

TROY TECHNICAL

OPTICAL COMMUNICATIONS & PHOTONICS WEEKLY REPORT

October 3 - October 9, 2026 | Issue 2026-W40

**Link budget, packaging, thermal reliability and
deployable supply**

12

ANALYZED ARTICLES

10

PRIORITY THEMES

Issued 2026-10-09 | English edition | fixed approved cover artwork



Optical Communications & Photonics: the boundary moved from isolated claims to qualification evidence

12 English articles were reviewed independently from the Japanese edition. Ten themes are retained for management decisions.

SIGNAL 01 | LIGHT SOURCE

Thin-Film LiNbO3: UMC and HyperLight volume production specs

Thin-film lithium niobate (TFLN) technology is rapidly advancing, with HyperLight collaborating with UMC and Wavetek to accelerate volume production of its modulators.

SIGNAL 02 | PHOTONIC DEVICE

Broadcom Baily Switch: 51.2Tb CPO power reduction specs

Broadcom has commenced mass production of its 51.2 Terabit (Tb) Baily switch, achieving up to 70% power consumption reduction through Co-Packaged Optics (CPO) compared to traditional pluggable transceivers.

SIGNAL 03 | PACKAGING

Texas Instruments: 400G to 1.6T optical transceiver roadmap

Texas Instruments (TI) emphasizes that optical communication is driving the future of data center networking, a rapidly growing investment area.

SIGNAL 04 | NETWORK SYSTEM

TFLN Modulators: Murata's CPO roadmap for 2028 production

Murata Manufacturing has unveiled its Co-Packaged Optics (CPO) technology roadmap for AI data centers, centering on high-performance Thin-Film Lithium Niobate (TFLN) optical modulators.

DECISION

Focus this week's management attention on link budget, packaging, thermal reliability and deployable supply. Move only when the linked evidence can be reproduced under your own operating conditions.

THREE MANAGEMENT MOVES

1. Buyers: evaluate optical architectures by reach, power, latency, interoperability and field service.
2. Suppliers: pair optical performance with packaging yield, test, thermal drift and qualified capacity.
3. Executives: track yield, coupling loss, temperature drift, interoperability, volume and cost.

EVIDENCE DISCIPLINE

FACT states what the collected English source reports. BUSINESS READ and ACTION are editorial interpretations. UNKNOWN lists information that the source file does not establish. A linked article is not treated as independent confirmation of every claim.

Five numbers or evidence anchors that require conditions

70% A06	SOURCE-BOUND CONDITION Broadcom has commenced mass production of its 51.2 Terabit (Tb) Baily switch, achieving up to 70% power consumption reduction through Co-Packaged Optics (CPO) compared to traditional pluggable transceivers.
400G A11	SOURCE-BOUND CONDITION Optical transceivers have seen a fourfold increase in performance, evolving from 400G designs to 1.6T, with projections for a future scale-up to 12.8T (12.8 trillion bits per second).
800G A05	SOURCE-BOUND CONDITION The OIF announced its 800ZR Coherent Implementation Agreement in October 2024, highlighting the hidden complexities in the transition to 400G and 800G infrastructure driven by AI-led network upgrades.
1.3 kg A02	SOURCE-BOUND CONDITION Utilizing their NXTR SMT pick and place machines, Fuji demonstrated precise handling and placement of the approximately 1.3 kg component, which integrates a switch ASIC and optical engines within a single package.
256 G A08	SOURCE-BOUND CONDITION At OFC 2026, NVIDIA announced a 256 Gb/s DWDM optical I/O, achieved through the intimate 3D stacked packaging of Electronic ICs (EIC) and Photonic ICs (PIC), delivering an ultra-low power consumption of 2.6 pJ/bit.

HOW TO USE THESE NUMBERS

Each value is shown with the condition stated in the collected English article. Do not compare values across different test methods, device sizes, operating temperatures or commercial stages without normalizing those conditions.

Where the value chain moved

Each stage combines one leading theme and one supporting theme. The report treats handoffs as evidence transfers, not decorative arrows.

1

Light source

P01/P06

Thin-film lithium niobate (TFLN) technology is rapidly advancing, with HyperLight collaborating with UMC and Wavetek to.

2

Photonic device

P02/P07

Broadcom has commenced mass production of its 51.2 Terabit (Tb) Bailly switch, achieving up to 70% power consumption.

3

Packaging

P03/P08

Texas Instruments (TI) emphasizes that optical communication is driving the future of data center networking, a rapidly.

4

Network system

P04/P09

Murata Manufacturing has unveiled its Co-Packaged Optics (CPO) technology roadmap for AI data centers, centering on.

5

Operations

P05/P10

The OIF announced its 800ZR Coherent Implementation Agreement in October 2024, highlighting the hidden complexities in the.

EVIDENCE REQUIRED AT EACH HANDOFF

Light source to Photonic device	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck yield, coupling loss, temperature drift, interoperability, volume and cost.
Photonic device to Packaging	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck yield, coupling loss, temperature drift, interoperability, volume and cost.
Packaging to Network system	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck yield, coupling loss, temperature drift, interoperability, volume and cost.
Network system to Operations	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck yield, coupling loss, temperature drift, interoperability, volume and cost.

COMMON CONTROL POINT

No theme moves to a commercial commitment until a named owner has reproduced the relevant evidence and documented yield, coupling loss, temperature drift, interoperability, volume and cost.

Priority map: maturity and next decision locus



REPRODUCIBLE POSITIONING RULE

Horizontal position uses five discrete stages: Research, Prototype, Joint evaluation, Product adoption and Scale. The vertical bands classify the next decision locus: Tech / process, System / adoption, or Business / rules (including supply systems). Article-title cues take precedence; the overview is used only when the title is silent. These are categories, not impact ranks. Bubble diameter reflects the editorial score.

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P01

Thin-Film LiNbO3: UMC and HyperLight volume production specs

A12 | Light source | Product adoption

Date not stated

FACT

Thin-film lithium niobate (TFLN) technology is rapidly advancing, with HyperLight collaborating with UMC and Wavetek to accelerate volume production of its modulators. Concurrently, Huazhong University of Science and Technology's research team achieved a new record for broad-spectrum coverage with a single TFLN modulator. While TFLN demonstrates superior bandwidth, thermal independence, and low power, significant industrial volume production challenges persist despite these technological victories.

WHY IT MATTERS

The decision relevance is link budget, packaging, thermal reliability and deployable supply. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, pair optical performance with packaging yield, test, thermal drift and qualified capacity.

ACTION

Within 30 days, evaluate optical architectures by reach, power, latency, interoperability and field service; record the evidence and owner against P01.

UNKNOWN

Still unproven or incompletely stated: yield, coupling loss, temperature drift, interoperability, volume and cost.

EDITORIAL SCORE 21/25 | MATURITY 4/5 | DECISION LOCUS: TECH / PROCESS

P02

Broadcom Bailly Switch: 51.2Tb CPO power reduction specs

A06 | Photonic device | Scale-up

Date not stated

FACT

Broadcom has commenced mass production of its 51.2 Terabit (Tb) Bailly switch, achieving up to 70% power consumption reduction through Co-Packaged Optics (CPO) compared to traditional pluggable transceivers. Concurrently, NVIDIA has begun shipping its Spectrum-X switches, leveraging TSMC's advanced packaging technology, to deliver up to 400 Tbps switching capacity. Both companies are at the forefront of AI data center upgrades, deploying CPO strategies to lead the industry.

WHY IT MATTERS

The decision relevance is link budget, packaging, thermal reliability and deployable supply. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, pair optical performance with packaging yield, test, thermal drift and qualified capacity.

ACTION

Within 30 days, evaluate optical architectures by reach, power, latency, interoperability and field service; record the evidence and owner against P02.

