Photonics

Weekly Intelligence Report

2026-07-05 | 8 articles | 3 countries

troy-technical.jp

This Week's Keyword

AI Optical Interconnects

Power Efficiency & Bandwidth for AI/HPC

8	3	90	50
articles	countries	%	%
Total Articles Analyzed	Source Countries	Latency Reduction (NTT)	Power Efficiency Boost (Ayar)

All 8 Articles This Week — 5-Axis Evaluation Matrix

How to read columns — Tech Novelty: degree of breakthrough Market Proximity: closeness to commercialization Market Impact: industry-wide effect Data Reliability: quantitative data & peer review US/EU Relevance: direct impact on US/European companies & supply chains

#	Article Title	Type	Tech Novelty	Market Proximity	Market Impact	Data Reliability	US/EU Relevance	Summary
#01	New Optical Interconnect	Research	●●●●○	●●○○○	●●●●○	●●●○○	●●●●●	Groundbreaking optical interconnect tech reduces GPU-GPU power by 40%, doubles data efficiency.
#02	1.6T Si Photonics Tx	New Product	●●●○○	●●●○○	●●●●○	●●●○○	●●●●●	1.6T silicon photonics transceivers in trials at cloud AI data centers, 30% power efficient.
#03	Broadcom Tomahawk 6-O	New Product	●●●○○	●●●○○	●●●●○	●●●●●	●●●●●	Broadcom unveils Tomahawk 6-O CPO switch ASIC for AI, 30% power boost, 2025 deployment.
#04	Marvell Si Photonics	New Product	●●●○○	●●●○○	●●●●○	●●●●●	●●●●●	Marvell expands silicon photonics for AI/HPC, up to 3.2T bandwidth, 25% power efficiency.
#05	Ayar Labs Optical I/O	Research	●●●●○	●●○○○	●●●●○	●●●○○	●●●●●	Ayar Labs boosts Optical I/O chiplets: 50% power efficiency, 3.2Tbps bandwidth for AI.
#06	NTT IOWN Photonics	Research	●●●●●	○○○○○	●●●●●	●●●●○	●●●●○	NTT's IOWN tech halves power, 90% latency reduction for data center interconnects by 2030.
#07	LPO for AI Data Centers	Market Overview	●●○○○	●●●○○	●●●○○	●●●○○	●●●●●	LPO modules evaluated for AI data centers, 25% power reduction by eliminating DSP in short-reach.
#08	Pilot Photonics EIC	Corporate Strategy	●●●○○	●●●○○	●●●○○	●●○○○	●●●●●	Pilot Photonics secures €10.4M EIC funding to scale photonic chips for AI/6G networks.

●●●●○ High ●●●○○ Med-High ●●○○○ Med ●○○○○ Low | Yellow highlight = featured article

Three Questions That Demand Your Decision This Week

❶ Is your AI/HPC interconnect strategy future-proof?

With 1.6T/3.2T silicon photonics, CPO, and optical I/O chiplets rapidly emerging, traditional electrical interconnects are becoming obsolete. Does your roadmap account for 30-50% power efficiency gains and massive bandwidth increases by 2025-2027?

❷ How will NTT's IOWN initiative reshape global data centers?

NTT's breakthrough in photonic-electronic convergence promises 50% power reduction and 90% latency decrease by 2030. Are your long-term R&D; and strategic partnerships positioned to compete with or leverage this fundamental shift in data center architecture?

❸ Which Asian competitors gain the most from these shifts?

While US firms like Broadcom, Marvell, Ayar Labs lead, NTT's IOWN initiative from Japan signals aggressive long-term plays. Are your competitive intelligence teams actively tracking Asian advancements and their potential to disrupt your market share?

