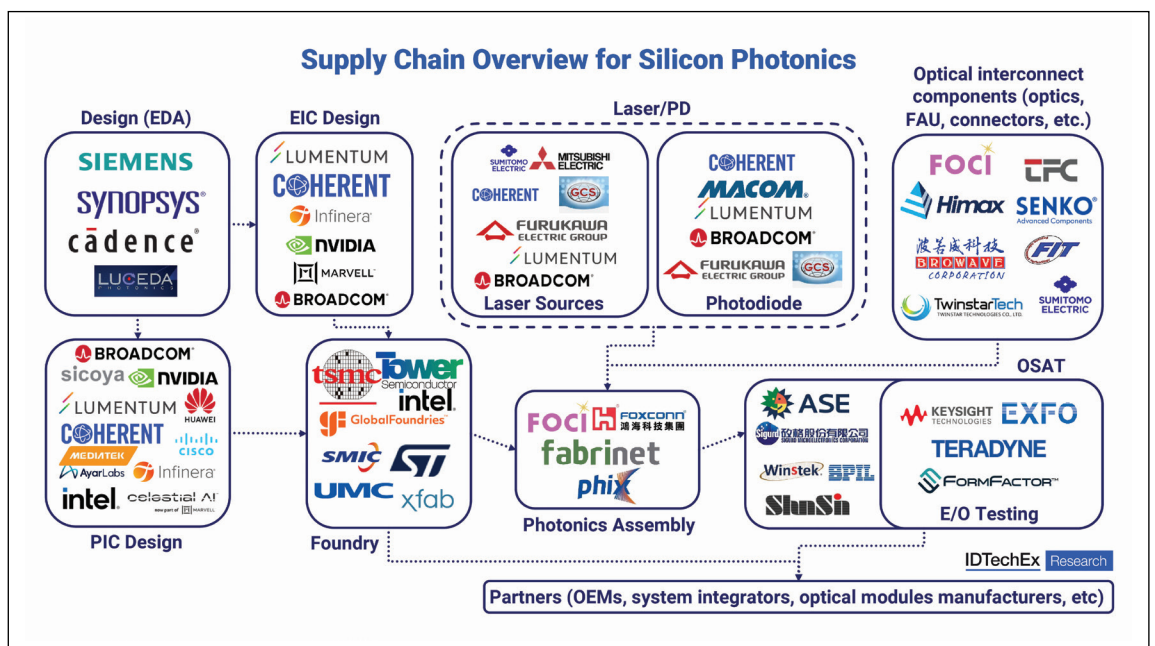
As AI drives the demand for interconnect bandwidth beyond the limits of copper, photonics is emerging as a critical part of the AI hardware stack. This article delves into the photonic supply chain for transceivers and CPO.

BY MIKA TAKAHASHI, SENIOR TECHNOLOGY ANALYST AT IDTECHEX

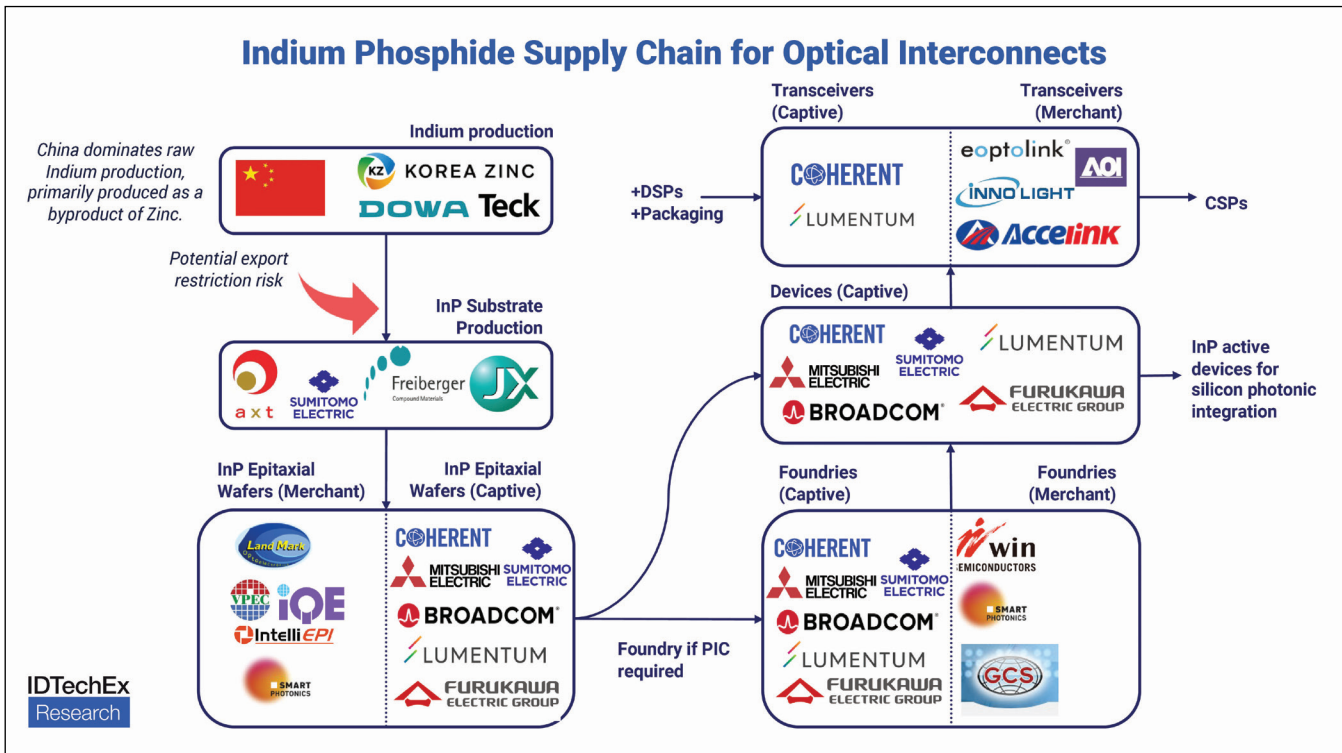
FOR DECADES, computational performance has been improved mainly by adding greater numbers of transistors on a single chip. The empirical observation that the transistor density doubles every two years (Moore’s law as its often known) has taken the integrated circuit (IC) from the early microchips with tens of devices to the NVIDIA Rubin with 336 billion transistors per chip.

However, in the age of AI, as demand for computing accelerates, the raw processing power is no longer the only limiting factor. How quickly and efficiently

data can be transmitted between chips, racks, and clusters is emerging as a structural bottleneck within data center design. Photonic integrated circuits (PICs) process and handle both optical and electrical signals on a single chip and sit at the heart of the optical shift within interconnect technologies. However, the advent of the (PIC) industry represents a transformational moment in the semiconductor industry, with new materials, new players, and entirely new processes leading to a major shakeup of the supply chain.



➤ IDTechEx has tracked the emerging silicon photonics supply chain across key segments. Source: IDTechEx



Copper domain is shrinking; optical domain is growing

The training and inferencing of large language models (LLMs) are too large to fit on a single chip, and so data centers string together thousands of graphical processing units (GPUs) together to run in parallel – creating a super cluster able to process the workload. In order to run efficiently, this requires an extraordinarily low-latency and high-bandwidth transfer of information between GPUs and network switches, creating a complex mesh of interconnects. Historically, interconnections within a data center have relied on a clear dichotomy: copper for short reach and fiber optics for long reach. Copper cabling is mature and reliable for short distances; however as data rates climb, the energy efficiency begins to degrade as electrons encounter signal loss in the cabling, suffer from crosstalk, and electromagnetic interference (EMI). This is compounded by the second major shift in data center design – increasingly physically large clusters, which in turn requires longer interconnects between pods. While active cabling can mitigate these effects to some extent by boosting the signal, as data rates climb, the power consumption required becomes unfeasible, and the industry approaches the so-called ‘Copper Wall’. The solution is to switch to optics. Photons do not experience resistive losses, travel 3x faster than electrons and can be manipulated in ways that allow for greater bandwidth.