UNKNOWN

Still unproven or incompletely stated: yield, coupling loss, temperature drift, interoperability, volume and cost.

EDITORIAL SCORE 21/25 | MATURITY 5/5 | DECISION LOCUS: BUSINESS / RULES

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P03

Texas Instruments: 400G to 1.6T optical transceiver roadmap

A11 | Packaging | Product adoption

Date not stated

FACT

Texas Instruments (TI) emphasizes that optical communication is driving the future of data center networking, a rapidly growing investment area. Optical transceivers have seen a fourfold increase in performance, evolving from 400G designs to 1.6T, with projections for a future scale-up to 12.8T (12.8 trillion bits per second).

WHY IT MATTERS

The decision relevance is link budget, packaging, thermal reliability and deployable supply. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, pair optical performance with packaging yield, test, thermal drift and qualified capacity.

ACTION

Within 30 days, evaluate optical architectures by reach, power, latency, interoperability and field service; record the evidence and owner against P03.

UNKNOWN

Still unproven or incompletely stated: yield, coupling loss, temperature drift, interoperability, volume and cost.

EDITORIAL SCORE 20/25 | MATURITY 4/5 | DECISION LOCUS: BUSINESS / RULES

P04

TFLN Modulators: Murata's CPO roadmap for 2028 production

A04 | Network system | Scale-up

Date not stated

FACT

Murata Manufacturing has unveiled its Co-Packaged Optics (CPO) technology roadmap for AI data centers, centering on high-performance Thin-Film Lithium Niobate (TFLN) optical modulators. The company plans to complete the first TFLN modulator samples by Q4 2026, aiming for mass production in 2028.

WHY IT MATTERS

The decision relevance is link budget, packaging, thermal reliability and deployable supply. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, pair optical performance with packaging yield, test, thermal drift and qualified capacity.

ACTION

Within 30 days, evaluate optical architectures by reach, power, latency, interoperability and field service; record the evidence and owner against P04.

UNKNOWN

Still unproven or incompletely stated: yield, coupling loss, temperature drift, interoperability, volume and cost.

EDITORIAL SCORE 20/25 | MATURITY 5/5 | DECISION LOCUS: SYSTEM / ADOPTION

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P05

800ZR: OIF's 2024 coherent implementation agreement specs

A05 | Operations | Product adoption

Date not stated

FACT

The OIF announced its 800ZR Coherent Implementation Agreement in October 2024, highlighting the hidden complexities in the transition to 400G and 800G infrastructure driven by AI-led network upgrades. This agreement defines the 800ZR coherent interface for single-span amplified DWDM links used in Data Center Interconnect (DCI).

WHY IT MATTERS

The decision relevance is link budget, packaging, thermal reliability and deployable supply. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, pair optical performance with packaging yield, test, thermal drift and qualified capacity.

ACTION

Within 30 days, evaluate optical architectures by reach, power, latency, interoperability and field service; record the evidence and owner against P05.

UNKNOWN

Still unproven or incompletely stated: yield, coupling loss, temperature drift, interoperability, volume and cost.

EDITORIAL SCORE 20/25 | MATURITY 4/5 | DECISION LOCUS: BUSINESS / RULES

P06

OFC 2026: Broadcom and Coherent photonic interconnect specs

A01 | Light source | Research

Date not stated

FACT

An OFC 2026 panel highlighted the critical challenge of ultra-low-power photonic scale-up interconnects for AI clusters, showcasing diverse technological approaches. Broadcom demonstrated approximately 1 pJ/bit with its VCSEL Near-Package Optics (NPO) solution, while Coherent achieved an impressive 0.18 pJ/bit of core optoelectronic energy.

WHY IT MATTERS

The decision relevance is link budget, packaging, thermal reliability and deployable supply. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, pair optical performance with packaging yield, test, thermal drift and qualified capacity.

ACTION

Within 30 days, evaluate optical architectures by reach, power, latency, interoperability and field service; record the evidence and owner against P06.

UNKNOWN

Still unproven or incompletely stated: yield, coupling loss, temperature drift, interoperability, volume and cost.

EDITORIAL SCORE 20/25 | MATURITY 1/5 | DECISION LOCUS: SYSTEM / ADOPTION

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P07

Tomahawk 6: Fuji's first SMT placement of Broadcom CPO

A02 | Photonic device | Pilot

Date not stated

FACT

Fuji Corporation has successfully performed the industry's first SMT placement evaluation of Broadcom's Tomahawk® 6 - Davisson Co-Packaged Optics (CPO) Ethernet switch. Utilizing their NXTR SMT pick and place machines, Fuji demonstrated precise handling and placement of the approximately 1.3 kg component, which integrates a switch ASIC and optical engines within a single package.

WHY IT MATTERS

The decision relevance is link budget, packaging, thermal reliability and deployable supply. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, pair optical performance with packaging yield, test, thermal drift and qualified capacity.

ACTION

Within 30 days, evaluate optical architectures by reach, power, latency, interoperability and field service; record the evidence and owner against P07.

UNKNOWN

Still unproven or incompletely stated: yield, coupling loss, temperature drift, interoperability, volume and cost.

EDITORIAL SCORE 19/25 | MATURITY 2/5 | DECISION LOCUS: SYSTEM / ADOPTION

P08

Intel OCI Chiplet: AI infrastructure power specs for 2026

A07 | Packaging | Research

Date not stated

FACT

At OFC 2026, Intel introduced its Optical Compute Interconnect (OCI) chiplet, deemed essential for AI scale-up infrastructure. This OCI chiplet demonstrates remarkable energy efficiency of approximately 5 pJ/b and a shoreline density of 1 Tbps/mm, translating to over 70% power reduction compared to conventional 800G pluggable modules (approx. 21 pJ/b).

WHY IT MATTERS

The decision relevance is link budget, packaging, thermal reliability and deployable supply. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, pair optical performance with packaging yield, test, thermal drift and qualified capacity.

ACTION

Within 30 days, evaluate optical architectures by reach, power, latency, interoperability and field service; record the evidence and owner against P08.

UNKNOWN

Still unproven or incompletely stated: yield, coupling loss, temperature drift, interoperability, volume and cost.

EDITORIAL SCORE 19/25 | MATURITY 1/5 | DECISION LOCUS: BUSINESS / RULES

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P09

NVIDIA Optical I/O: 256 Gb/s DWDM silicon photonics specs

A08 | Network system | Research

Date not stated

FACT

At OFC 2026, NVIDIA announced a 256 Gb/s DWDM optical I/O, achieved through the intimate 3D stacked packaging of Electronic ICs (EIC) and Photonic ICs (PIC), delivering an ultra-low power consumption of 2.6 pJ/bit. Based on TSMC's COUPE platform and utilizing Cu-Cu hybrid bonding (SoIC), this technology significantly reduces parasitic capacitance, leading to an impressive 30% to 40% reduction in laser power requirements.

WHY IT MATTERS

The decision relevance is link budget, packaging, thermal reliability and deployable supply. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, pair optical performance with packaging yield, test, thermal drift and qualified capacity.

ACTION

Within 30 days, evaluate optical architectures by reach, power, latency, interoperability and field service; record the evidence and owner against P09.

UNKNOWN

Still unproven or incompletely stated: yield, coupling loss, temperature drift, interoperability, volume and cost.

EDITORIAL SCORE 19/25 | MATURITY 1/5 | DECISION LOCUS: SYSTEM / ADOPTION

P10

Frequency Combs: Optica's 2026 AI optical I/O roadmap

A03 | Operations | Research

Date not stated

FACT

Alexander Gaeta's article in Optica details groundbreaking advancements in commercializing frequency combs for integrated photonics. This technology allows for the generation of low-power, compact microcombs compatible with existing photonic foundry materials, paving the way for large-scale deployment.