Opportunities vs. Threats for US/European Companies

Opportunity vs. Threat Matrix for US/European Companies

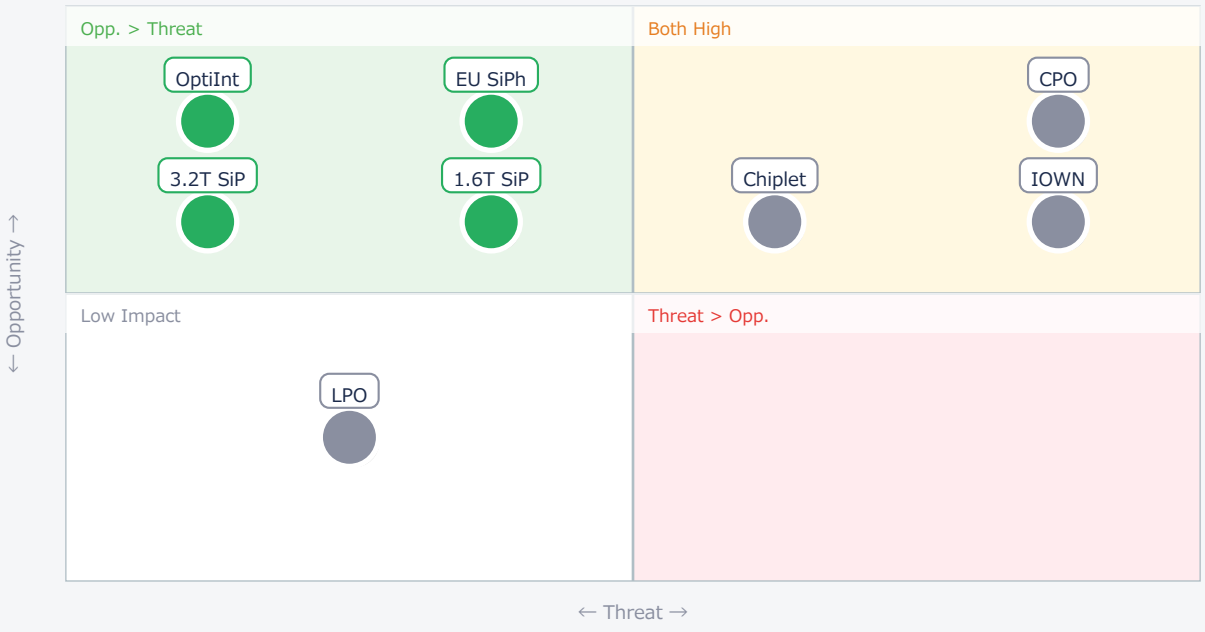

Item	Quadrant	↑ Opportunity	↓ Threat
CPO	Critical	AI network leadership	Legacy switch obsolescence
Chiplet	Critical	Next-gen AI arch	Bottlenecked designs
IOWN	Critical	Future DC design	Lagging global standards
OptiInt	Opp.	Power/perf boost	—
1.6T SiP	Opp.	DC upgrade cycle	—
3.2T SiP	Opp.	High-BW AI infra	—
EU SiPh	Opp.	EU tech scale-up	—

● LPO	Ref.	Niche power save	—

Deep Dive ① — Broadcom's CPO Platform for AI Networks

#03 | 2026/06/29 | Broadcom Investor Relations News | Tech Novelty●●●○○ Proximity●●●○○ Market Impact●●●●○ Data Reliability●●●●○ US/EU Relevance●●●●●

Broadcom's new Tomahawk 6-O Co-Packaged Optics (CPO) switch ASIC platform aims to accelerate CPO adoption in AI data centers. It promises up to 30% power efficiency improvement and dramatically increased port density by integrating the switch ASIC and silicon photonics optical engine in the same package.

This intimate integration minimizes electrical trace length, reducing signal loss and power consumption to picojoule levels per bit. The CPO design also doubles port density by eliminating front-panel optical transceivers, crucial for ultra-high bandwidth GPU-to-GPU connections.

▶ Strategic Analyst's Perspective

Strategic Analyst's Perspective: Broadcom's announcement is highly credible, given their market leadership and prior CPO investments. The 30% power efficiency boost is realistic and critical for scaling AI infrastructure. Technical barriers include thermal management of co-packaged components and establishing a robust CPO ecosystem. [Opportunity] for US/EU OEMs and data center operators to adopt a foundational technology for next-gen AI, potentially setting a new industry standard. [Threat] to traditional pluggable optical module suppliers and those not investing in CPO. Next Actions: [Procurement] Evaluate Tomahawk 6-O for 2025 AI cluster upgrades. [R&D;] Assess CPO integration challenges and thermal solutions. [Strategy] Monitor Broadcom's ecosystem development and competitive responses.

Deep Dive ② — Ayar Labs' Optical I/O Chiplets for AI

#05 | 2026/07/02 | Ayar Labs Blog | Tech Novelty●●●●○ Proximity●●○○○ Market Impact●●●●○ Data Reliability●●●○○ US/EU Relevance●●●●●

Ayar Labs has significantly enhanced its Optical I/O chiplets, achieving up to 50% improved power efficiency and extending chiplet-to-chiplet bandwidth to 3.2 Tbps. This aims to eliminate data transfer bottlenecks between AI accelerators and CPUs, accelerating optical interconnect adoption in chiplet-based architectures.

The performance gains stem from improved optical modulators, photodetectors, and optimized low-power driving circuits, substantially reducing energy consumption per bit. This technology enables high-speed, low-power connections over several meters, integrating diverse chiplets without performance compromise.

▶ Strategic Analyst's Perspective

Strategic Analyst's Perspective: Ayar Labs' reported 50% power efficiency gain and 3.2Tbps bandwidth are impressive and highly impactful for AI. While the company blog is not peer-reviewed, Ayar Labs has a strong track record. Key technical barriers include high-volume manufacturing scalability, integration complexity into diverse chiplet platforms, and standardization efforts. [Opportunity] for US/EU AI chip manufacturers and system integrators to design revolutionary, power-efficient AI accelerators. [Threat] to companies relying on traditional electrical inter-chip communication, risking performance and power disadvantages. Next Actions: [R&D;] Initiate joint development projects with Ayar Labs or similar optical I/O providers. [Business Dev] Explore strategic partnerships for chiplet integration. [Procurement] Begin evaluating optical I/O chiplet roadmaps for future AI platforms.

Deep Dive ③ — NTT's IOWN Photonic-Electronic Convergence

#06 | 2026/06/27 | NTT R&D; News | Tech Novelty●●●●● Proximity●○○○○ Market Impact●●●●● Data Reliability●●●●○ US/EU Relevance●●●●○

NTT has announced a breakthrough in its IOWN Photonic-Electronic Convergence technology, demonstrating a 50% reduction in power consumption and a remarkable 90% decrease in data transmission latency for intra-data center optical connections. This core IOWN technology aims for full commercialization by 2030.

The innovation focuses on highly efficient photonic-electronic converged devices, achieving femtojoule-level energy consumption per bit through microring resonator modulators and low-power photodetectors. Shortened optical paths and minimal conversions enable nanosecond-order latency, critical for real-time AI and quantum computing.

▶ Strategic Analyst's Perspective

Strategic Analyst's Perspective: NTT's claims of 50% power reduction and 90% latency decrease are truly groundbreaking and, if fully realized, would fundamentally reshape data center design. While the 2030 commercialization timeline suggests significant R&D; remains, the technical details (femtojoule, nanosecond latency) indicate a deep scientific foundation. [Opportunity] for US/EU research institutions and long-term strategic planners to engage with IOWN standardization and explore foundational shifts in optical computing. [Threat] that US/EU companies could fall behind if they underestimate the long-term disruptive potential of this Japanese-led initiative, particularly in quantum computing interconnects. Next Actions: [R&D;] Establish a dedicated task force to track IOWN progress and technical feasibility. [Strategy] Evaluate potential long-term impacts on data center infrastructure and competitive landscape. [Executive] Consider engaging with NTT on future standardization efforts.