Optical communication has long been used in the telecom and long-haul aspects of data transmission, but the AI revolution is expanding the optical domain closer to the chip at the expense of copper. However, switching from copper wiring to

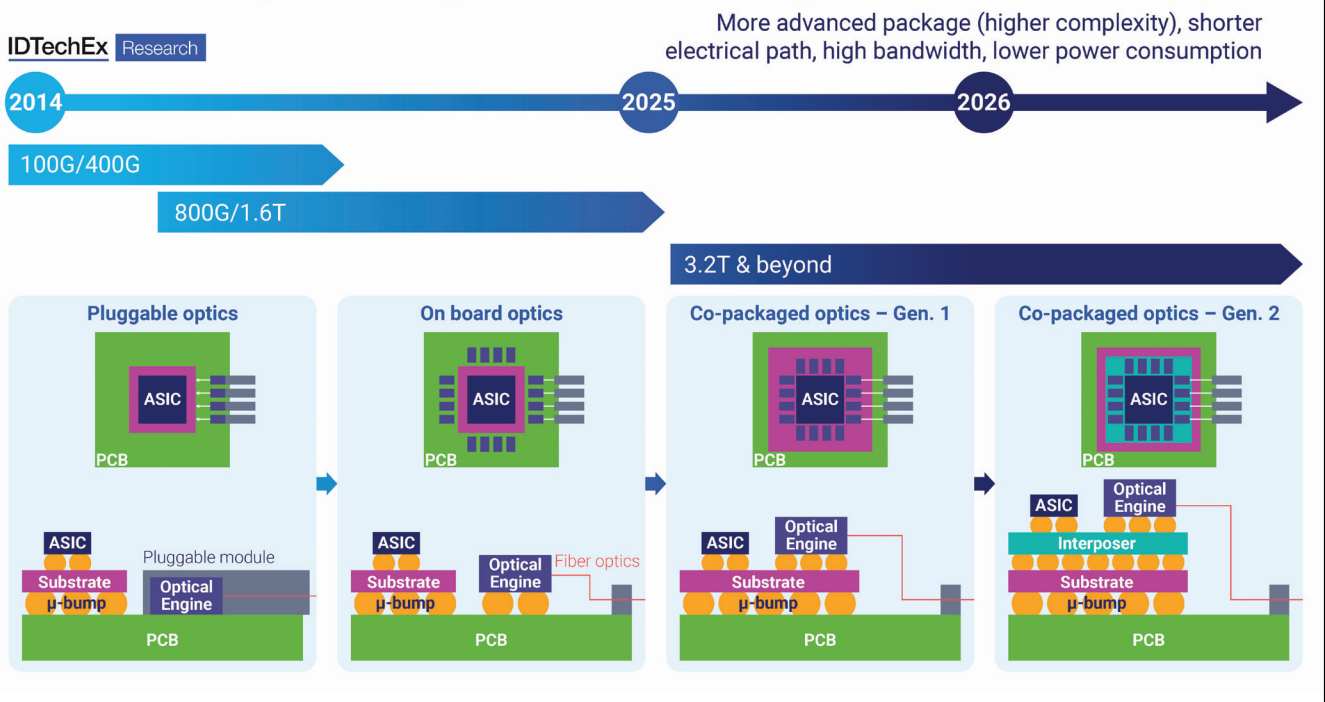
optical fiber is not straightforward, as the on-chip communications still all happen electrically due to the technical challenges of integrating optics on an xPU. Therefore, there needs to be a conversion between electric and optical and vice versa. This is where PICs have found their greatest use-case.

How exactly are PICs integrated into data center networking?

Optical transceivers handle the conversion from electrical signals on-chip to optical signals for longer distance transmission in fiber through the data hall. To do this, the transceiver must receive an electrical signal, generate an optical signal, encode the digital signal and transmit it to the network. It must also do the reverse on the receive side. Traditionally, these were built out of discrete sub-assemblies consisting of lasers, diodes, and lenses, attached manually. However, on a photonic platform, all these components are placed on an integrated circuit, allowing much greater component density and ease of manufacture using semiconductor processes such as lithography. In an ideal world, the industry would use silicon. Silicon is already used for the vast majority of logic IC applications and has a decades-old mature ecosystem with advanced process nodes. Silicon can also be used to build waveguides to transport light around the circuit and modulators to encode the light. Silicon, however, has a problem. It is an indirect bandgap semiconductor, and as such, cannot lase. This is the critical material challenge for PICs, requiring a non-silicon material to be integrated as a light source. Instead of silicon, a III-V semiconductor material (very often Indium Phosphide (InP)) is used either in conjunction with silicon or standalone.

➤ IDTechEx research has tracked the global InP supply chain from indium production to transceiver assembly. A significant portion of supply from epitaxial growth to device manufacturing is captive. Source: IDTechEx

Key Trend of Optical Transceiver in High-End Data Center



➤ Co-packaged optics (CPO) is an emerging technology that places the optical engine directly on the substrate or interposers. Due to thermal constraints current approaches to CPO require external InP lasers, which places unique demands on the lasers.

Indium Phosphide vs Silicon Photonics, a false dichotomy

There are two main types of PICs used in optical transceivers: silicon photonic and monolithic indium phosphide. Although there are future material platforms explored by IDTechEx, these are still to be commercialised at the same scale as silicon and InP. These are Monolithic InP builds the entire PIC – modulator, laser, photodiode, waveguides – out of InP. This avoids the complexity of bonding the InP laser to a silicon photonic chip (and the interface losses that entail) but requires a lot of a relatively expensive and supply-constrained material and a smaller wafer that has a much lower yield than silicon processes. Silicon photonics, on the other hand, uses silicon wherever possible and uses III-V semiconductors just for the lasers and photodiodes. One key misunderstanding often associated with photonics is the split between InP and silicon photonics. As of 2026, nearly all silicon

photonics devices will contain an InP laser, and as such, a shift towards greater silicon photonics adoption does not occur at the expense of demand for InP components, but it does change the nature of InP demand. InP and silicon also have fundamentally different supply chains and ecosystem maturities.

Photonics drives demand for indium phosphide

Indium is produced as a byproduct of zinc, and according to IDTechEx research, the majority of global supply (around 70%) is concentrated in China. This regional concentration adds an additional layer of risk, as China imposed export controls on a suite of materials – including indium – in April 2025. So far, these export restrictions only amounted to the requirement for licensing rather than a full export ban. However, it indicates the precarity of the global supply chain for photonics. Once processed, Indium is then combined with phosphide to make indium phosphide (InP) substrates, a highly specialised task dominated by Sumitomo Electric, AXT, JX Nippon, and Freiburger. These wafers are then sent to epitaxy foundries where crystalline layers are grown either by metal-organic chemical vapor deposition (MOCVD) or molecular beam epitaxy (MBE) reactors.

Demand for high-quality InP substrates has skyrocketed in recent years, and while producers are ramping up production, this is an emerging bottleneck in the AI hardware rollout. After this, a foundry fabricates the devices on the wafer, for example, lasers or modulators, and these are either sold as discrete devices, packaged into an InP optical transceiver, or integrated into silicon

IDTechEx research indicates that many of the leading InP players have a high degree of vertical integration. Coherent and Lumentum both control the supply from wafer to final device, be that a laser module or an optical transceiver

photonic devices. In general, yields are much lower than for silicon devices, and the wafer sizes available are much smaller. While silicon foundries have matured the 12" wafer size, for InP 2,3, and 4" are standard. With an increasingly supply-constrained market, there have been moves to increase the wafer platform size (and thus improve throughput). Coherent has been gradually shifting to a 6" platform, and several pure-play foundries such as the Netherland's SMART Photonics are also looking to shift to 6", but broadly speaking InP is still a far more specialised and boutique material platform than silicon. As a result, InP devices, wafers, and processes are still more expensive and in shorter supply than is possible with silicon photonics, driving an incentive to minimise the footprint of InP devices.

InP supply mostly captive

IDTechEx research indicates that many of the leading InP players have a high degree of vertical integration. Coherent and Lumentum both control the supply from wafer to final device, be that a laser module or an optical transceiver. In contrast, it is common for silicon photonic circuits to be manufactured at an external foundry (such as TSMC, Tower, or GlobalFoundries) before being shipped to an outsourced assembly and test (OSAT) partner.