WHY IT MATTERS

The decision relevance is link budget, packaging, thermal reliability and deployable supply. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, pair optical performance with packaging yield, test, thermal drift and qualified capacity.

ACTION

Within 30 days, evaluate optical architectures by reach, power, latency, interoperability and field service; record the evidence and owner against P10.

UNKNOWN

Still unproven or incompletely stated: yield, coupling loss, temperature drift, interoperability, volume and cost.

EDITORIAL SCORE 19/25 | MATURITY 1/5 | DECISION LOCUS: SYSTEM / ADOPTION

Playbook and 0-5 year roadmap

STAKEHOLDER	NOW 0-30 DAYS	NEXT 1-12 MONTHS	LATER 1-5 YEARS
Operating companies	evaluate optical architectures by reach, power, latency, interoperability and field service	Run production-like qualification with explicit acceptance and stop criteria.	Build a portfolio and architecture that can absorb technology and supplier changes.
Materials, equipment and technology suppliers	pair optical performance with packaging yield, test, thermal drift and qualified capacity	Create shared reliability, process-window and failure-analysis data with customers.	Standardize interfaces, traceability, lifecycle support and recovery services.
Trading, investment and legal teams	Map yield, coupling loss, temperature drift, interoperability, volume and cost.	Contract for capacity, quality, change notice, liability, alternatives and exit.	Diversify regional, technical and customer concentration risk.

ROADMAP

NOW 0-30 DAYS Evidence ledger Owner: Strategy + quality Deliverable: claim, source, condition, owner	NOW 0-30 DAYS Risk screen Owner: Procurement + legal Deliverable: unknowns, alternatives, stop	NEXT 1-12 MONTHS Joint qualification Owner: R&D + supplier Deliverable: process window and failure data
NEXT 1-12 MONTHS Commercial gate Owner: Business + finance Deliverable: unit economics and capacity	LATER 1-5 YEARS Platform strategy Owner: Executive team Deliverable: portfolio and interface roadmap	LATER 1-5 YEARS Service layer Owner: Operations + channel Deliverable: monitoring, maintenance and recovery

GO / NO-GO GATES

GATE	OWNER	REQUIRED EVIDENCE	NO-GO CONDITION
G1 Evidence	Quality	Source, test conditions and reproducibility	No traceable evidence
G2 Integration	R&D	Interface, process window and failure analysis	Cannot meet system conditions
G3 Commercial	Business	Capacity, total cost, contract and alternatives	Economics or supply cannot be supported
G4 Lifecycle	Operations	Monitoring, maintenance, change notice and exit	No accountable operating plan

ANALYZED ARTICLES | ITEMS 01-12

Every title and URL is displayed in full. URLs wrap without shortening and remain clickable. A verified-reference label means the translated HTML did not retain its original source URL.

ID	FULL ARTICLE TITLE	FULL SOURCE / VERIFIED REFERENCE URL - CLICKABLE
A01	OFC 2026: Broadcom and Coherent photonic interconnect specs	https://www.simpletechtrend.com/en/post/ofc2026-ai-scale-up-cpo-vcs-el-photonics-analysis-en
A02	Tomahawk 6: Fujii's first SMT placement of Broadcom CPO	https://smttoday.com/2026/10/07/fuji-achieves-industry-first-placement-of-broadcoms-tomahawk-6-davisson/
A03	Frequency Combs: Optica's 2026 AI optical I/O roadmap	https://www.optica-opn.org/home/articles/volume_37/october_2026/departments/frequency_combs_for_integrated_photonics/
A04	TFLN Modulators: Murata's CPO roadmap for 2028 production	https://www.thelec.net/news/articleView.html?idxno=14362
A05	800ZR: OIF's 2024 coherent implementation agreement specs	https://www.digitaljournal.com/article/ai-driven-network-upgrades-expose-hidden-complexity-in-the-shift-to-400g-and-800g-infrastructure/
A06	Broadcom Bailly Switch: 51.2Tb CPO power reduction specs	https://www.simpletechtrend.com/en/post/cpo_nvidia_vs_broadcom-en
A07	Intel OCI Chiplet: AI infrastructure power specs for 2026	https://www.simpletechtrend.com/en/post/ofc2026-intel-oci-siph-chiplet-scaling-en
A08	NVIDIA Optical I/O: 256 Gb/s DWDM silicon photonics specs	https://www.simpletechtrend.com/en/post/ofc2026-nvidia-dwdm-optical-io-3d-stacked-en
A09	Intel CPO: InP on Silicon integrated laser roadmap	https://intellionaire.substack.com/p/the-intellionaire-ep-43-intel-and
A10	NTT & Chunghwa Telecom: 100Gbps All-Photonics Network specs	https://en.bloguru.com/ken/569294/the-future-of-data-transfer-is-light
A11	Texas Instruments: 400G to 1.6T optical transceiver roadmap	https://www.ti.com/about-ti/behind-chip/articles/whats-next-for-the-future-of-data-center-networking-optics.html
A12	Thin-Film LiNbO3: UMC and HyperLight volume production specs	https://www.simpletechtrend.com/en/post/tfln-optical-cpo-opportunity-walls-2026-en

Photonics — Selected Articles

Date: 2026-10-11

Articles: 12

Table of Contents

- #01 OFC 2026: Broadcom and Coherent photonic interconnect specs
- #02 Tomahawk 6: Fuji's first SMT placement of Broadcom CPO
- #03 Frequency Combs: Optica's 2026 AI optical I/O roadmap
- #04 TFLN Modulators: Murata's CPO roadmap for 2028 production
- #05 800ZR: OIF's 2024 coherent implementation agreement specs
- #06 Broadcom Bailly Switch: 51.2Tb CPO power reduction specs
- #07 Intel OCI Chiplet: AI infrastructure power specs for 2026
- #08 NVIDIA Optical I/O: 256 Gb/s DWDM silicon photonics specs
- #09 Intel CPO: InP on Silicon integrated laser roadmap
- #10 NTT & Chunghwa Telecom: 100Gbps All-Photonics Network specs
- #11 Texas Instruments: 400G to 1.6T optical transceiver roadmap
- #12 Thin-Film LiNbO3: UMC and HyperLight volume production specs

#01 OFC 2026: Broadcom and Coherent photonic interconnect specs

Published October 06, 2026 Simple Tech Trend (OFC Panel Report) USA



OVERVIEW

An OFC 2026 panel highlighted the critical challenge of ultra-low-power photonic scale-up interconnects for AI clusters, showcasing diverse technological approaches. Broadcom demonstrated approximately 1 pJ/bit with its VCSEL Near-Package Optics (NPO) solution, while Coherent achieved an impressive 0.18 pJ/bit of core optoelectronic energy. Marvell presented a more comprehensive silicon photonics Co-Packaged Optics (CPO) solution with analog DSP, delivering 3 pJ/bit for the total link via WDM, signaling an intense race to optimize power efficiency and bandwidth for next-generation AI infrastructure.

Key Findings

At the OFC 2026 panel, leading industry players presented their latest advancements in achieving ultra-low-power photonic scale-up interconnects, a critical requirement for next-generation AI data centers. Broadcom's VCSEL Near-Package Optics (NPO) solution attained approximately 1 pJ/bit, while Coherent demonstrated an even lower 0.18 pJ/bit of core optoelectronic energy, setting new benchmarks for efficiency. Marvell's silicon photonics Co-Packaged Optics (CPO) approach, integrating analog DSP, achieved 3 pJ/bit for the total link, providing a more complete solution for bandwidth scaling through WDM.