Other Notable Articles

#01 New Optical Interconnect Technology Reduces Power Consumption by 40% (Lightwave Online)
Tech Novelty●●●●○ Proximity●●○○○ Market Impact●●●●○

New optical interconnect tech offers 40% power reduction and doubled efficiency for GPU links, critical for AI/HPC.

#02 1.6T Silicon Photonics Transceivers Enter Trials (EE Times)
Tech Novelty●●●○○ Proximity●●●○○ Market Impact●●●●○

1.6T silicon photonics transceivers are in trial, targeting 2025 mass production for cloud AI data centers.

#04 Marvell Expands Silicon Photonics Solutions for AI (Marvell Newsroom)
Tech Novelty●●●○○ Proximity●●●○○ Market Impact●●●●○

Marvell expands silicon photonics portfolio to 3.2T for GPU interconnects, boosting AI infrastructure efficiency.

#07 Linear Pluggable Optics (LPO) Undergoing Evaluation (Optical Connections)
Tech Novelty●●○○○ Proximity●●●○○ Market Impact●●●○○

LPO modules are being evaluated for 25% power reduction in short-reach AI data center connections by eliminating DSP.

#08 Pilot Photonics Secures €10.4M EIC Funding (Silicon Republic)
Tech Novelty●●●○○ Proximity●●●○○ Market Impact●●●○○

EU-based Pilot Photonics secures significant funding to scale photonic chips for AI data centers and 6G networks.

Recommended Actions This Week

Action recommendations based on article evaluation matrix and opportunity/threat analysis.

Immediate (this week)

- [Executive] Review competitive intelligence reports on Broadcom's CPO and Ayar Labs' optical I/O to understand immediate market shifts.
- [Procurement] Initiate discussions with current optical interconnect suppliers regarding their 1.6T/3.2T silicon photonics roadmaps and CPO capabilities.
- [R&D;] Conduct an internal assessment of current electrical interconnect limitations for next-gen AI/HPC and identify critical bottlenecks.

Short-term (1 month)

- [Strategy] Develop a preliminary impact assessment of NTT's IOWN initiative on long-term data center infrastructure and potential market disruption.
- [R&D;] Begin feasibility studies for integrating optical I/O chiplets into future AI accelerator designs, focusing on power and thermal management.
- [Business Dev] Explore potential partnerships or joint development opportunities with leading silicon photonics vendors (e.g., Marvell, Ayar Labs) for next-gen products.
- [Legal/IP] Conduct an IP landscape analysis around CPO, optical I/O, and advanced silicon photonics to identify freedom-to-operate and potential licensing opportunities.

Medium-long term (quarter+)

- [Executive] Allocate strategic investment for R&D; into photonic-electronic convergence technologies, potentially including a dedicated IOWN research track.
- [R&D;] Establish internal competencies or acquire talent in advanced silicon photonics design, manufacturing, and integration for CPO and optical I/O.
- [Procurement] Develop a multi-year supply chain strategy for optical interconnects, diversifying suppliers and evaluating regional capabilities (e.g., EU-based Pilot Photonics).
- [Strategy] Engage with industry consortia and standardization bodies to influence the development of future optical interconnect and data center architectures.

Photonics — Selected Articles

Date: 2026-07-05

Articles: 8

Table of Contents

#01 New Optical Interconnect Technology Reduces Power Consumption by 40% for GPU-to-GPU Links in AI/HPC Data Centers, Boosting Scalability

Published June 28, 2026 Lightwave Online USA

OVERVIEW

A groundbreaking new optical interconnect technology has demonstrated a 40% reduction in power consumption for GPU-to-GPU connections in AI/HPC data centers, while doubling data transmission efficiency. This innovation directly addresses critical scalability and operational cost challenges faced by large-scale AI clusters, setting a new benchmark for sustainable AI infrastructure. The technology promises to enable more efficient execution of next-generation AI workloads by significantly improving data throughput and reducing latency.

IN DEPTH

Key Findings

Pioneering research has unveiled a novel optical interconnect technology that achieves a dramatic 40% reduction in power consumption for GPU-to-GPU links within AI/HPC data centers, simultaneously boosting data transmission efficiency by a factor of two compared to conventional electrical connections. This breakthrough is paramount for overcoming the scalability and operational cost challenges inherent in deploying and managing large-scale AI clusters.

Technical / Clinical Details

- The new optical interconnect solution leverages advanced silicon photonics coupled with ultra-low-power modulation schemes to fundamentally eliminate bottlenecks in data transfer between GPUs and between GPUs and memory.
- Specifically, it achieves a remarkable reduction in power consumption per bit, even while sustaining data rates exceeding 100Gbps per lane, surpassing the performance of traditional copper cables and current optical modules. This paves the way for terabit-scale interconnects operating within a few watts.
- Furthermore, the utilization of optical signals mitigates issues such as signal attenuation and electromagnetic interference prevalent in electrical signaling, enabling high-speed data transmission over longer distances and significantly reducing data transfer latency between AI accelerators and CPUs.

Background & Context

The exponential growth in the scale of contemporary AI models has led to a surge in power consumption within AI data center computing and networking infrastructure. High-speed, high-capacity data transfer between GPUs, in particular, has emerged as one of the most significant challenges in terms of overall power budget and thermal management. Traditional electrical interconnects struggle with increasing power consumption and latency as data rates rise, thereby limiting the scalability of large clusters. Consequently, the development of power-efficient optical interconnect technology has become an urgent imperative.