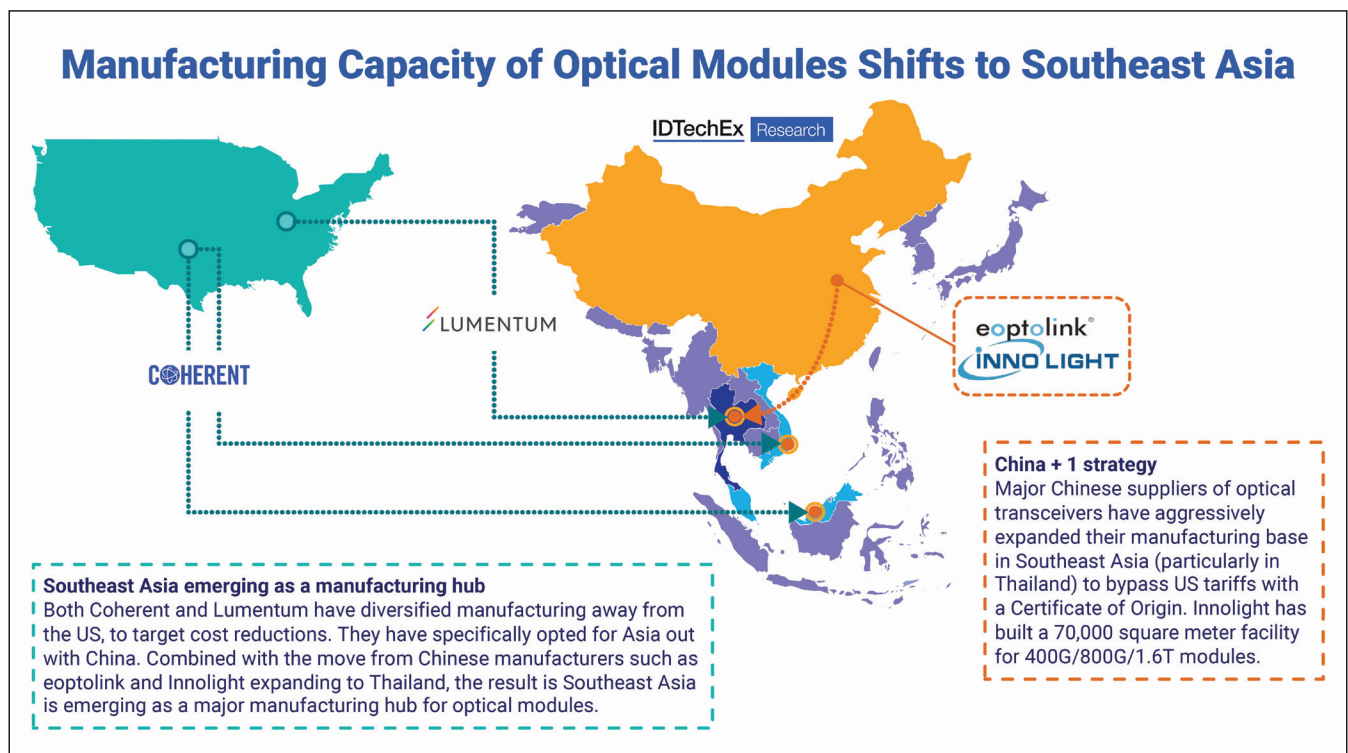
The InP industry has long been for specialised low-volume and high-value telecoms applications and thus has never been driven to develop a disaggregated global ecosystem. There is also a strong amount of proprietary knowledge in InP epitaxial growth that keeps device manufacturing

in-house. However, this bespoke low-volume production ecosystem is facing historically high demand. Leaders Coherent and Lumentum have both reported demand drastically outstripping supply, as major players such as Nvidia flex their leverage to secure long-term supply of InP devices. A downstream consequence of this IDTechEx has identified is the acceleration of silicon photonics adoption, which requires only simpler and CW (continuous wave) lasers that are externally modulated by a silicon modulator.

Ultra-high-powered lasers for co-packaged optics

One particular class of InP devices is emerging into the limelight with the advent of co-packaged optics (CPO). CPO takes the idea of optics a step further, eliminating the lossy copper trace between the pluggable optics and the ASIC entirely, placing the optical engine directly on the substrate. The benefits

There is also a strong amount of proprietary knowledge in InP epitaxial growth that keeps device manufacturing in-house. However, this bespoke low-volume production ecosystem is facing historically high demand



➤ Southeast Asia is emerging as a key manufacturing hub, particularly for optical transceivers.

are significant: lower latency, better signal integrity, and lower power consumption. However, there are substantial challenges, primarily associated with the thermal stability of the laser. Bonding a heat-sensitive laser on the same substrate as an ASIC that generates upwards of 80 degrees Celsius is a major engineering challenge, one that the industry has opted to avoid by using external laser sources (ELS).

Co-packaged optics does not yet entail co-packaged lasers. ELS places the laser in an external, air-cooled package at the face plate rather than within the xPU/ASIC package – mitigating the thermal issues but generating an entirely new issue of distance. Compared with conventional lasers for transceivers, for CPO, the ELS must generate a powerful enough light beam to traverse the extra distance and interfaces from the faceplate through to the ASIC before it can even begin to carry data, crossing additional connectors, fibers, and couplers. As a result, the 70mW class continuous wave lasers for silicon photonics are not enough – instead, ~300mW+ ELS are required. These are extremely challenging to manufacture, as increasing the power generally increases the signal noise, which, to mitigate, requires extremely high quality and low defects.

As of 2026, only a handful of companies can produce high-power low-noise (CPO capable) InP lasers. A key player in this space is Lumentum, which has leveraged its expertise in producing lasers for subsea telecoms applications – and according to IDTechEx research, is one of the leading players in the low-linewidth high power laser market. Lumentum, together with American giant Coherent, have both received a \$2 billion investment from NVIDIA, another sign of the AI hardware giant is seeking to secure strategic supply in advance of its rollout of CPO-based switches.

For the optical signals to avoid interface losses, the lenses, fibers, arrays, and light sources must all be bonded with sub-micron level accuracy

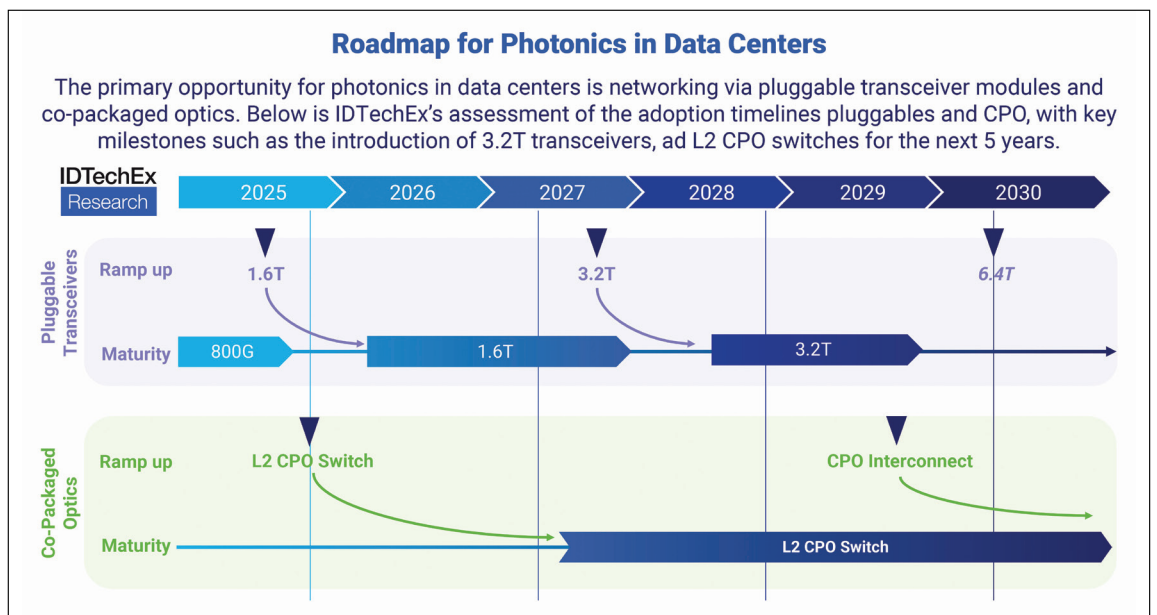
Emergence of new manufacturing hubs and a shift in the value chain

Another key trend of the photonics era is movement in the value and manufacturing chain. Optical transceiver assembly is driven by margin reductions, with major Chinese and American players offshoring to Southeast Asia to avoid potential tariffs and reduce costs, respectively. The region as a whole is emerging as a key region in the optics supply chain.

Silicon components themselves have a low unit cost, and the majority of the value added comes from the need for precise and high-throughput packaging and alignment. For the optical signals to avoid interface losses, the lenses, fibers, arrays, and light sources must all be bonded with sub-micron level accuracy. With the advent of CPO, the optical engine shifts from the optical transceivers to the network switch (and eventually the xPU) manufacturer, and as such, the value shifts to manufacturers of these devices.

With the advent of AI demand for compute has grown and begun to outpace conventional networking technologies. PICs play a critical role in the optical transformation and are unlocking huge advances in latency, energy efficiency, and aggregate bandwidth. Behind this transition, there is also a complex ecosystem bringing new materials, processes, and players into the spotlight.