Technical Details

- **Broadcom's VCSEL NPO:** Broadcom's strategy focuses on 100G/200G VCSELs due to their cost-effectiveness and manufacturability. Their NPO implementation delivered an energy efficiency of approximately 1 pJ/bit, positioning VCSELs as a strong contender for raw power efficiency in certain applications.
- **Coherent's Optoelectronic Energy:** Coherent showcased a remarkable 0.18 pJ/bit in core optoelectronic energy. This figure represents the energy consumption within the core optical-to-electrical conversion components, indicating highly optimized device-level performance.
- **Marvell's Silicon Photonics CPO:** Marvell's solution leverages silicon photonics with external lasers and analog DSP, offering a scalable approach via WDM. While its total link energy consumption was 3 pJ/bit, this includes the full link budget, providing a practical solution for deploying high-bandwidth AI clusters, particularly post-2027.
- **NVIDIA's Direction:** NVIDIA is also actively pursuing high-bandwidth-density silicon photonics CPO with external lasers and analog DSP for its top-tier AI clusters, aligning with Marvell's vision for future AI networking architectures.

Background & Context

The explosive growth of AI workloads demands unprecedented bandwidth and ultra-low power consumption for interconnects within data centers. As AI clusters scale to thousands of accelerators, even marginal improvements in pJ/bit energy efficiency can translate into significant operational cost savings and reduced environmental impact. The discussion at OFC 2026 highlighted a divergence in technological roadmaps, with some favoring VCSELs for their established manufacturing and cost advantages, while others champion silicon photonics CPO for its integration density and scalability for advanced AI fabrics. The evolving AI infrastructure increasingly seeks open, multi-vendor Scale-Up fabrics, creating a layered coexistence model where proprietary and open Ethernet-based technologies compete and complement each other.


Strategic Significance & Outlook

The panel discussion underscores the intense innovation in photonic interconnects, driven by the insatiable demands of AI. The competing technologies—VCSEL-based NPO versus silicon photonics CPO with external lasers and analog DSP—each offer distinct advantages. VCSELs excel in raw power efficiency and cost for shorter reaches, while silicon photonics CPO promises higher integration density and bandwidth for larger, more complex AI clusters. The industry is moving towards a future where hybrid solutions and application-specific optimizations will likely dictate market adoption. The continuous push for sub-1 pJ/bit interconnects will be pivotal in enabling the next generation of AI supercomputers, impacting everything from data center design to sustainable computing practices.

Source: <https://www.simpletechtrend.com/en/post/ofc2026-ai-scale-up-cpo-vcSEL-photonics-analysis-en>

#02 Tomahawk 6: Fuji's first SMT placement of Broadcom CPO

Published October 07, 2026 Davisson (Fuji Corporation Press Release) Japan



OVERVIEW

Fuji Corporation has successfully performed the industry's first SMT placement evaluation of Broadcom's Tomahawk® 6 – Davisson Co-Packaged Optics (CPO) Ethernet switch. Utilizing their NXTR SMT pick and place machines, Fuji demonstrated precise handling and placement of the approximately 1.3 kg component, which integrates a switch ASIC and optical engines within a single package. This breakthrough is critical for advancing high-density optical I/O solutions that address the escalating performance and power efficiency demands in AI data center networks.

Key Findings

Fuji Corporation has achieved an industry first by successfully demonstrating Surface Mount Technology (SMT) placement of Broadcom's Tomahawk® 6 – Davisson Co-Packaged Optics (CPO) Ethernet switch. This significant milestone was accomplished using Fuji's advanced NXTR SMT pick and place machines, proving the capability to precisely handle and integrate this substantial component, weighing approximately 1.3 kg, onto circuit boards. The Tomahawk® 6 – Davisson is engineered to meet the demanding requirements for higher performance and lower power consumption in AI data center networks by integrating a switch ASIC and optical engines into a single package.

Technical Details

The Tomahawk® 6 – Davisson is a next-generation CPO solution designed to overcome the bandwidth and power limitations of traditional electrical interconnects in AI data centers. By co-packaging the switch ASIC and optical engines, it significantly shortens electrical trace lengths, leading to improved signal integrity and reduced power loss. The component's considerable weight of approximately 1.3 kg presented a unique challenge for conventional SMT equipment. Fuji's NXTR SMT pick and place machines, known for their high precision and capability to handle heavy and irregularly shaped components, were instrumental in achieving this successful placement. This demonstrates a critical manufacturing advancement for high-density optical interconnects, enabling their scalable production via SMT.

Background & Context

The explosion of AI workloads has led to an exponential increase in data transfer requirements between GPUs and CPUs within AI data centers. Traditional electrical interconnects are struggling to keep pace with these demands in terms of bandwidth, latency, and power consumption. Co-Packaged Optics (CPO) technology is widely recognized as a pivotal solution to these challenges, enabling significant performance boosts and power efficiency improvements by closely integrating optical engines with ASICs. However, the industrialization of CPO has been hampered by manufacturing complexities, particularly the reliable and cost-effective mass production of such large and sophisticated components. Fuji's successful SMT placement represents a major leap forward in overcoming these manufacturing hurdles, paving the way for broader adoption of CPO technology in AI infrastructure.


Strategic Significance & Outlook

Fuji's successful SMT placement of the Broadcom Tomahawk® 6 – Davisson is a critical enabler for the mass production of CPO technology, with profound implications for future AI data center designs. This achievement is expected to accelerate the adoption of CPO by a wider range of companies, driving the development of even denser and more power-efficient network equipment. Beyond AI, this technology holds promise for high-performance computing (HPC), cloud services, and any application demanding high bandwidth and low latency. The evolution of SMT technology to accommodate complex CPO components will be key to reducing manufacturing costs and facilitating widespread deployment, thereby shaping the future of optical communication infrastructure.

Source: <https://smttoday.com/2026/10/07/fuji-achieves-industry-first-placement-of-broadcoms-tomahawk-6-davisson/>

#03 Frequency Combs: Optica's 2026 AI optical I/O roadmap

Published October 01, 2026 Optica USA



OVERVIEW

Alexander Gaeta's article in Optica details groundbreaking advancements in commercializing frequency combs for integrated photonics. This technology allows for the generation of low-power, compact microcombs compatible with existing photonic foundry materials, paving the way for large-scale deployment. Crucially, it enables the generation of multiple wavelengths from a single laser, integrated as Co-Packaged Optics (CPO) on a chip, promising to revolutionize high-bandwidth communication between GPUs and CPUs in AI clusters.

Key Findings

An article in *Optica* by Alexander Gaeta highlights significant advancements towards the commercialization of frequency comb technology for integrated photonics. This innovation enables the creation of low-power, compact, and lightweight microcombs that are compatible with existing photonic foundry materials, suggesting a clear path to large-scale adoption. The most pivotal breakthrough is the ability to generate multiple wavelengths from a single laser source, which can then be integrated as Co-Packaged Optics (CPO) on a chip. This is expected to fundamentally transform high-bandwidth communication within AI clusters, specifically between GPUs and CPUs.

Technical Details

Frequency combs are specialized lasers that emit a spectrum of light comprising many discrete, equally spaced frequency components, often referred to as 'teeth.' Unlike traditional systems that require multiple lasers for various wavelengths, frequency combs achieve this from a single, compact laser chip. Microcombs, fabricated using standard silicon-on-insulator (SOI) or other photonic foundry materials, are particularly notable for their miniaturization and low power consumption. Gaeta describes this technology as a 'miniature rainbow,' emphasizing its potential to drastically simplify complex optical systems by integrating them onto a chip. This capability is crucial for meeting the demands of AI clusters for ultra-high-bandwidth data transfer, potentially exceeding hundreds of terabits per second, by providing high-density and energy-efficient optical I/O solutions.