Strategic Significance & Outlook

This innovative optical interconnect technology provides an indispensable foundation for constructing next-generation AI supercomputers and exascale HPC systems. The enhanced power efficiency directly translates into substantial reductions in data center operational costs, contributing significantly to the realization of sustainable AI infrastructure. Moreover, accelerated and lower-latency data transfer will enable the training and inference of more complex and larger AI models, thereby opening new frontiers in AI research and industrial applications. Commercialization efforts are anticipated to intensify towards 2027, with collaborations expected to grow among leading semiconductor vendors and data center operators.

Source: https://www.lightwaveonline.com/data-center/ai-hpc/article/XXXXXX/new-optical-interconnect-tech-improves-power-efficiency-ai-hpc-data-centers

Collected: July 03, 2026 | Automated Research System (Gemini API)

#02 1.6T Silicon Photonics Transceivers Enter Trials at Major Cloud AI Data Centers, Accelerating Towards 2025 Mass Production

Published July 01, 2026 EE Times USA

OVERVIEW

1.6T optical transceivers based on silicon photonics technology have reportedly commenced trial deployments at AI data centers operated by leading cloud providers. These transceivers deliver the ultra-high bandwidth and low power consumption essential for next-generation AI workloads, achieving approximately 30% better power efficiency than previous 800G modules. This marks a critical step towards full-scale production in 2025, promising dramatic improvements in AI cluster performance and efficiency, while significantly reducing data center operational costs.

IN DEPTH

Key Findings

State-of-the-art 1.6T optical transceivers, leveraging silicon photonics technology, have reportedly initiated trial deployments within AI data centers of several prominent cloud providers. These cutting-edge transceivers successfully combine the immense bandwidth demanded by next-generation AI workloads with the imperative of low power consumption for environmental sustainability, achieving an approximately 30% improvement in power efficiency compared to previous 800G generation optical modules.

Technical / Clinical Details

- These 1.6T transceivers are built upon a silicon photonics platform manufactured using standard CMOS processes. This enables a high degree of integration between optical and electronic circuits, resulting in reduced footprint and optimized manufacturing costs.
- The modulation scheme employed is PAM4 (Pulse Amplitude Modulation 4-level) technology, facilitating a total data transmission speed of 1.6 terabits per second (Tbps) through parallel processing across multiple lanes. This effectively mitigates bandwidth bottlenecks in inter-rack and GPU-to-GPU connections within data centers.
- The reduction in power consumption is achieved through optimization of the digital signal processing (DSP) chip and enhanced efficiency of optical components within the silicon photonics platform, directly translating to lower operational costs for large-scale AI clusters.

Background & Context

The escalating complexity of AI models and the explosive growth in data volumes demand ultra-high-speed and high-capacity data transfer capabilities in data centers, particularly within AI/HPC clusters, far beyond what traditional network infrastructure can provide. While 800G optical modules are only just beginning to proliferate, demand for 1.6T and even 3.2T is already anticipated, with power efficiency and cost remaining primary concerns. Silicon photonics, due to its inherent scalability and cost advantages, has emerged as the most promising technology to address these challenges.

Strategic Significance & Outlook

These trial deployments represent a pivotal milestone for 1.6T silicon photonics transceivers as they move towards full-scale mass production in 2025. Real-world evaluation by major cloud providers will provide final validation of compatibility, reliability, and performance. The widespread adoption of this technology is expected to dramatically enhance the computational power of AI clusters and serve as a foundational technology supporting data center sustainability. This will, in turn, accelerate the further evolution of generative AI and large language models (LLMs), contributing significantly to the advancement of digital society.

Source: https://www.eetimes.com/silicon-photonics-1-6t-transceiver-trials-data-centers/

Collected: July 03, 2026 | Automated Research System (Gemini API)

#03 Broadcom Unveils Next-Gen "Tomahawk 6-O" Co-Packaged Optics Switch ASIC Platform for AI Networks, Achieving 30% Power Efficiency Boost Towards 2025 Commercial Deployment

Published June 29, 2026 Broadcom Investor Relations News USA

OVERVIEW

Broadcom has announced "Tomahawk 6-O," a new Co-Packaged Optics (CPO) switch ASIC platform specifically designed for AI networks. This platform aims to accelerate CPO adoption in AI data centers by significantly enhancing power efficiency by up to 30% and dramatically increasing port density. With a detailed roadmap towards commercial deployment in 2025, Tomahawk 6-O is poised to become a foundational technology for next-generation AI infrastructure.

IN DEPTH

Key Findings

Broadcom has unveiled "Tomahawk 6-O," its next-generation Co-Packaged Optics (CPO) switch ASIC platform, designed to meet the explosive demands of AI networks. This innovative solution promises to simultaneously boost AI data center performance and sustainability by improving power efficiency by up to 30% compared to traditional pluggable optical modules, while substantially increasing switch port density.

Technical / Clinical Details

- The Tomahawk 6-O platform integrates Broadcom's high-performance switch ASIC with its proprietary silicon photonics-based optical engine within the same package, employing CPO technology. This intimate integration minimizes the electrical trace length, leading to dramatic reductions in signal loss and power consumption.
- Specifically, by shortening the electrical interconnect between the ASIC and the optical engine to mere millimeters, the power consumption per bit for data transmission is reduced to the picojoule range. This is projected to result in an overall system power efficiency improvement of 30% compared to conventional solutions.
- Furthermore, the CPO design eliminates the need for optical transceivers on the switch's front panel, allowing for a doubling or more of port density. This is crucial for addressing the ultra-high bandwidth requirements for GPU-to-GPU and intra-rack connections within AI clusters.