➤ IDTechEx’s roadmap for photonics in data center networking anticipates that 3.2T will reach commercial maturity by the late 2020s, with CPO being implemented first for the L2 switches. Source: IDTechEx





**SAVE THE
DATE
12-14
APRIL
2027**

**Sheraton Brussels
Airport Hotel**

To find out more about our sponsor and speaker opportunities, contact us today on:

+44 (0) 2476 718970


or email:

info@csinternational.net

info@picinternational.net

info@peinternational.net

info@angel-tech.net



Integrating thin film lithium niobate on low loss PIC platform

Ligentec and UCSD have demonstrated 110 GHz modulators by integrating thin-film lithium niobate with low-loss silicon nitride waveguides on standard silicon wafers, addressing key processing incompatibilities that have previously constrained the scalability of photonic circuits.

BY MYRIAM LABIDI, CAMIEL OP DE BEECK, LIGEN TEC

PHOTONIC INTEGRATED CIRCUITS have progressed from research demonstrations to foundational components of commercial infrastructure. Applications in communications and quantum computing are driving demand for manufacturing platforms that simultaneously deliver high device performance and process scalability.

Lithium niobate offers outstanding electro-optic coefficients for high-speed modulation, while silicon nitride provides ultra-low-loss passive waveguides suitable for complex photonic routing, spot size converters and offers scalability using standard CMOS environment. Monolithic integration of these two materials on a single chip, however, presents substantial fabrication challenges.

The incompatibility is process-level. Stoichiometric SiN waveguides with

propagation losses that can achieve below 0.2 dB/m require LPCVD deposition, rendering this thermal budget inaccessible once LN is on the substrate. Additionally, lithium is a forbidden material in CMOS foundries due to its ability to diffuse into, and destroy, transistors on electronic ICs. This inhibits the installation of a monolithic process involving lithium niobate on the majority of today's semiconductor manufacturing infrastructure.

Material Integration: The temperature challenge

Unlike integration approaches that compromise on material quality, advanced hybrid platforms demand process-level compatibility. Stoichiometric silicon nitride's exceptional optical performance—propagation losses below 0.2 dB/m depends on high-temperature LPCVD deposition. Lithium niobate, however, begins to degrade above 550 °C,

creating a fundamental thermal incompatibility that cannot be resolved within a single sequential process.

Prior approaches using lower-temperature PECVD silicon nitride or reduced film thicknesses compromise the optical and confinement properties that make these materials attractive. The wafer bonding strategy described here resolves this constraint by processing each material independently under its optimal thermal conditions, then combining the two substrates through a low-temperature bonding step that preserves the properties of both.

Prior integration efforts have each involved trade-offs that limit overall platform performance. Depositing silicon nitride by plasma-enhanced CVD at reduced temperatures avoids the thermal constraint but yields films with higher propagation loss than stoichiometric LPCVD material.

Researchers from UC San Diego and LIGENEC have addressed these constraints through a process-partitioned integration scheme reported in the *Journal of Lightwave Technology*. Each material is processed under its respective optimal conditions and subsequently bonded at room temperature into a unified platform. The approach yields Mach-Zehnder modulators with 3-dB electro-optic bandwidths up to 110 GHz on standard silicon substrates.

A modular manufacturing strategy

The integration strategy is founded on process partitioning: each material is fabricated under its respective optimal conditions, and the two wafers are subsequently united through low-temperature direct bonding. Realising this approach at wafer scale required addressing a series of challenges, ultimately achieving defect-free bonding across approximately 95% of the wafer area.

LIGENEC's proprietary LPCVD process deposits SiN on a 200mm silicon wafer under optimal thermal conditions. LNOI wafers are prepared in parallel. Contact at room temperature initiates bonding through van der Waals and hydrogen-bond interactions.

The LNOI silicon handle and buried oxide are subsequently removed, leaving the lithium niobate film directly bonded to the underlying silicon nitride structures. Inspection confirms that approximately 95% of the bonded area is free from voids or delamination, a yield consistent with the requirements of volume photonic manufacturing.

The hybrid-mode innovation

A key design consideration in conventional thin-film lithium niobate modulators is the need to etch the LN layer to form ridge waveguides. Although etching provides strong optical confinement, plasma processing damages the crystal structure at the sidewall surfaces, introducing scattering losses and degrading the electro-optic coefficients that determine modulation efficiency. Additionally, etched lithium niobate waveguides often have non-vertical sidewalls, which can lead to unwanted polarisation conversion and mixing.

The present platform avoids lithium niobate etching entirely. Optical mode confinement and routing are determined by the geometry of the silicon nitride waveguides beneath the unpatterned LN film, a configuration referred to here as the hybrid-mode approach. The LN film is left in its as-deposited state, preserving the crystalline quality and electro-optic coefficients of the bonded material.

In passive regions, wider waveguides formed in the SiN layer confine the majority of optical power within the silicon nitride, rendering the overlying lithium niobate slab optically inactive. These waveguides exhibit low propagation loss and support tight bends, enabling compact routing of passive photonic structures.

In active regions, the mode is transitioned vertically through adiabatic tapers. At the end of the transition, approximately 53% of the guided power resides in the LN film, forming the hybrid mode that couples to the applied electric field via the Pockels effect.

Transition loss measured to be a few m dB per taper, a low value for a vertical inter-layer mode converter. The dual-layer silicon nitride architecture enables these low-loss transitions while remaining fully compatible with standard photolithography and dry-etch processing.


In passive regions, wider waveguides formed in the SiN layer confine the majority of optical power within the silicon nitride, rendering the overlying lithium niobate slab optically inactive

Performance that delivers

Mach-Zehnder modulators with phase-shift lengths of 6 mm, 9 mm, and 15 mm were characterised across multiple dies and wafer locations. Key performance metrics showed good uniformity, confirming process consistency at the wafer scale.

The voltage-length product $V_{\pi L}$ measured 3.8 V-cm. Agreement with predictions based on reported Pockels coefficients for thin-film LN at 1550 nm confirms the bonded film retains its intrinsic electro-optic activity.

The modulator extinction ratio exceeded 35 dB across all devices, indicating well-balanced splitting in both the couplers and the push-pull interferometric arms consistent across multiple dies and wafer locations.



➤ Schematic cross-section of the integrated platform showing SiN waveguides and bonded LN film, and electrode structure

Electro-optic bandwidth measurements yielded the most significant results. The 6 mm modulator achieved a 3-dB electro-optic bandwidth of 110 GHz; the 9 mm and 15 mm devices reached 70 GHz and 20 GHz, respectively. Bandwidth scaled inversely with phase-shift length, as expected for a travelling-wave electrode architecture. These values exceed prior demonstrations of silicon nitride–lithium niobate integration, providing quantitative validation of both the substrate removal strategy and the electrode design.

Looking ahead

This work demonstrates that material incompatibilities, which have historically constrained photonic integration, can be resolved through deliberate process partitioning. By processing each material under its optimal conditions and uniting them via controlled low-temperature wafer bonding, the researchers have established a viable pathway to monolithic platforms that combine best-in-class passive and active components on standard, cost-effective silicon substrates.

The platform’s modular architecture offers clear advantages for commercial

development: SiN photonic circuits can be independently designed and qualified using established foundry infrastructure, while LN is incorporated only where active functionality is required. The bonding interface also confers forward compatibility should superior electro-optic materials emerge; they can, in principle, be substituted without redesigning the base SiN platform.

The platform addresses requirements across multiple domains: high-speed coherent transceivers in telecommunications; phased-array antenna remoting and wideband signal processing in RF photonics; and large-scale quantum photonic circuits requiring low propagation loss and precise electro-optic control.

The commercial trajectory for this technology accelerated in January 2026, when LIGEN-TEC and X-FAB announced an expanded partnership targeting the industrialisation of thin-film lithium niobate integration on both silicon nitride and silicon-on-insulator platforms. The collaboration commits to scaling TFLN-on-SiN and TFLN-on-SOI technologies within X-FAB’s foundry ecosystem using


200 mm wafer processes precisely the manufacturing approach validated by the collaborative demonstration.

“By expanding our partnership with X-FAB, we are giving our customers early access to a broader technology stack that now includes high-volume thin-film lithium niobate integration. This technology enables cutting-edge performance and helps our partners stay at the forefront of innovation in sensing, connectivity and computing applications.” - Thomas Hessler, CEO, LIGEN-TEC

The partnership establishes a defined pathway from advanced research to industrial-scale production, reinforcing Europe’s position in next-generation photonic integration and signalling broader industry confidence in the wafer-bonding approach.