Background & Context

The relentless progression of AI and high-performance computing (HPC) has led to an exponential increase in data movement within data centers, pushing the limits of conventional electrical interconnects. Co-Packaged Optics (CPO) and other optical interconnect solutions are seen as key to overcoming this bottleneck, but their realization hinges on the availability of compact, power-efficient, and mass-producible optical components. Frequency comb technology holds immense promise in meeting these requirements, particularly in enhancing the density and reducing the cost of optical modules. Its compatibility with existing semiconductor manufacturing processes streamlines the path to industrial-scale deployment. This technology directly contributes to scaling AI training models and accelerating inference speeds, making it a foundational technology for future digital infrastructure.


Strategic Significance & Outlook

The commercialization of frequency combs in integrated photonics is poised to have a transformative impact across diverse fields, including data communication, sensing, metrology, and quantum computing, extending beyond just AI clusters. Low-power, compact microcombs are expected to dramatically boost optical I/O performance, becoming a critical determinant of power efficiency and scalability in future computing architectures. As Gaeta notes, the ability to leverage existing foundry technologies will expedite time-to-market and accelerate widespread adoption. The evolution of this technology not only shapes the future of optical communication but also fosters the creation of entirely new applications, further solidifying its role as a cornerstone of next-generation technological innovation.

Source: https://www.optica-opn.org/home/articles/volume_37/october_2026/departments/frequency_combs_for_integrated_photonics/

#04 TFLN Modulators: Murata's CPO roadmap for 2028 production

Published October 01, 2026 The Elec Japan



OVERVIEW

Murata Manufacturing has unveiled its Co-Packaged Optics (CPO) technology roadmap for AI data centers, centering on high-performance Thin-Film Lithium Niobate (TFLN) optical modulators. The company plans to complete the first TFLN modulator samples by Q4 2026, aiming for mass production in 2028. These modulators are designed to achieve an electro-optical (EO) bandwidth exceeding 100GHz and transmission speeds over 400Gbps per channel, enabling broad applications beyond CPO, including coherent communication, WDM, and quantum communication.

Key Findings

Murata Manufacturing has publicly released its technology roadmap for Co-Packaged Optics (CPO), designed to meet the escalating demands of AI data centers. Central to this roadmap is a groundbreaking optical modulator utilizing Thin-Film Lithium Niobate (TFLN) technology. The first samples of this TFLN modulator are slated for completion by the fourth quarter of 2026, with mass production targeted for 2028, signifying a major leap forward in next-generation optical communication performance.

Technical / Clinical Details

Murata's TFLN optical modulator is engineered to achieve an electro-optical (EO) bandwidth exceeding 100GHz and ultra-high transmission speeds of over 400Gbps per channel. Lithium niobate is intrinsically a material with excellent electro-optical properties, but its thin-film form factor facilitates hybrid integration with silicon photonics while simultaneously reducing power consumption and miniaturizing the device. This high bandwidth and speed are critically important for alleviating data bottlenecks between CPU/GPU and optical modules within AI data center CPO configurations. Beyond CPO, this technology is also being considered for a wide range of advanced applications, including coherent communication for long-haul, high-capacity transmission, Wavelength Division Multiplexing (WDM) for simultaneous multi-wavelength transmission, and future quantum communication systems.

Background & Context

The explosive growth in power consumption and data processing volumes in AI data centers has exposed the limitations of conventional pluggable optical transceivers. CPO, as a next-generation integration technology, dramatically reduces signal loss and power consumption by shortening electrical interconnects through placing optical components in close proximity to, or within the same package as, electrical chips (CPUs/GPUs). Murata's integration of TFLN technology into its CPO roadmap underscores its belief that this technology will be an indispensable component of AI-era data center infrastructure. The high-precision component manufacturing and novel material development capabilities of Japanese companies, in particular, are key to establishing competitive advantages in this field.

Strategic Significance & Outlook

The mass production of Murata's TFLN optical modulators will significantly accelerate the practical implementation of CPO technology, elevating the performance and energy efficiency of AI data centers to new levels. This technology is expected to have a profound impact across the optical communication industry, not only enabling higher performance in existing technologies like coherent communication and WDM but also contributing to the realization of optical interconnects in nascent areas such as quantum computing and quantum networks. The targeted mass production date of 2028 represents a critical milestone that will shape the future direction of optical communication technology.

Source: <https://www.thelec.net/news/articleView.html?idxno=14362>

Collected: October 09, 2026 | Automated Research System (Gemini API)

#05 800ZR: OIF's 2024 coherent implementation agreement specs

Published October 03, 2026 Digital Journal USA



OVERVIEW

The OIF announced its 800ZR Coherent Implementation Agreement in October 2024, highlighting the hidden complexities in the transition to 400G and 800G infrastructure driven by AI-led network upgrades. This agreement defines the 800ZR coherent interface for single-span amplified DWDM links used in Data Center Interconnect (DCI). Coherent technology provides an efficient method to transmit high-capacity signals over amplified optical systems, particularly for connecting data centers over metro or long distances, resolving bandwidth and efficiency challenges.

Key Findings

In October 2024, the Optical Internetworking Forum (OIF) announced its 800ZR Coherent Implementation Agreement, detailing the latent complexities in the transition to 400G and 800G infrastructure amidst AI-driven network upgrades. This agreement establishes the technical specifications for the 800ZR coherent interface, intended for single-span amplified DWDM (Dense Wavelength Division Multiplexing) links utilized in Data Center Interconnect (DCI). Coherent technology offers the capability to efficiently transmit ultra-high-capacity signals over amplified optical systems, particularly for connecting data centers across metropolitan or long distances, playing an indispensable role in expanding bandwidth and improving operational efficiency.

Technical Details

The OIF's 800ZR Coherent Implementation Agreement defines the key requirements for realizing an 800Gb/s single-wavelength coherent optical interface in a pluggable form factor. This technology is designed to efficiently handle massive traffic within data centers and enable connectivity between data centers over distances up to approximately 120km. Coherent optical transmission, by employing advanced encoding and decoding of light signal amplitude, phase, and polarization, offers superior noise immunity and spectral efficiency compared to traditional direct detection methods. This allows for higher bitrate transmission while maximizing the utilization of existing optical fiber infrastructure. When combined with DWDM, multiple 800Gb/s wavelengths can be multiplexed onto a single fiber pair, dramatically increasing network capacity. This standardization aims to ensure multi-vendor interoperability and foster competition in the optical transceiver market.

Background & Context

The explosive growth of AI/ML workloads and the widespread adoption of cloud computing have pushed data center traffic to unprecedented levels. This has led to exceptionally high demands for bandwidth and energy efficiency in Data Center Interconnects (DCI). While upgrading networks to 400G and 800G is essential, its implementation presents challenges such as compatibility with existing infrastructure, power consumption, thermal management, and complex network design. Coherent technology is recognized as the most efficient and scalable solution to these challenges, and implementation agreements promoted by standardization bodies like OIF are crucial for accelerating industry-wide unification and innovation.


Strategic Significance & Outlook

The OIF's 800ZR Implementation Agreement marks a significant milestone in the optical communications industry. This standardization will accelerate the adoption of 800Gb/s coherent optical technology, enabling data center operators and communication providers to meet the immense data demands of the AI era. Ensuring interoperability enhances supply chain flexibility and promotes cost competition, thereby lowering barriers to technology adoption. In the future, coherent technologies like 800ZR are expected to evolve to even higher speeds (e.g., 1.6T), forming the foundation for optical networks to sustainably support the exponential growth of data traffic. This will powerfully drive the development of the global digital economy and enable new technological innovations.