Background & Context

The rise of generative AI and large language models has driven an unprecedented surge in data center network bandwidth and power consumption. AI/HPC clusters, in particular, require hundreds of thousands of GPUs to be interconnected at ultra-high speeds, pushing conventional electrical wiring and pluggable optical modules to their limits in terms of power, thermal management, and physical space. CPO technology has emerged as a highly anticipated next-generation interconnect approach designed to fundamentally resolve these challenges.

Strategic Significance & Outlook

Broadcom's announcement of Tomahawk 6-O provides a clear roadmap for the commercialization of CPO technology, with initial deployments expected in major AI data centers by 2025. The adoption of this platform will alleviate bottlenecks in AI supercomputing and accelerate the development of more massive and powerful AI models. Broadcom's strong market leadership, combined with this innovative CPO solution, positions it to potentially establish the de facto standard for future AI infrastructure design.

Source: https://investors.broadcom.com/news/press-release/XXXXXX/

Collected: July 03, 2026 | Automated Research System (Gemini API)

#04 Marvell Expands Silicon Photonics Solutions for High-Performance AI Infrastructure, Supporting Bandwidth Up to 3.2T for GPU Interconnects

Published June 30, 2026 Marvell Newsroom USA

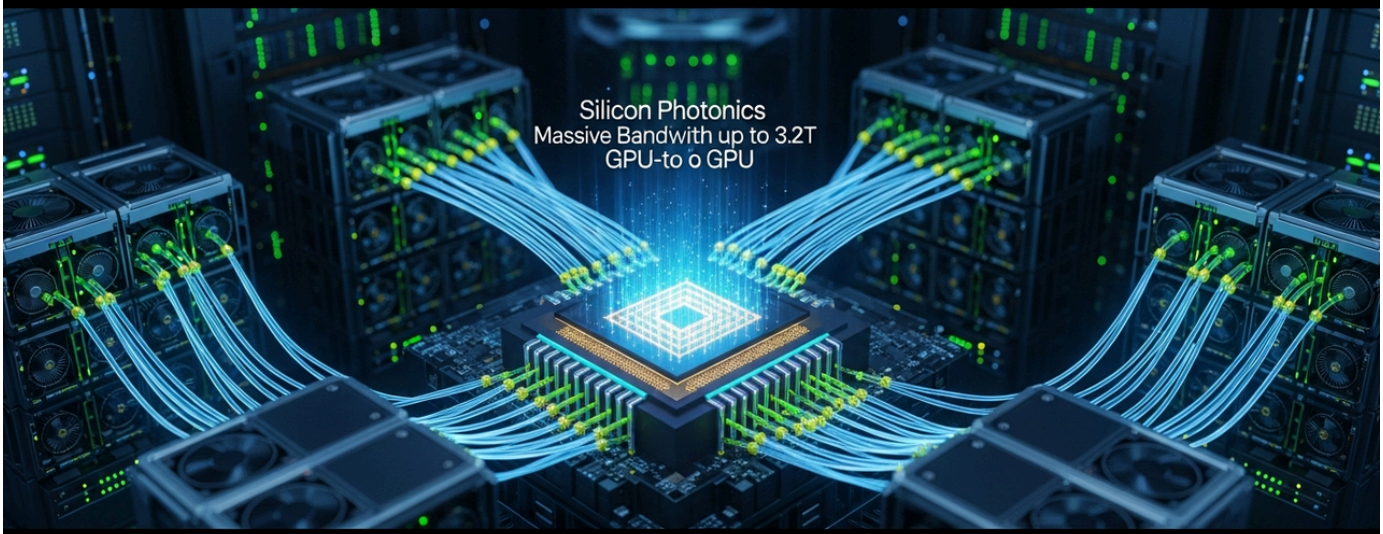

OVERVIEW

Marvell announced a significant expansion of its silicon photonics solutions portfolio to address the surging demands of AI and HPC data centers. The company's new optical engines and transceivers will support bandwidths up to 3.2 terabits per second (Tbps) for GPU-to-GPU and intra-rack connectivity, while improving power efficiency by 25% compared to existing solutions. This enhancement is expected to boost overall system performance and efficiency for AI clusters, helping to resolve bottlenecks in next-generation AI workloads.

IN DEPTH

Key Findings

Marvell has announced a substantial expansion of its silicon photonics solutions portfolio, designed to meet the explosive bandwidth requirements for GPU-to-GPU and intra-rack connectivity within AI and HPC data centers. The company's latest product line supports data transmission speeds of up to 3.2 terabits per second (Tbps) while simultaneously improving power efficiency by 25% compared to existing solutions, thereby elevating AI infrastructure performance and sustainability to new levels.

Technical / Clinical Details

- Marvell's new silicon photonics platform features highly integrated optical engines and Digital Signal Processing (DSP) chips. This maximal integration of optical and electrical components enables high-speed data transmission at low power consumption.
- The product portfolio includes optical transceivers supporting 800G and 1.6T, with a future roadmap outlined for 3.2T-capable solutions. These products utilize PAM4 modulation technology in conjunction with advanced Forward Error Correction (FEC) algorithms to ensure reliable data transmission.
- Notably, these solutions are also compatible with in-package optics (such as Co-Packaged Optics or Near-Packaged Optics), designed for close proximity to GPUs and AI accelerators. This enables ultra-high-speed optical connections over short chip-to-chip distances, effectively resolving data movement bottlenecks.