Continued refinement is already yielding improvements. Long-term reliability testing across temperature and environmental conditions will be essential for qualifying the technology for deployment.

The results presented here confirm that material incompatibilities previously regarded as fundamental barriers to photonic integration can be overcome through disciplined process engineering and a systematic re-evaluation of fabrication sequences. The approach achieves both the device performance required by next-generation applications and the process compatibility required for scalable manufacturing outcomes that have historically been difficult to realise simultaneously on a single integrated platform.



➤ Photograph of a bonded wafer showing successful integration across the full diameter

FURTHER READING

- A. Rahman et al., “Integration of Hybrid Thin-Film Lithium Niobate Electro-optic Modulators on a Wafer-scale Silicon Nitride Photonics Platform,” *Journal of Lightwave Technology*, vol. 44, no. 1, pp. 1–11, January 2026. DOI: 10.1109/JLT.2025.3646212

Available at: <https://ieeexplore.ieee.org/document/11304539>



SMC STRATEGIC MATERIALS CONFERENCE

July 13-15, 2026 | San Jose, CA



Register at www.semi.org/smc



**SEMICON
WEST**

OCT 13-15, 2026

Moscone Center, San Francisco, CA

TRANSFORM TOMORROW



Register at www.semiconwest.org

Electro-optic polymers move from promise to production

As AI data centres push networking to its limits, electro-optic polymers are moving beyond laboratory demonstrations. Advances in reliability, foundry integration and packaging now position them as a production-ready technology for next-generation optical interconnects.

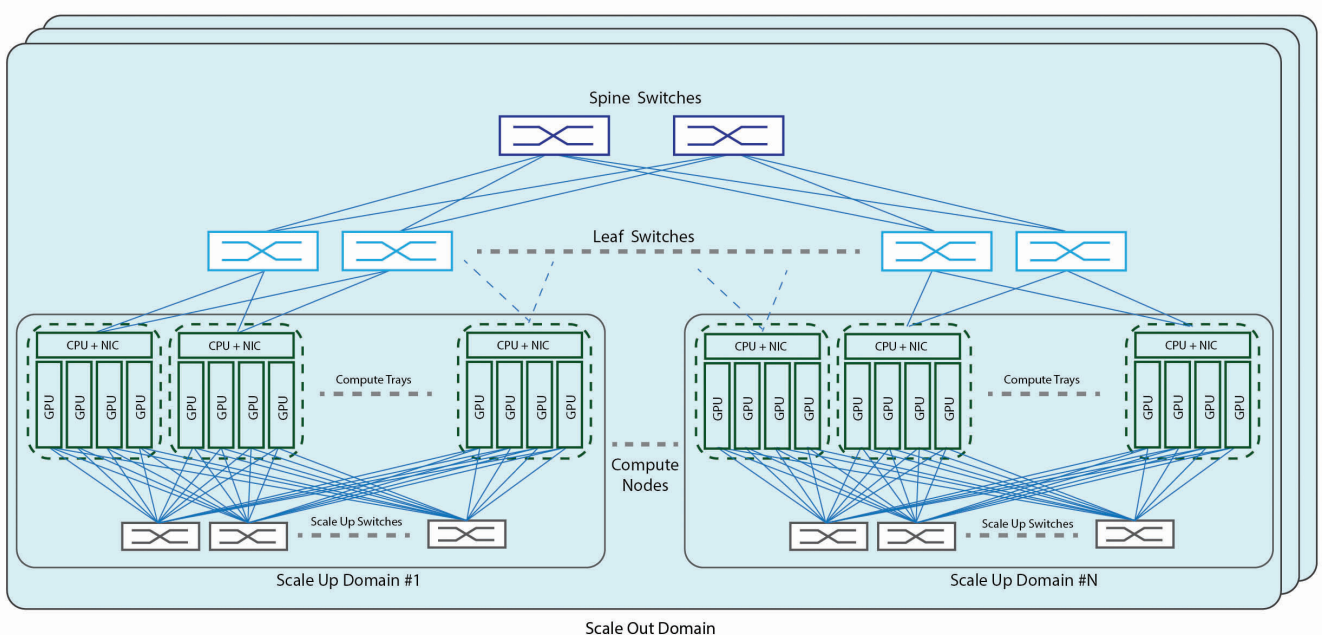
BY ROBERT BLUM, SENIOR VICE PRESIDENT, SALES AND MARKETING, LIGHTWAVE LOGIC

ARTIFICIAL INTELLIGENCE continues to drive unprecedented demand for computing, but it is increasingly clear that processing capability alone no longer defines overall system performance. As AI training clusters scale to tens of thousands of GPUs, efficiency depends just as much on how rapidly data can move between processors. In large-scale deployments, the network is no longer a background consideration; it has become a first-order design constraint.

This shift is already reshaping requirements for optical interconnects. Transceivers and the modulators within them must operate at ever higher

speeds while consuming less power and occupying smaller footprints. Electrical lane rates are moving from 200 Gbit/s towards 400 Gbit/s, and optical platforms must evolve accordingly. Incremental improvements are proving insufficient. Instead, the industry is being forced to re-evaluate materials and device concepts that can sustain this scaling without imposing unacceptable penalties in power, size, or manufacturability.

Electro-optic (EO) polymers sit squarely in this discussion. Long regarded as an intriguing laboratory technology, they are now being assessed in a much more pragmatic context: whether they can



➤ AI data-centre networking bottleneck: As GPU clusters scale, overall system performance increasingly depends on the ability to move data efficiently through optical interconnects, elevating the importance of modulator speed, power and density

meet the reliability, integration and volume-manufacturing requirements demanded by modern data-centre infrastructure.

A long history of promise – and scepticism

EO polymers are by no means new. Researchers have explored their use in optical modulators for several decades, drawn by their ability to exhibit large and rapid changes in refractive index under an applied electric field. Unlike many inorganic materials, EO polymers do not rely on the transport of charge carriers to achieve modulation. Instead, their response is dominated by electronic polarisation, enabling extremely fast operation.

Despite these advantages, adoption has historically been limited. The reasons are familiar to anyone who has followed the transition of new materials from laboratory to factory. Questions around long-term stability, sensitivity to temperature, oxygen and humidity, and compatibility with standard semiconductor manufacturing flows all contributed to hesitation. In an industry built on yield, reliability, and repeatability, such caution is not only understandable – it is essential.


However, there is a useful parallel to the evolution of organic light-emitting diodes (OLEDs). OLED materials were once widely viewed as too fragile for mainstream deployment. Their eventual success did not stem from a fundamental change in organic chemistry but from learning how to protect the materials effectively. Encapsulation technologies transformed OLEDs from laboratory curiosities into mass-market products found in smartphones, televisions, and wearables.

EO polymers are now following a similar trajectory.

What makes an electro-optic polymer different?

At the heart of an EO polymer lies a deliberately engineered molecular structure. Dipolar chromophores embedded within a polymer matrix respond directly to an applied electric field by changing the refractive index of the material. Because this process relies primarily on electronic polarisation rather than carrier transport, it can occur extremely quickly, enabling modulation well beyond the limits of conventional semiconductor approaches.

One practical consequence is a very large electro-optic coefficient, typically on the order of several hundred picometres per Volt. In device terms, this allows a substantial optical phase shift to be achieved with drive voltages well below one Volt. For system architects, the implications are straightforward: lower electrical power consumption,



simplified driver electronics and compact modulator geometries that help preserve bandwidth density at the chip edge.

At Lightwave Logic, these considerations have guided the development of the Perkinamine electro-optic polymer platform. Rather than optimising for a single headline performance metric, Perkinamine has been engineered with system-level deployment in mind. The polymer matrix exhibits a high glass-transition temperature, enabling compatibility with standard solder-reflow assembly processes. At the same time, the material is designed to integrate cleanly with silicon photonics, complementing a proven high-volume wafer-level manufacturing platform rather than requiring a fundamentally new one.