Source: <https://www.digitaljournal.com/article/ai-driven-network-upgrades-expose-hidden-complexity-in-the-shift-to-400g-and-800g-infrastructure/>

#06 Broadcom Bailly Switch: 51.2Tb CPO power reduction specs

Published October 05, 2026 Simple Tech Trend USA



OVERVIEW

Broadcom has commenced mass production of its 51.2 Terabit (Tb) Bailly switch, achieving up to 70% power consumption reduction through Co-Packaged Optics (CPO) compared to traditional pluggable transceivers. Concurrently, NVIDIA has begun shipping its Spectrum-X switches, leveraging TSMC's advanced packaging technology, to deliver up to 400 Tbps switching capacity. Both companies are at the forefront of AI data center upgrades, deploying CPO strategies to lead the industry.

Key Findings

In the competitive landscape of AI data center performance enhancement, Broadcom and NVIDIA have unveiled groundbreaking products, each deploying aggressive Co-Packaged Optics (CPO) strategies. Broadcom has initiated mass production of its 51.2 Terabit (Tb) Baily switch, which integrates CPO to achieve an impressive power consumption reduction of up to 70% compared to systems using traditional pluggable transceivers. Meanwhile, NVIDIA has begun shipping its Spectrum-X switches, utilizing TSMC's advanced packaging technology, to provide an extraordinary 400 Tbps switching capacity, directly addressing the massive bandwidth requirements of AI infrastructure.

Technical / Clinical Details

Broadcom's 51.2Tb Baily switch dramatically improves power efficiency by combining industry-leading switching capacity with CPO. CPO is a technology that tightly integrates electrical chips and optical engines within the same package, minimizing electrical signal paths. This integration significantly reduces signal degradation and power consumption. The 70% power reduction has profound implications for the operational costs and environmental footprint of large-scale data centers. NVIDIA's Spectrum-X switch, with its 400 Tbps switching capacity, efficiently manages the immense data flows characteristic of AI workloads. NVIDIA's strategic adoption of advanced packaging technologies, such as TSMC's Co-packaged Optics Using an Embedded Die (COUPE), is key to achieving both high performance and high density. This switch provides essential infrastructure for AI training clusters, particularly those requiring high-speed interconnects between GPUs, thereby eliminating bottlenecks and maximizing computational power.

Background & Context

The rapid advancement of AI necessitates a revolution in data center architecture. Traditional data centers are encountering limitations in terms of power consumption, cooling, and physical space. AI workloads, in particular, demand vast data processing capabilities and ultra-fast communication between GPUs. Against this backdrop, innovative optical interconnect technologies like CPO are viewed as crucial for breaking through the 'power wall' and 'I/O wall' in the AI era. Broadcom has long been a leader in the network switch market, and its CPO strategy represents a significant leap forward for next-generation data centers. NVIDIA, as a premier GPU vendor, enhances the entire AI ecosystem by offering networking solutions that seamlessly integrate with its computing platforms. The actions of both companies demonstrate that competition in the AI infrastructure market is shifting beyond mere chip performance to encompass integrated optical communication technologies.


Strategic Significance & Outlook

The aggressive adoption of CPO technology by Broadcom and NVIDIA is profoundly shaping the future of AI data centers. CPO will enhance data center power efficiency and scalability, enabling the deployment of larger and more complex AI models. This will, in turn, accelerate innovation across various fields, including generative AI, autonomous driving, and scientific research. However, the widespread adoption of CPO also entails challenges such as supply chain maturation, standardization efforts, and the complexity of integration processes. The competition spearheaded by these two industry giants is expected to accelerate the maturation and proliferation of CPO technology, ultimately leading to greater cost efficiency and performance in AI infrastructure. In the future, these technological innovations will form the foundation for building massive computational resources, often referred to as 'AI factories.'

Source: https://www.simpletechtrend.com/en/post/cpo_nvidia_vs_broadcom-en

#07 Intel OCI Chiplet: AI infrastructure power specs for 2026

Published October 05, 2026 Simple Tech Trend USA



OVERVIEW

At OFC 2026, Intel introduced its Optical Compute Interconnect (OCI) chiplet, deemed essential for AI scale-up infrastructure. This OCI chiplet demonstrates remarkable energy efficiency of approximately 5 pJ/b and a shoreline density of 1 Tbps/mm, translating to over 70% power reduction compared to conventional 800G pluggable modules (approx. 21 pJ/b). Intel's innovation is underpinned by its silicon photonics (SiPh) heterogeneous integration technology, directly bonding Indium Phosphide (InP) lasers onto silicon wafers.

Key Findings

At OFC 2026, Intel announced its Optical Compute Interconnect (OCI) chiplet, a critical component for scaling AI infrastructure. This groundbreaking OCI chiplet achieves exceptional energy efficiency of approximately 5 picojoules per bit (pJ/b) and a shoreline density of 1 Terabit per second per millimeter (Tbps/mm). This represents a significant power reduction of over 70% when compared to existing 800G pluggable modules, which typically consume around 21 pJ/b. This achievement is made possible by Intel's unique silicon photonics (SiPh) heterogeneous integration technology, which directly bonds Indium Phosphide (InP) lasers onto silicon wafers.

Technical / Clinical Details

Intel's OCI chiplet is designed to eliminate data transfer bottlenecks between AI accelerators and CPUs within data centers. AI workloads demand the high-speed processing of vast amounts of data, with associated power consumption presenting a major challenge. The OCI chiplet utilizes optical signals for data transmission, achieving performance levels beyond the capabilities of electrical signals. Specifically, the low energy consumption of 5 pJ/b dramatically reduces the overall operational costs and heat generation of data centers. Furthermore, a shoreline density of 1 Tbps/mm signifies extremely high bandwidth density from the chip edge, ensuring maximum data transfer capability within limited physical space. The core of this technology lies in 'heterogeneous integration,' Intel's approach to directly integrating disparate InP lasers onto a standard silicon photonics platform. InP lasers can generate high-quality light that cannot be produced by silicon alone, thereby enabling high-performance and highly efficient optical interconnects through this integration.

Background & Context

The explosive growth of AI is placing unprecedented demands on data center computing and networking infrastructure. Specifically, training large-scale AI models requires numerous processing units (GPUs, AI accelerators, etc.) to collaborate and efficiently exchange petabit-scale data. Traditional electrical wiring and pluggable optical modules are increasingly reaching their limits in terms of power consumption, latency, and physical space. Intel has invested extensively in silicon photonics research and development for many years, and the OCI chiplet is a significant outcome of these efforts. Heterogeneous integration technology, which combines the advantages of different materials to achieve performance unattainable with a single material, is a crucial trend in the semiconductor industry. Intel's approach is positioned as a strategic step to overcome the 'power wall' and 'bandwidth wall' of data centers in the AI era.

Strategic Significance & Outlook


Intel's OCI chiplet holds immense potential to significantly impact the future design and operation of AI data centers. The reduction of over 70% in power consumption will dramatically improve the Total Cost of Ownership (TCO) of AI infrastructure and enable sustainable, large-scale AI computing. The proliferation of this technology will foster the development of more powerful and efficient AI accelerators, accelerating innovation across a wide range of fields, including generative AI, autonomous driving, and scientific computing. Future challenges include further mass production and cost optimization of the heterogeneous integration process, as well as building an ecosystem around the OCI chiplet with other data center components. Intel's leadership in this domain is expected to shape the future of computing in the AI era.

Source: <https://www.simpletechtrend.com/en/post/ofc2026-intel-oci-siph-chiplet-scaling-en>

Collected: October 09, 2026 | Automated Research System (Gemini API)

#08 NVIDIA Optical I/O: 256 Gb/s DWDM silicon photonics specs

Published October 05, 2026 Simple Tech Trend USA



OVERVIEW

At OFC 2026, NVIDIA announced a 256 Gb/s DWDM optical I/O, achieved through the intimate 3D stacked packaging of Electronic ICs (EIC) and Photonic ICs (PIC), delivering an ultra-low power consumption of 2.6 pJ/bit. Based on TSMC's COUPE platform and utilizing Cu-Cu hybrid bonding (SoIC), this technology significantly reduces parasitic capacitance, leading to an impressive 30% to 40% reduction in laser power requirements.