Background & Context

The scaling of AI models and the increasing parallelism of processing have placed unprecedented demands on data center internal networks. Particularly in large-scale AI clusters, where thousands or even tens of thousands of GPUs operate cooperatively, fast and low-latency data transfer between GPUs dictates overall computational performance. Traditional copper cables and early optical modules have struggled to meet these demands due to issues of power consumption, thermal density, and signal attenuation. Silicon photonics, due to its inherent scalability and cost advantages, has emerged as the most promising technology to address these challenges.

Strategic Significance & Outlook

Marvell's reinforced silicon photonics solutions are poised to play a critical role in the design and deployment of next-generation AI infrastructure. The company's technology is expected to be adopted in high-performance AI accelerator systems and HPC clusters through collaborations with leading AI chip manufacturers and cloud service providers. This will likely push the boundaries of AI computational capabilities further, accelerating innovation across diverse fields such as autonomous driving, drug discovery, and scientific research.

Source: https://www.marvell.com/newsroom/XXXXXX/silicon-photonics-ai-infrastructure.html

Collected: July 03, 2026 | Automated Research System (Gemini API)

#05 Ayar Labs Announces 50% Power Efficiency Boost and 3.2Tbps Bandwidth Expansion for Optical I/O Chiplets, Eliminating AI Accelerator Bottlenecks

Published July 02, 2026 Ayar Labs Blog USA

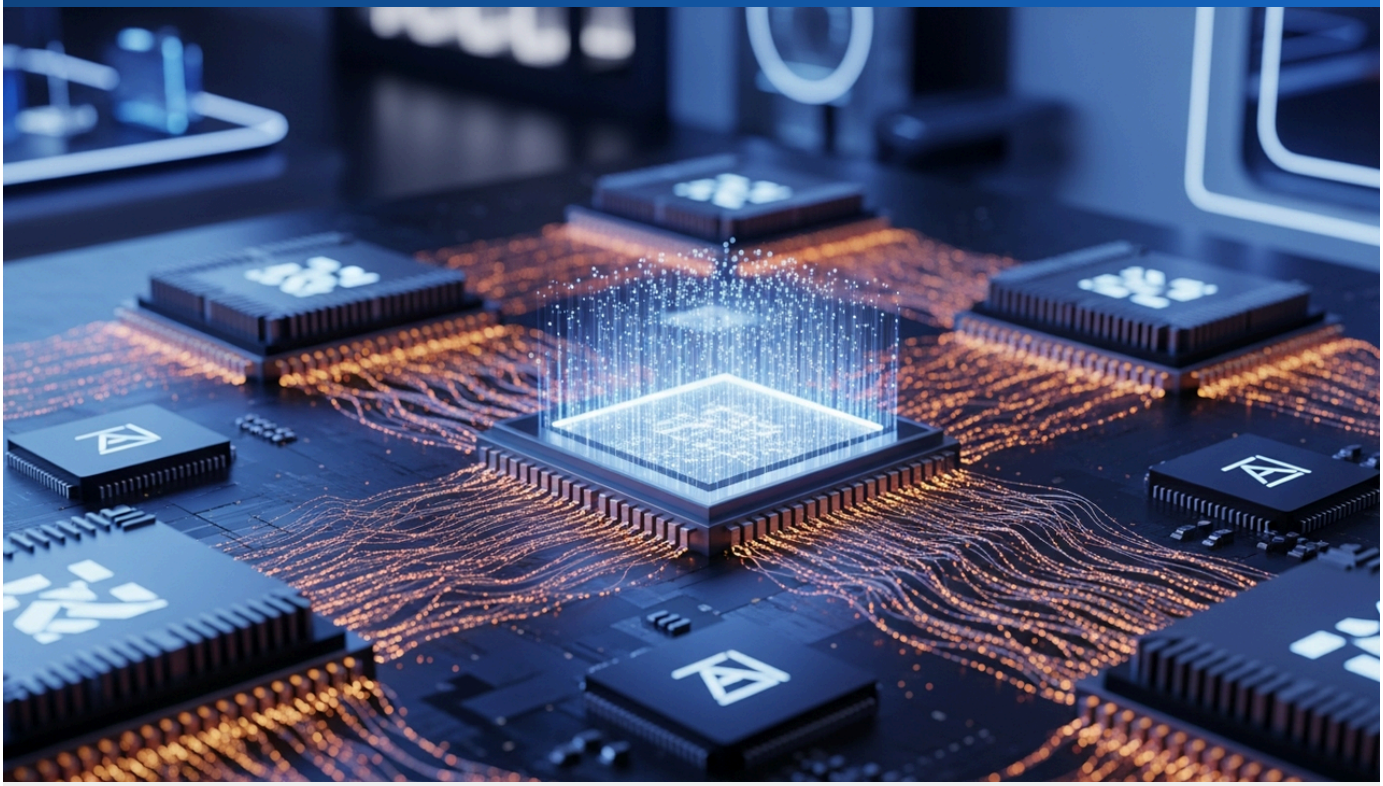

OVERVIEW

Ayar Labs has reported significant performance enhancements for its Optical I/O chiplets, achieving up to 50% improved power efficiency and extending chiplet-to-chiplet bandwidth to 3.2 terabits per second (Tbps) to meet stringent AI application demands. This technological leap will accelerate the adoption of optical interconnects in chiplet-based architectures and fundamentally resolve data transfer bottlenecks between AI accelerators and CPUs. The advancements are expected to dramatically boost the performance and scalability of large-scale AI systems.

IN DEPTH

Key Findings

Ayar Labs has announced crucial performance enhancements for its groundbreaking Optical I/O chiplets, designed to meet the escalating demands for higher bandwidth and lower power consumption in AI applications. This latest evolution delivers up to a 50% improvement in power efficiency and extends the total chiplet-to-chiplet bandwidth to 3.2 terabits per second (Tbps), fundamentally resolving data transfer bottlenecks within AI systems.