Redefining “production readiness”

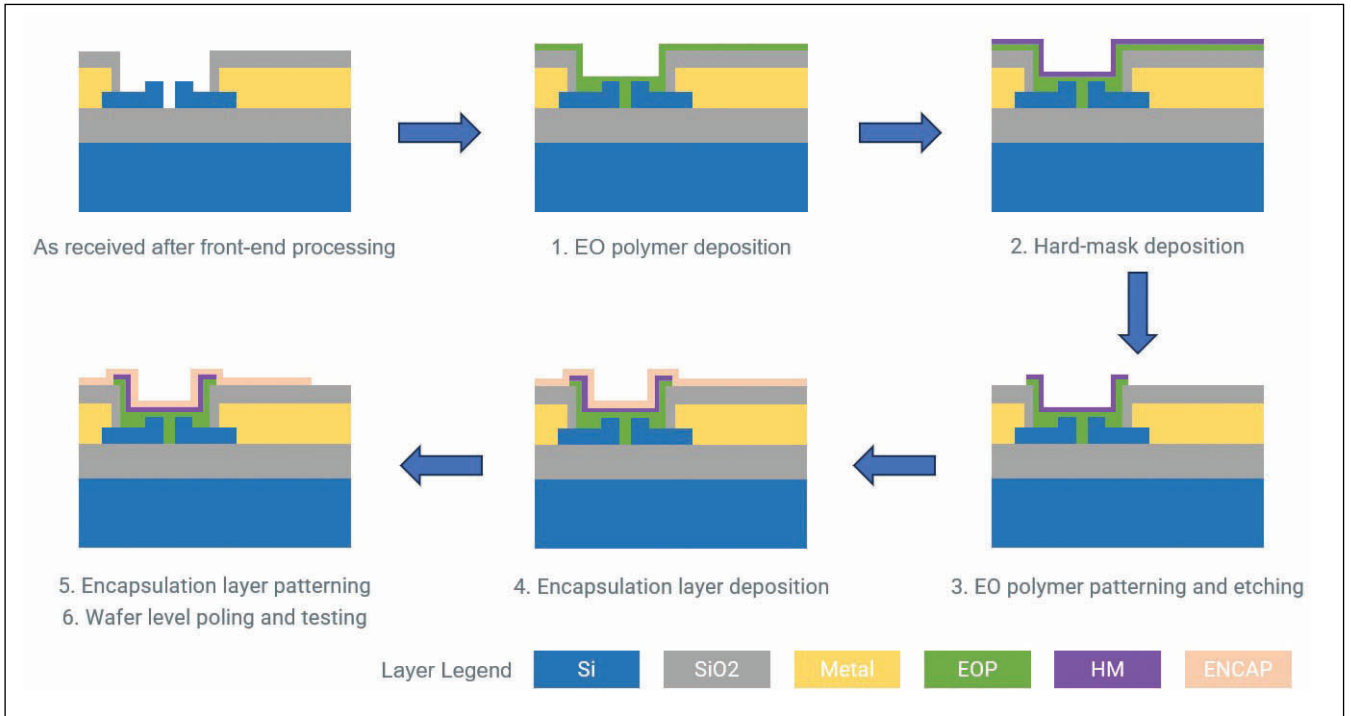
When engineers talk about production readiness, they rarely mean a single attribute. Readiness emerges from the alignment of several factors: material performance, long-term reliability, process compatibility, and ecosystem integration. For EO polymers, each of these dimensions has advanced significantly in recent years.

Reliability under real operating conditions

Optical modules in data centres operate at elevated temperatures and are expected to function reliably for many years. EO polymers must therefore maintain their electro-optic properties under

➤ Schematic of an electro-optic polymer modulator: In an EO polymer modulator, an applied electric field changes the refractive index of the polymer, enabling high-speed optical modulation with low drive voltage

Rather than optimising for a single headline performance metric, Perkinamine has been engineered with system-level deployment in mind. The polymer matrix exhibits a high glass-transition temperature, enabling compatibility with standard solder-reflow assembly processes



➤ Integration of EO polymers in a silicon photonics foundry flow: Wafer-level deposition and encapsulation of EO polymers can be incorporated into standard silicon photonics process design kits, supporting scalable manufacturing

sustained thermal stress – often around 85 °C – while resisting degradation from environmental exposure.

Advances in molecular design have improved intrinsic stability to the point where projected lifetimes, even at high operating temperatures, extend well beyond typical product requirements. Equally important, the industry now has a much clearer understanding of the dominant failure mechanisms. Experience has shown that exposure to oxygen and moisture, rather than inherent material fragility, drives much of the long-term degradation observed in early devices. Once these mechanisms are understood, effective mitigation becomes possible.

Encapsulation as an enabler, not an afterthought
Encapsulation has emerged as a critical enabler for EO polymer deployment. By surrounding the active material with robust barrier layers, engineers can prevent oxygen and humidity from reaching the polymer, preserving performance over time.

Here, the broader semiconductor industry has provided valuable tools. Over the past decade, thin-film deposition and barrier technologies have matured significantly and are now widely available in foundries. Techniques that were once specialised can be integrated into standard process flows, allowing EO polymer encapsulation to be treated as part of routine manufacturing rather than an exotic add-on. As with OLEDs, production readiness ultimately depends as much on system-level engineering as on the intrinsic properties of the material itself.

Compatibility with silicon photonics manufacturing
Even a reliable material will struggle to gain traction

if it requires disruptive processing steps. One of the strengths of modern EO polymer platforms is their compatibility with established fabrication techniques. Deposition processes can closely resemble those used for photoresists, allowing polymers to be applied at the wafer level using familiar equipment.

This compatibility lowers the barrier to adoption. Foundries and device designers can integrate EO polymers into existing silicon photonics flows without re-architecting their manufacturing infrastructure. For an industry that values continuity and predictability, this is a significant advantage.

The importance of the process design kit

In silicon photonics, the process design kit (PDK) plays a central role in translating materials into manufacturable products. A PDK encapsulates the knowledge required to design for volume: process steps, design rules, performance models and verified layouts. Without this framework, even promising technology remains confined to R&D demonstrations.

The emergence of EO polymer PDK modules therefore represents a meaningful inflection point. By defining how materials such as Perkinamine are deposited, encapsulated and integrated alongside silicon photonic components, these kits allow designers to treat EO polymers as part of the standard design toolbox.

At Lightwave Logic, this effort extends beyond material deposition. The PDK also includes guidance on assembly and test, including reference assembly flows for different packaging concepts. This ensures

that once a design is proven, customers can transition efficiently towards volume manufacturing rather than encountering new integration challenges at a later stage.

Power and density: why polymers matter now

As data rates climb, power consumption becomes a dominant system-level concern. Every additional milliwatt dissipated by a modulator is multiplied across thousands of links and racks, contributing to cooling challenges and operational cost. Lower drive voltage and reduced power dissipation are therefore not merely attractive features; they are fundamental requirements for scalable AI infrastructure.

EO polymers address this challenge directly. Their high electro-optic coefficient enables efficient modulation with minimal electrical overhead. In parallel, their material properties support compact device geometries, allowing higher bandwidth density at the edges of chips and packages.

What is particularly notable is that this shift is already visible at 200 Gbit/s. At these lane rates, silicon photonics devices enhanced with electro-optic polymers can offer measurable advantages in performance and power efficiency, prompting customers and foundries alike to engage earlier in the design cycle than in previous technology transitions. As a result, evaluation and integration efforts are no longer being deferred until the 400 Gbit/s inflection point.

This density advantage becomes increasingly important as the industry explores advanced packaging concepts such as co-packaged optics. In these architectures, optical interfaces sit in close proximity to GPUs or switches, and space becomes a scarce resource. Smaller, lower-power modulators simplify integration and improve overall system efficiency.

Competing technologies and practical differentiation

Several alternative electro-optic technologies have demonstrated impressive performance and, in some cases, have reached a higher level of maturity in early deployments. However, these approaches often introduce structural challenges as systems move towards volume production. Integration outside standard silicon photonics foundry flows, increased process complexity and reliance on relatively concentrated supply chains all become more significant considerations as lane rates increase and deployment scales.

In this context, electro-optic polymers appear better aligned with the industry's long-term direction, where manufacturability, supply-chain resilience and seamless foundry integration increasingly outweigh early-mover advantage.

A business model aligned with the ecosystem

Another indicator of maturity lies in how EO polymer suppliers engage with the market. Rather than competing directly with transceiver or system vendors, companies such as Lightwave Logic focus on providing materials, intellectual property, and reference designs.

This approach allows partners across the value chain to adopt EO polymers without conflict, accelerating experimentation and deployment. By positioning EO polymers as an enabling technology rather than an end product, suppliers can support a broad range of applications while leveraging the expertise of established device and component manufacturers. This ecosystem-centric model reflects lessons learned across the wider semiconductor industry.

Milestones that signal readiness

Scepticism around new materials will not disappear overnight, nor should it. In high-volume manufacturing, confidence is built through qualification data, repeatable processes and multiple independent design cycles.

For EO polymers, the most meaningful signals of maturity include adoption within foundry PDKs, customers progressing through advanced modulator designs, successful completion of industry-standard reliability testing and a clearly defined path to volume wafer processing and packaging. These developments are now beginning to align, shifting the discussion from whether EO polymers can be deployed to when and where they will first see large-scale adoption.