Key Findings

At OFC 2026, NVIDIA unveiled a groundbreaking 256 Gigabits per second (Gb/s) Dense Wavelength Division Multiplexing (DWDM) optical I/O technology tailored for AI infrastructure. This novel technology achieves an astonishingly low power consumption of 2.6 picojoules per bit (pJ/bit) by intimately integrating Electronic Integrated Circuits (EIC) and Photonic Integrated Circuits (PIC) through advanced 3D stack packaging. Furthermore, leveraging TSMC's Co-packaged Optics Using an Embedded Die (COUPE) platform and Cu-Cu hybrid bonding (SoIC) technology, this innovation dramatically reduces parasitic capacitance, resulting in a remarkable 30% to 40% reduction in laser power requirements.

Technical / Clinical Details

As AI systems scale, inter-chip data transfer speed and energy efficiency become paramount. NVIDIA's 256 Gb/s DWDM optical I/O provides a cutting-edge solution to this challenge. DWDM technology allows multiple wavelengths of optical signals to be transmitted simultaneously over a single optical fiber, substantially increasing effective bandwidth. The 3D stack packaging of EIC and PIC dramatically shortens signal paths, enhancing the conversion efficiency between electrical and optical signals. The ultra-low power consumption of 2.6 pJ/bit represents an order-of-magnitude improvement in efficiency compared to conventional optical modules, fundamentally addressing the operational costs and thermal management issues of large-scale AI data centers. TSMC's COUPE platform and SoIC technology enable high-precision chip bonding, optimizing the electrical characteristics of interconnects. This minimizes signal degradation and significantly reduces the power required to drive lasers, thereby maximizing overall system energy efficiency.

Background & Context

The proliferation of AI, particularly large language models (LLMs) and generative AI, is placing unprecedented demands on data transfer between GPUs and AI accelerators in data centers. Conventional electrical wiring and packaging technologies are struggling to meet these requirements, accelerating the shift towards optical interconnects, especially Co-Packaged Optics (CPO) and Near-Packaged Optics (NPO). As a leader in AI hardware, NVIDIA is pioneering innovations in optical interconnect technology, with its close collaboration with TSMC being a critical factor in achieving these breakthroughs. TSMC's advanced packaging technologies are key to overcoming the limits of semiconductor scaling across the industry, and when combined with NVIDIA's expertise, they unlock new possibilities for computing in the AI era.

Strategic Significance & Outlook


NVIDIA's 256 Gb/s DWDM optical I/O technology is poised to revolutionize the future design of AI data centers. The breakthrough low power consumption of 2.6 pJ/bit will dramatically enhance the sustainability and scalability of AI infrastructure, enabling the development and deployment of more powerful and efficient AI models. The widespread adoption of this technology is expected to accelerate innovation in diverse fields such as drug discovery, autonomous driving, and climate modeling. Future challenges include establishing mass production capabilities for this 3D-stacked silicon photonics platform, optimizing cost-efficiency, and ensuring its adoption as an industry standard. The collaboration between NVIDIA and TSMC will play a decisive role in solving one of the most challenging problems facing computing in the AI era and in laying the foundation for next-generation AI factories.

Source: <https://www.simpletechtrend.com/en/post/ofc2026-nvidia-dwdm-optical-io-3d-stacked-en>

Collected: October 09, 2026 | Automated Research System (Gemini API)

#09 Intel CPO: InP on Silicon integrated laser roadmap

Published October 06, 2026 The Intellionaire USA



OVERVIEW

Intel is distinguishing itself in the Co-Packaged Optics (CPO) market as the sole company to design and manufacture its own integrated InP on Silicon lasers. Its Optical Compute Interconnect (OCI) solution leverages Foveros technology for hybrid bonding of Electrical Integrated Circuits (EIC) to Photonic Integrated Circuits (PIC), co-packaged with the XPU via an EMIB bridge. This approach aims to establish a performance, power efficiency, and cost advantage over competitors reliant on external laser sources for their CPO solutions.

Key Findings

Intel is establishing a clear differentiation in the Co-Packaged Optics (CPO) market by incorporating its proprietary integrated laser technology, InP on Silicon, into its CPO solutions. The company stands as the world's only entity capable of designing and manufacturing InP (Indium Phosphide) lasers directly on a silicon platform in-house. This capability ensures a significant technological advantage over competitors who rely on external laser sources, guaranteeing superior performance, supply stability, and cost-efficiency. This unique strategy forms the cornerstone of Intel's Optical Compute Interconnect (OCI) solutions.

Technical / Clinical Details

Intel's OCI solution employs a distinctive architecture that utilizes Foveros technology for hybrid bonding of Electrical Integrated Circuits (EIC) to Photonic Integrated Circuits (PIC). This enables electrical-to-optical signal conversion to occur at the closest possible chip level, minimizing signal latency and power consumption. Furthermore, this EIC-PIC stack is co-packaged with the main XPU (processor, such as CPU/GPU) via an Embedded Multi-die Interconnect Bridge (EMIB). This integrated approach facilitates high-bandwidth, low-power optical interconnects at scale, effectively resolving performance bottlenecks in data centers and AI acceleration. Intel's ability to integrate its own lasers is particularly advantageous for optimizing performance, ensuring supply chain stability, and reducing overall costs.

Background & Context

As AI and High-Performance Computing (HPC) workloads surge, the volume of data transfer between chips has dramatically increased, pushing the limits of traditional electrical interconnects. Co-Packaged Optics (CPO) is seen as the next-generation technology to address these challenges, but many companies currently depend on external laser suppliers. This dependency introduces complexities related to laser module size, cost, power efficiency, and supply chain management. Intel's unique InP on Silicon laser integration has the potential to overcome these obstacles and accelerate the widespread adoption of CPO technology. This strategy is crucial for Intel to maintain its leadership in the data center market and lay the foundation for future computing paradigms.

Strategic Significance & Outlook

Intel's integrated laser CPO technology is poised to become an indispensable component in future AI accelerators, HPC systems, and cloud infrastructure. The company's comprehensive control over the entire optical I/O stack positions it to deliver superior performance and more cost-effective solutions to the market. This is expected to significantly improve data center power efficiency, reduce operational costs, and dramatically enhance AI processing capabilities. Through this technology, Intel aims to establish new standards in optical computing and silicon photonics, driving innovation across a wide range of industries.

Source: <https://intellionaire.substack.com/p/the-intellionaire-ep-43-intel-and>

Collected: October 09, 2026 | Automated Research System (Gemini API)

#10 NTT & Chunghwa Telecom: 100Gbps All-Photonics Network specs

Published October 09, 2026 Japan Talk - Bloguru Japan



OVERVIEW

NTT, in collaboration with Chunghwa Telecom, has successfully established the world's first international All-Photonics Network between Japan and Taiwan. This groundbreaking connection provides a 100 Gbps optical path across approximately 3,000km with an ultra-low one-way latency of just 17 milliseconds. Through its IOWN (Innovative Optical and Wireless Network) initiative, NTT aims to minimize electrical-to-optical conversion, with a future target of achieving 100 times greater power efficiency. This achievement marks a significant milestone towards realizing next-generation networks demanding ultra-low latency and high-capacity data communication.

Key Findings

NTT, in collaboration with Chunghwa Telecom, a leading Taiwanese telecommunications carrier, has successfully constructed the 'world's first international All-Photonics Network' connecting Japan and Taiwan. This innovative network has demonstrated the capability to provide a 100 Gbps optical path over a long distance of approximately 3,000 km with an ultra-low one-way latency of just about 17 milliseconds. This achievement represents concrete progress in NTT's IOWN (Innovative Optical and Wireless Network) concept, which aims to minimize electrical-to-optical signal conversion.