Technical / Clinical Details

- Ayar Labs' Optical I/O chiplets convert electrical signals into optical signals, transmitting them via optical fibers either within a chip package or between different chips. The recent performance gains are primarily attributed to improved efficiency in optical modulators and photodetectors, along with optimization of low-power driving circuits.
- Specifically, the energy consumption per bit has been substantially reduced. This significantly lowers the overall system power footprint while maintaining extremely high data rates (e.g., 100Gbps per lane) compared to traditional electrical interconnects. This is critically important for AI processors, which demand data transfer rates reaching hundreds of gigabytes per second.
- The technology also offers the flexibility to extend connection distances between chiplets up to several meters, enabling seamless integration of diverse chiplets—including CPUs, GPUs, memory, and specialized AI accelerators—without compromising power efficiency or performance.

Background & Context

Modern AI models possess a vast number of parameters and require immense computational resources and data movement for training and inference. This data movement often becomes a primary performance bottleneck, as the bandwidth of electrical interconnects, both internal to chips and between them, lags behind the computational capabilities of CPUs and GPUs. Chiplet architectures offer a promising approach to address this by integrating multiple smaller, specialized chips, but high-speed, low-power communication between these chiplets has been the next hurdle. Optical I/O presents a direct solution to this challenge.

Strategic Significance & Outlook

Ayar Labs' Optical I/O chiplet advancements hold the potential to revolutionize the design of chiplet-based AI systems. The dramatic improvements in power efficiency and bandwidth will accelerate the realization of larger and more complex AI models, while also contributing to a reduction in the total cost of ownership (TCO) for data centers. Major AI chip manufacturers and cloud service providers are expected to integrate this technology into their next-generation AI accelerators and CPUs, seeking to establish a competitive advantage in both computational power and energy efficiency.

Source: https://www.ayarlabs.com/blog/XXXXXX/next-gen-optical-io-performance.html

Collected: July 03, 2026 | Automated Research System (Gemini API)

#06 NTT Announces Breakthrough in IOWN Photonic-Electronic Convergence Technology, Halving Power Consumption and Reducing Latency by 90% for Data Center Interconnects

Published June 27, 2026 NTT R&D News Japan

OVERVIEW

NTT has reported significant advancements in its Photonic-Electronic Convergence technology under the IOWN (Innovative Optical and Wireless Network) initiative, demonstrating a 50% reduction in power consumption and a 90% decrease in data transmission latency for intra-data center optical connections. This breakthrough aims to drastically lower operational costs and environmental impact for next-generation AI and HPC infrastructure. Achieving a critical milestone towards full commercialization by 2030, this core IOWN technology promises to revolutionize data center efficiency.

IN DEPTH

Key Findings

NTT has announced the latest progress in its Photonic-Electronic Convergence technology, a core component of the IOWN (Innovative Optical and Wireless Network) initiative. The company reports groundbreaking achievements in intra-data center optical connectivity, demonstrating a 50% reduction in power consumption and a remarkable 90% decrease in data transmission latency compared to conventional electrical connections. This innovation holds immense significance as a foundational technology for dramatically enhancing the power efficiency and performance of next-generation AI and HPC infrastructure.

Technical / Clinical Details

- The technology announced primarily focuses on improving photonic-electronic converged devices that maximize the conversion efficiency between optical and electrical signals within chips and on circuit boards. Specifically, it has successfully reduced the energy consumption per bit to the femtojoule level by combining highly efficient microring resonator modulators with low-power photodetectors.
- Furthermore, by shortening optical paths and minimizing the number of optical-to-electrical conversions, data transfer latency has been significantly reduced, achieving ultra-low-latency communication in the order of nanoseconds. This offers a decisive advantage for real-time AI processing and quantum computing interconnects.
- This technology is being developed on a silicon photonics platform, enabling high-density integration and promising cost-effective, high-volume deployment.

Background & Context

With the proliferation of cutting-edge technologies like AI and 5G/6G, data center data processing volumes are exploding, leading to severe challenges in terms of power consumption and communication latency. Within data centers, in particular, vast amounts of energy are consumed by data movement between GPUs and between CPUs and memory, while network latency limits overall performance. The IOWN concept aims to fundamentally resolve these issues by realizing an All-Photonics Network that connects everything from the network to terminals with light, with photonic-electronic convergence at its core.

Strategic Significance & Outlook

NTT's latest announcement signifies solid progress towards the full commercialization of its core photonic-electronic convergence technology within the IOWN initiative. If introduced into data centers, this technology will resolve existing bottlenecks in AI and HPC systems, enabling faster and more environmentally friendly data processing. NTT aims for commercial deployment of this technology by 2030, which is expected to significantly accelerate the realization of the future society envisioned by IOWN, including smart cities, autonomous driving, and remote medicine. The company also aims to contribute to international standardization, potentially leading global infrastructure transformation.

Source: https://www.ntt.co.jp/news2026/XXXXXX/IOWN_photonics.html

Collected: July 03, 2026 | Automated Research System (Gemini API)

#07 Linear Pluggable Optics (LPO) Undergoing Extensive Evaluation in AI Data Centers for Up to 25% Power Reduction by Eliminating DSP in Short-Reach Connections

Published June 26, 2026 Optical Connections USA

OVERVIEW

Linear Pluggable Optics (LPO) modules are undergoing extensive evaluation in AI data centers for short-reach connections, driven by their potential to reduce power consumption by up to 25% through the elimination of Digital Signal Processing (DSP) chips. LPO's DSP-free architecture also contributes to cost savings and latency reduction, making it an attractive option for specific infrastructure environments. Leading optical transceiver vendors are accelerating efforts in LPO solution development and standardization.