Looking ahead

The demand for bandwidth shows no sign of slowing. If anything, the pace of increase is accelerating as AI workloads continue to scale. Optical technologies must therefore deliver not only the next generation of performance, but a sustainable roadmap beyond it.

EO polymers are increasingly well positioned in this context. Their combination of speed, power efficiency and manufacturability addresses immediate system constraints while leaving room for future growth. Perhaps most importantly, the industry now has the tools, experience and infrastructure required to deploy them reliably at scale.

After many years of development, electro-optic polymers are approaching a practical inflection point. What was once primarily a laboratory technology is now emerging as a production-ready option for next-generation optical interconnects, shaped as much by advances in manufacturing and system design as by materials science itself.

Why optical scale-up must be multiplexed

Single-wavelength CPO works for scale-out, but future AI system scale-up demands a fundamentally different architecture built around multi-wavelength slow and wide networks.

BY MATT CROWLEY, CEO, SCINTIL PHOTONICS

➤ Scale-up clusters demand ten times the bandwidth of scale-out networks, creating interconnect bottlenecks only high-speed, dense, and low-power optical can solve. Source: Scintil Photonics.

IN MARCH 2026, AMD, Broadcom, Meta, Microsoft, NVIDIA, and OpenAI announced a joint consortium to establish an open specification for optical scale-up interconnects. The architecture they aligned on is slow and wide NRZ modulation with wavelength-division multiplexing. When six of the most consequential companies in AI infrastructure publicly agree on the same optical interconnect design for the networks that tie GPUs together into a single scale-up network, then the question of direction is settled.

This announcement definitively settled one of the longest-running debates in datacenter optics: “When will CPO reach the scale-up network?” The next question is no longer whether optical scale-up networks will use wavelength multiplexing, but how fast the industry implements this architecture and how fast it evolves from four wavelengths today to eight, sixteen, and beyond.

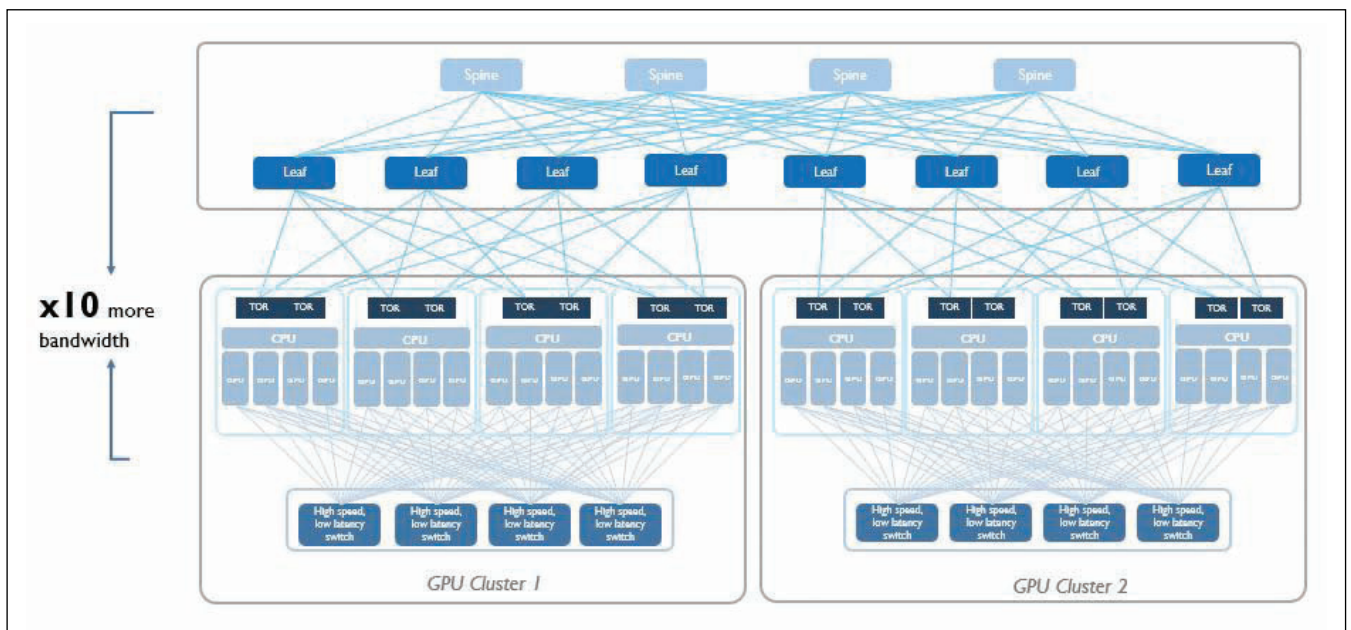
Wavelength multiplexing is not a new idea. DWDM was deployed in the telecom industry twenty-five

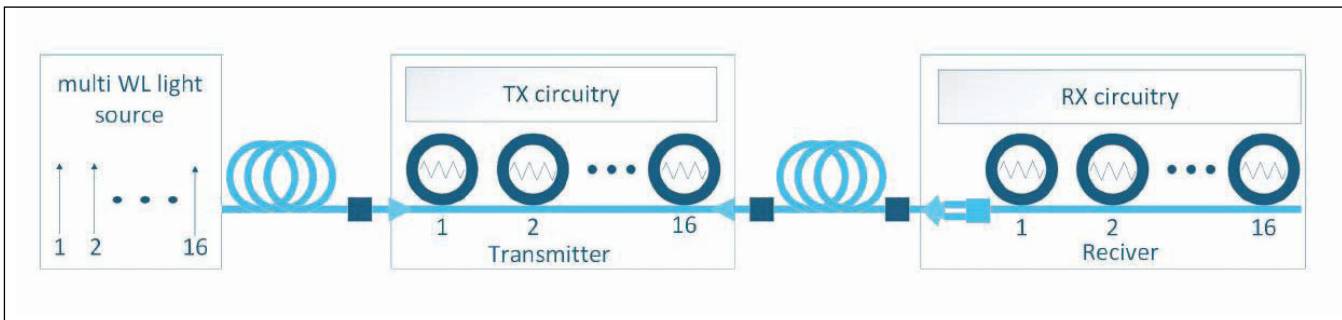
years ago. In principle, the idea is simple: instead of pushing a single channel per fiber, running many channels and scaling bandwidth through multiplexing. Applying this successful optical networking approach to AI scale-up networks is about to reshape AI datacenter architecture the same way DWDM reshaped long-haul telecommunications.

Slow and wide goes fast

The conventional approach to increasing interconnect bandwidth is brute force: push symbol rates higher, move from simple 2-voltage NRZ encoding to more complex 4-voltage-level PAM-4 (up to 400Gbps), and carry more data per fiber. This works in theory, but each step up in modulation bears a cost. The increased complexity incurs higher latency, power, and costs, as well as increased heat and space requirements. Single wavelength per fiber solutions do not scale up the scale.

Wavelength multiplexing resolves this by scaling bandwidth per fiber through increasing channel





count rather than increasing modulation rates. Each wavelength runs at a rate well within NRZ reach, typically 25 to 100 Gbps, depending on bandwidth needs and power budget. At these low rates, the encoding is simple; the signal processing is minimal, and forward error correction can be shallow or eliminated entirely. Aggregate bandwidth scales through multiplication: double the wavelengths, double the bandwidth, without changing the per-channel electronics. The network operator can now choose to run anywhere from 400 Gbps (50 Gbps X 8 λ) to 1.6 Tbps (100 Gbps X 16 λ), depending on bandwidth needs and power budget.

DWDM scale-up is the architecture NVIDIA described in their published work on a “slow and parallel” interface, [A Roadmap Toward Sub-1 pJ/b Optical Interconnect](#). This landmark paper models a 16-wavelength DWDM microring interconnect and projects a baseline of 3.8 pJ/bit today with a path toward sub-1 pJ/bit through improved receiver sensitivity, integrated optical gain, and more thermally stable rings. It is now the architecture codified by 6 of the industry’s leading companies in the recent OCI MSA specification.

The tail latency tax

Tail latency is the silent killer of GPU cluster economics. When a thousand GPUs execute a collective operation, they move at the speed of the slowest processor. A single delayed bit leaves the rest of the cluster to idle, which hurts utilisation. As cluster size grows from 72 to hundreds, then thousands of processors, the statistical likelihood of a tail-latency event on any given operation explodes, as does the risk of an underutilised datacenter.