Technical / Clinical Details

The international All-Photonics Network is built upon 'all-photonics' technology, which maintains optical signals throughout the transmission path, minimizing conversion to electrical signals. This dramatically reduces latency and power loss typically associated with electrical-optical conversions in traditional networks. Specifically, for the 3,000 km link between Japan and Taiwan, data transmission latency is suppressed to 17 milliseconds one-way, a groundbreaking low-latency performance for international communications. The 100 Gbps optical path provided more than adequately meets current data communication needs and lays the groundwork for future high-capacity communication. NTT has set an ambitious goal of achieving 100 times greater power efficiency with this technology compared to current networks.

Background & Context

The global digital economy's expansion, particularly driven by AI, IoT, and cloud services, demands communication networks with higher capacity, lower latency, and more efficient data transmission capabilities. Traditional electrical communication infrastructures have faced fundamental physical and technical limitations in meeting these demands. The IOWN concept is NTT's vision for a next-generation network, proposing to overcome information transmission bottlenecks by thoroughly leveraging optical technologies. The establishment of this international APN between Japan and Taiwan is a critical demonstration that the IOWN concept is viable on an international scale, contributing to enhanced digital connectivity in the Asia-Pacific region. This technology also facilitates high-speed inter-data center collaboration, robust communication during disasters, and the creation of new international cooperation models.

Strategic Significance & Outlook

The success of the international All-Photonics Network by NTT and Chunghwa Telecom will significantly influence global communication infrastructure development. This technology has the potential to accelerate the deployment of ultra-low latency, high-capacity networks not only in Asia but also on a global scale. In the future, further advancements in this all-photonics technology are expected to enable international optical paths of several hundred Gbps to Tbps, with even lower power consumption. This will bring closer a future where highly real-time applications such as remote medicine, autonomous driving, and the metaverse can be seamlessly provided across geographical boundaries. NTT is committed to further evolving this technology to bridge the global digital divide and contribute to a sustainable society.

Source: <https://en.bloguru.com/ken/569294/the-future-of-data-transfer-is-light>

#11 Texas Instruments: 400G to 1.6T optical transceiver roadmap

Published October 08, 2026 Texas Instruments USA



OVERVIEW

Texas Instruments (TI) emphasizes that optical communication is driving the future of data center networking, a rapidly growing investment area. Optical transceivers have seen a fourfold increase in performance, evolving from 400G designs to 1.6T, with projections for a future scale-up to 12.8T (12.8 trillion bits per second). This advancement is critical for meeting the escalating demands of AI and machine learning, enabling high-speed, high-capacity data transfer that will revolutionize data center architectures and power efficiency.

Key Findings

Texas Instruments (TI) asserts that optical communication is the definitive driver for the future of data center networking and is significantly accelerating its investments in this sector. Optical transceiver technology has undergone remarkable evolution, with current market designs already achieving a fourfold increase in speed, progressing from 400 gigabits per second (G) to 1.6 terabits per second (T). TI further anticipates the realization of 12.8T (12.8 trillion bits per second) data transfer capabilities for optical transceivers, a crucial step to meet the explosive demand from data-intensive applications like AI and machine learning.

Technical / Clinical Details

Current data center infrastructure faces inherent bottlenecks with electrical signal transmission, encountering limitations in power consumption, latency, and transmission distance. Optical transceivers overcome these challenges by converting electrical signals into optical signals for transmission over fiber optic cables. The advent of 1.6T transceivers delivers unprecedented bandwidth and efficiency for interconnects between servers, switches, and storage within data centers. This capability will enable faster AI model training times and enhance real-time analytics.

The core of this technology lies in the fusion of high-precision analog circuits, digital signal processing (DSP), and advanced integrated photonics. TI leverages its particular strengths in analog front-end (AFE) and DSP technologies to contribute significantly to the performance enhancement of optical transceivers. Future advancements in Co-Packaged Optics (CPO) and silicon photonics technologies are expected to further accelerate the development of ultra-high-speed transceivers like 12.8T.

Background & Context

The proliferation of cloud computing, AI, and big data has dramatically increased the demands on data centers. Global data center power consumption is a significant concern, making improvements in energy efficiency a pressing priority. Optical communication is positioned as a promising solution to this challenge, capable of transmitting large volumes of data with significantly lower power consumption compared to electrical communication. Hyperscale data center operators such as Google, Meta, and Microsoft are actively making substantial investments in optical communication technologies.


Strategic Significance & Outlook

TI's roadmap indicates a decisive shift in data center networking towards optical-based architectures. The realization of 12.8T capabilities will vastly exceed current data center capacities, laying the groundwork for next-generation AI factories and edge computing environments. This will enable a dramatic increase in processing power while reducing the physical footprint of data centers, contributing to the construction of sustainable and scalable digital infrastructure. This technological evolution is poised to create new business opportunities and competitive advantages for telecommunication providers, cloud service providers, and AI developers alike, marking a new era of data connectivity.

Source: <https://www.ti.com/about-ti/behind-chip/articles/whats-next-for-the-future-of-data-center-networking-optics.html>

#12 Thin-Film LiNbO₃: UMC and HyperLight volume production specs

Published October 05, 2026 STT (Strategic Technology & Trends) / Substack USA



OVERVIEW

Thin-film lithium niobate (TFLN) technology is rapidly advancing, with HyperLight collaborating with UMC and Wavetek to accelerate volume production of its modulators. Concurrently, Huazhong University of Science and Technology's research team achieved a new record for broad-spectrum coverage with a single TFLN modulator. While TFLN demonstrates superior bandwidth, thermal independence, and low power, significant industrial volume production challenges persist despite these technological victories.

Key Findings

Significant strides have been made in thin-film lithium niobate (TFLN) technology, with HyperLight announcing a collaboration with UMC and Wavetek to boost volume production of its TFLN modulators. Simultaneously, a research team from Huazhong University of Science and Technology has achieved a new record for covering a broad spectrum with a single TFLN modulator. These developments underscore TFLN's physical advantages, particularly in bandwidth, temperature independence, and low power consumption for pluggable optical components.

Technical Details

The partnership between HyperLight, UMC, and Wavetek aims to enhance TFLN's industrial manufacturing capabilities, which is crucial for delivering high-performance and cost-effective optical modules demanded by data centers and communication networks. The Huazhong University research, on the other hand, demonstrates that a single TFLN modulator can achieve broadband signal modulation previously deemed impossible, potentially revolutionizing the capacity and flexibility of optical communication systems. TFLN, compared to conventional silicon photonics and bulk lithium niobate, exhibits superior electro-optic coefficients and low-loss characteristics, positioning it as a critical material for next-generation co-packaged optics (CPO) and near-packaged optics (NPO).

Background & Context

With the explosive growth of AI and high-speed data traffic, data center interconnects face increasing challenges in power consumption and thermal management. TFLN, due to its exceptional physical properties, has emerged as a promising material to address these issues. In pluggable optical modules, low power consumption and high thermal stability directly translate to reduced operational costs and improved reliability. However, despite these technical advancements, the author, Goyal, cautions that industrial-scale volume production still confronts "three walls": material cost, manufacturing process complexity, and supply chain establishment.

Strategic Significance & Outlook

TFLN technology's performance potential is widely recognized, and it is expected to become an indispensable component for next-generation optical communication systems and co-packaged optics in AI accelerators. If the challenges of volume production can be overcome, TFLN will dramatically improve optical interconnect performance, serving as a core technology to support the further evolution of data-driven societies.

Source: <https://www.simpletechtrend.com/en/post/tfln-optical-cpo-opportunity-walls-2026-en>

Collected: October 09, 2026 | Automated Research System (Gemini API)