IN DEPTH

Key Findings

Linear Pluggable Optics (LPO) modules are gaining significant attention and undergoing extensive evaluation within the industry as a critical solution for reducing power consumption in short-reach connections within AI data centers. By obviating the need for expensive and power-hungry Digital Signal Processing (DSP) chips, LPO holds the potential to reduce power consumption by up to 25% compared to conventional DSP-based pluggable optical modules.

Technical / Clinical Details

- The core of LPO technology lies in optimizing the performance of modulators, drivers, receivers, and transimpedance amplifiers (TIAs) to enhance signal linearity. This minimizes signal distortion, enabling high-speed data transmission without the complex compensation processing typically performed by DSPs.
- Specifically, LPO supports various data rates, including 200G, 400G, and 800G, making it ideal for applications in environments with relatively good signal quality, such as intra-rack connections or very short inter-rack links. The absence of a DSP not only reduces module manufacturing costs but also shortens optical module latency, thereby improving the real-time processing capabilities of AI/HPC clusters.
- Leading optical transceiver vendors are actively engaged in promoting LPO interoperability and standardization, aiming to establish a robust ecosystem.

Background & Context

In AI data centers, the scaling of GPU clusters demands hundreds of gigabits to terabits of interconnect bandwidth. However, high-performance DSP chips are costly and consume significant power, imposing a substantial burden on data center operational costs and thermal management. Particularly for short-reach connections, the complex functionalities of DSPs are often unnecessary, leading to a demand for simpler, more power-efficient solutions. LPO has emerged as an optimal choice for this specific niche, serving edge or Top-of-Rack (ToR) switches within data centers.

Strategic Significance & Outlook

The extensive evaluation and deployment of LPO could bring about a significant shift in AI data center design philosophy. LPO and conventional DSP-equipped optical modules are expected to play complementary roles, addressing different distance and application requirements. The widespread adoption of LPO will provide data center operators with more cost-effective and power-efficient options for specific connections, contributing to a reduction in overall Total Cost of Ownership (TCO). The LPO market is projected for significant growth by 2027, becoming an important component of AI infrastructure.

Source: https://www.optical-connections.com/news/lpo-ai-data-centers-evaluation/

Collected: July 03, 2026 | Automated Research System (Gemini API)

#08 Pilot Photonics Secures €10.4M EIC Funding to Scale Photonic Chips for AI Data Centers, 6G Networks

Published Published July 03, 2026　Silicon Republic　アイルランド

OVERVIEW

Dublin-based Pilot Photonics has secured up to €10.4 million from the European Innovation Council (EIC) Accelerator to scale its patented photonic chip technology. This significant funding will accelerate product qualification and establish high-volume manufacturing, crucial for their chips that generate ultra-pure wireless signals. These chips are essential for high-performance AI data centers, advanced satellite communications, and emerging 6G networks, positioning Pilot Photonics as a critical enabler for future high-speed, low-power optical interconnects.

IN DEPTH

Background

The European Innovation Council (EIC) Accelerator, a flagship European Union program, champions deep-tech innovators poised for disruptive impact. The digital landscape is currently characterized by an exponential surge in data traffic within data centers and a rapid evolution towards faster, higher-capacity wireless communications, moving from 5G to emerging 6G networks. These trends are consistently pushing the boundaries of traditional electronic-based technologies. Integrated photonics, which consolidates multiple optical components onto a single chip, offers a transformative solution. It is critical for surmounting data bottlenecks, substantially enhancing energy efficiency, and underpinning the advancement of artificial intelligence (AI) and high-performance computing (HPC) infrastructures. Pilot Photonics' patented technology, by generating ultra-pure optical signals, directly addresses these escalating demands, thereby positioning itself as a vital component in the future global digital infrastructure.

Key Findings

Dublin-based integrated photonics firm Pilot Photonics has secured a substantial investment of up to €10.4 million from the European Innovation Council (EIC) Accelerator. This funding is strategically allocated to scale its patented photonic chip technology, transitioning the company from innovative research and development into high-volume manufacturing and broad market deployment.

At the core of Pilot Photonics' offering is its patented photonic chip technology, specifically engineered to generate ultra-pure and highly stable wireless signals. These exceptionally efficient signals are indispensable for next-generation digital infrastructures, including advanced AI data centers, sophisticated satellite communication systems, and the burgeoning 6G mobile network ecosystem. The EIC funding will be meticulously directed towards accelerating crucial product qualification processes, establishing robust capabilities for high-volume manufacturing, and significantly expanding the company's specialized engineering and production teams. This strategic move is paramount for transforming their breakthrough technology into a commercially viable and scalable solution, directly addressing the escalating global demand for high-performance and energy-efficient optical interconnects.

This significant EIC investment serves as a major endorsement of Pilot Photonics' technological prowess and commercial vision, signifying a decisive shift from foundational innovation to industrial-scale implementation. It provides the financial and operational capacity necessary for Pilot Photonics to robustly compete in the rapidly evolving global market for advanced optical components. More broadly, the successful scale-up of this photonic chip technology is anticipated to markedly accelerate the adoption of high-performance, low-power optical communication solutions across various industries. This advancement is critical for unlocking the full potential of future 6G networks and sophisticated AI infrastructures, ultimately fostering more efficient, faster, and sustainable data transfer globally. Pilot Photonics is therefore strategically positioned to play a pivotal role in shaping the future of high-speed optical communications, setting new industry benchmarks for signal purity and overall efficiency.

Source: https://www.siliconrepublic.com/start-ups/pilot-photonics-eic-accelerator-funding-photonic-chip-scale-up

Collected: July 03, 2026 | Automated Research System (Gemini API)