The slow-and-wide architecture attacks this directly. When per-channel bandwidth is below 100Gbps, NRZ suffices. The processing overhead in terms of time, power, and potential errors diminishes, and latency is reduced and made more deterministic. A low tail latency network is precisely what memory-coherent scale-up domains require for maximum performance and efficiency.

The wavelength staircase: 4 → 8 → 16 → beyond

The consortium’s first-generation specification starts at four wavelengths based on what is practically possible to make using traditional laser manufacturing methods. This is a pragmatic

beginning. It gets the ecosystem moving, proves the WDM CPO concept in production of silicon, and aligns the supply chain for the imminent surge in laser demand. But four wavelengths at 53.125 Gbps NRZ per channel deliver only 212.5 Gbps per fiber in each direction. That is meaningful, but it is not yet transformative for multi-rack scale-up clusters that need multi-terabit links at the package’s edge.

The real scaling begins at eight wavelengths. Double the bandwidth per fiber without increasing symbol rate or modulation complexity. At sixteen wavelengths, the architecture unlocks transformative bandwidth for scale-up networks. At this wavelength count, bandwidth density up to 1.6Tbps per fiber at the package edge reaches the regime where fundamentally larger clusters become architecturally feasible without a fiber-count explosion at the package boundary. And the same scaling law continues to thirty-two wavelengths and beyond until network bandwidth and latency are no longer the bottleneck in datacenter performance. Each wavelength increases the capacity on existing fiber infrastructure with no change to the fundamental link design. The per-channel electronics remain the same. The fiber plant remains the same. Only the light source scales.

Future proof your design


For teams designing co-packaged optics today, this trajectory has immediate practical implications. Architectures being laid out now will determine which AI systems can support thousand-GPU cluster configurations in 2028 and beyond. The optimal design is future-proof and can be optimised not only for 4 wavelengths but also beyond. These choices preserve a path to eight and sixteen wavelengths without requiring a major redesign and prevent developers from missing market cycles as wavelength counts jump. An architecture that cannot scale past four wavelengths will need to rework when the specification evolves, and it will evolve, because the physics of scale-up networking demand it.

The scaling law photonics has been waiting for

The semiconductor industry has been through this inflection point twice before.

Dennard scaling held for three decades: as transistors shrank, power density stayed constant, and each new process node delivered higher clock speeds

➤ Example DWDM optical interconnect from NVIDIA’s “A Roadmap Toward Sub 1 pJ/b Optical Interconnect” paper, with an external multi-wavelength light source feeding 16 microring modulators and 16 microring resonators at 100-GHz spacing. Source: NVIDIA.



➤ Optical spectrum from a LEAF Light™ single-chip DWDM light engine showing 16 simultaneous wavelength outputs (F0–F15) at approximately 100 GHz channel spacing, with ± 10 GHz wavelength accuracy across all channels. Source: Scintil Photonics.

at no power penalty. That relationship broke down around the 90nm node in the mid-2000s. Leakage currents overwhelmed the voltage scaling that had kept power density in check, and the industry's free ride on frequency ended. Clock speed plateaued. The industry adapted with multicore architectures, parallelism replacing raw speed.

Moore's Law continued to pack more transistors onto a chip, but the economic engine underneath it stalled at 28nm. Transistor cost scaling, the 0.7x cost reduction that had compounded with every generation since the 1960s, stopped. Every node since 28nm has been more expensive per transistor, not less. As confirmed by Google at IEDM 2023, cost per transistor is now flat to rising, generation over generation. The industry adapted again, this time with chiplets and heterogeneous integration, assembling systems from specialized dies rather than scaling monolithic chips.

Photonics is arriving at its own inflection, but from the opposite direction. The traditional photonics supply chain has been scaling-limited from the start. Each laser is a discrete device. Each wavelength requires a separate assembly step. Costs scale linearly with channel count because there has been no manufacturing mechanism to bend that curve. This is why, despite decades of investment in silicon photonics, the multi-wavelength laser source has remained the unsolved component: the one piece that could not ride a cost-reduction curve because no such curve existed.


Heterogeneous integration, a fundamental breakthrough in optical manufacturing process technology, changes this equation fundamentally. When III-V gain material is bonded directly onto a silicon photonics wafer in a foundry-aligned process, the laser ceases to be a special-purpose discrete component. It becomes just another circuit element on the die, fabricated, tested, and scaled using the same wafer-level infrastructure that produces millions of photonic devices today. The cost of adding the next wavelength decreases with each process

generation, not because the physics change, but because the manufacturing learning curve follows the same pattern as for every other semiconductor-adjacent technology. Importantly, heterogeneous integration not only scales active devices such as lasers, SOAs, and EAMs but also encompasses major photonic circuit elements. This paradigm shift in ultra-integration will allow the photonics industry to finally achieve the scale and design flexibility that enabled CMOS to thrive for decades.

This is what makes wavelength count the relevant scaling law for optical interconnects. Where Dennard scaling gave the industry a free ride on frequency, and Moore's Law gave it a free ride on transistor cost, wavelength scaling through heterogeneous integration gives photonics its own compounding curve: more bandwidth per fiber, per millimetre of package edge, per microwatt, with each generation improving upon the prior one.

The semiconductor industry learned, both times, that when a scaling law breaks, the replacement is not incremental. It is architectural. Heterogeneous integration is poised to transform photonics from a scaling-limiting technology into a scaling-enabling one. The teams that recognize this transition and design for it now will hold the architectural position when optical scale-up interconnects become the standard configuration for AI infrastructure.

This is the problem we set out to solve at Scintil Photonics. Our SHIP™ heterogeneous integration technology bonds III-V material directly onto silicon photonics in a foundry-aligned process flow validated at Tower Semiconductor. We have already demonstrated our LEAF Light™ single-chip DWDM laser sources in eight- and sixteen-wavelength configurations compatible with micro ring-based CPO transceivers. The manufacturing path runs on established production lines. The cost curve bends with volume. For teams designing scale-up optical interconnects for what comes after the first generation, that path is open.



➤ The LEAF Light™ single-chip DWDM laser source; a foundry-aligned heterogeneous photonics integration die producing 16 multiplexed wavelengths for co-packaged optics applications. Source: Scintil Photonics.

To promote your products and services contact:

Shehzad Munshi

T: +44 (0)1923 69015

E: shehzad.munshi@angelbc.com



Tel: +1 215-441-0107
eillg@avophotonics.com
www.avophotonics.com



www.epic-assoc.com



Tel: +1 418-683-0211
francois.couny@exfo.com
www.exfo.com/en



Tel: +49 4202 51160-435
info@ficontec.com
www.ficontec.com



http://www.iqep.com




info@ligentec.com
www.ligentec.com



Tel: +32 (0) 52 339 838
erwin@lucedaphotonics.com
www.lucedaphotonics.com/en



Tel +1(604) 733-9006
info@lumerical.com
www.lumerical.com



Tel: +1-613-224-4700*249
sawyer.ge@optiwave.com
www.optiwave.com



Tel: +49 721 4846 1802
j.denzer@pi.de
www.physikinstrumente.com/en



Tel: +1-781-762-3901 / +49-30-398058-0
info_pic@VPIphotonics.com
www.vpi-photonics.com



To advertise here call
Shehzad Munshi
+44 (0)1923 690215
shehzad.munshi@angelbc.com

Showcase your presence to the global PIC industry for 12 months

PIC MAGAZINE is the primary information resource for professionals working with PIC materials and devices.

Each issue of the magazine is distributed digitally to over 28,000+ professionals worldwide.

With the largest and most diverse global audience of qualified buyers served by any photonic focused magazine, secure a corporate partnership and buyers guide entry and receive 12 months of promotion and exposure to the PIC industry across the magazine portfolio including:

- 12 months of promotion in quarterly issues of the digital and PDF magazine featuring your company logo and company details
- 12 months of promotion and exposure on the PIC Magazine website www.picmagazine.net featuring a Buyers Guide entry and hyperlinked company logo
- Sponsorship of one e-news alerts sent out to over 15,000 industry professionals with text and a standard web banner.

To find out more, contact: Shehzad Munshi
E: shehzad.munshi@angelbc.com
T: +44 (0)1923 690215



Scale with Confidence Infrastructure Test for the AI Revolution

VIAVI MAP-300 is optimized for compact, high-throughput silicon photonics testing - from development through manufacturing - from wafer through packaging devices.

- Superior testing speed
- Lowest cost-of-ownership
- High-density hardware with multithreaded test execution
- Global footprint with expert local support
- Triggering with deep buffer capture



Learn More:

<https://www.viavisolutions.com/en-us/products/map-300-multiple-application-platform>