

TROY TECHNICAL

OPTICAL COMMUNICATIONS & PHOTONICS WEEKLY REPORT

October 17 - October 23, 2026 | Issue 2026-W42

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

12

ANALYZED ARTICLES

10

PRIORITY THEMES

Issued 2026-10-23 | 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

GaN MicroLEDs: 2.75 GHz Bandwidth Record and Scaling Roadmap

Researchers at the University of Illinois at Urbana-Champaign have achieved a world-record modulation bandwidth of 2.75 GHz in a blue GaN microLED.

SIGNAL 02 | PHOTONIC DEVICE

Linktel 1.6T LPO: Performance Specs and AI Network Roadmap

Linktel Technologies demonstrated 1.6T Linear Pluggable Optics (LPO) at OCP 2026, featuring power consumption as low as 10W and ultra-low latency.

SIGNAL 03 | PACKAGING

NVIDIA OOI Packaging: CPO Roadmap and 2028 Feynman Platform

NVIDIA has developed a prototype "Optics-on-Interposer" (OOI) chip that integrates optical engines directly into the silicon interposer to eliminate bandwidth bottlenecks in AI computing.

SIGNAL 04 | NETWORK SYSTEM

Silicon Photonics Interconnect: 3.2 Tbps AI Data Traffic Chip

Researchers at Zhejiang University have developed a scalable silicon photonic interconnect architecture achieving an aggregate bandwidth of 3.2 terabits per second in the O-band.

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

2.75 GHz A07	SOURCE-BOUND CONDITION Researchers at the University of Illinois at Urbana-Champaign have achieved a world-record modulation bandwidth of 2.75 GHz in a blue GaN microLED.
10W A10	SOURCE-BOUND CONDITION Linktel Technologies demonstrated 1.6T Linear Pluggable Optics (LPO) at OCP 2026, featuring power consumption as low as 10W and ultra-low latency.
100 mm A01	SOURCE-BOUND CONDITION This technology significantly reduces component footprint from 100 mm to 8 mm, paving the way for the 2028 Rosa Feynman platform.
400 G A08	SOURCE-BOUND CONDITION NVIDIA and Cisco have developed Ultra-Dense Small Form Factor Pluggable (USFP) optics to support 400 Gb/s-per-lane connectivity for AI factories.
10 GHz A05	SOURCE-BOUND CONDITION The system achieves a 10 GHz modulation bandwidth and reduces power consumption by over 75% compared to conventional current-based methods.
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.



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

GaN MicroLEDs: 2.75 GHz Bandwidth Record and Scaling Roadmap

A07 | Light source | Pilot

Date not stated

FACT

Researchers at the University of Illinois at Urbana-Champaign have achieved a world-record modulation bandwidth of 2.75 GHz in a blue GaN microLED. By analyzing devices from 100µm down to 2µm, the team established a device-level roadmap to replace copper links with high-bandwidth optical interconnects for AI computing.

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 2/5 | DECISION LOCUS: TECH / PROCESS

P02

Linktel 1.6T LPO: Performance Specs and AI Network Roadmap

A10 | Photonic device | Product adoption

Date not stated

FACT

Linktel Technologies demonstrated 1.6T Linear Pluggable Optics (LPO) at OCP 2026, featuring power consumption as low as 10W and ultra-low latency. The company has moved these modules into volume production to address bandwidth and thermal constraints in AI data centers.

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 20/25 | MATURITY 4/5 | DECISION LOCUS: TECH / PROCESS

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

P03

NVIDIA OOI Packaging: CPO Roadmap and 2028 Feynman Platform

A01 | Packaging | Joint evaluation

Date not stated

FACT

NVIDIA has developed a prototype "Optics-on-Interposer" (OOI) chip that integrates optical engines directly into the silicon interposer to eliminate bandwidth bottlenecks in AI computing. This technology significantly reduces component footprint from 100 mm to 8 mm, paving the way for the 2028 Rosa Feynman platform.

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 3/5 | DECISION LOCUS: SYSTEM / ADOPTION

P04

Silicon Photonics Interconnect: 3.2 Tbps AI Data Traffic Chip

A03 | Network system | Pilot

Date not stated

FACT

Researchers at Zhejiang University have developed a scalable silicon photonic interconnect architecture achieving an aggregate bandwidth of 3.2 terabits per second in the O-band. By combining hybrid wavelength-port multiplexing with adaptive stabilization, the system addresses the energy and bandwidth bottlenecks of copper-based electrical interconnects in AI clusters.

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

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

P05

USFP Optics: NVIDIA and Cisco's 12.8 Tb/s AI Factory Specs

A08 | Operations | Product adoption

Date not stated

FACT

NVIDIA and Cisco have developed Ultra-Dense Small Form Factor Pluggable (USFP) optics to support 400 Gb/s-per-lane connectivity for AI factories. The technology delivers four times the bandwidth density of current 200 Gb/s-per-lane pluggables and improves power efficiency by approximately 1.3 times.

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 19/25 | MATURITY 4/5 | DECISION LOCUS: TECH / PROCESS

P06

Celestica DS6005: 102.4Tbps CPO Switch Specs and Design

A02 | Light source | Product adoption

Date not stated

FACT

Celestica has introduced the DS6005, a liquid-cooled data center switch featuring 102.4Tbps switching capacity and 64 1.6TbE ports. By utilizing co-packaged optics (CPO) and an Open Rack v3 chassis, the system targets power-efficient AI/ML scale-up and scale-out fabrics.

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 19/25 | MATURITY 4/5 | DECISION LOCUS: BUSINESS / RULES

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

P07

TDK Photonic Chip: Voltage-Led Control for Smart-Glass Displays

A05 | Photonic device | Pilot

Date not stated

FACT

TDK Corporation has developed a voltage-led photonic chip using thin-film lithium niobate to improve the efficiency and resolution of laser beam scanning displays. The system achieves a 10 GHz modulation bandwidth and reduces power consumption by over 75% compared to conventional current-based methods. This advancement enables high-resolution 4K imaging within the compact form factor required for wearable smart glasses.

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

Ge/SiGe-AIN Core-Shell Crystals: E-band Emission for Optical Comms

A06 | Packaging | Pilot

Date not stated

FACT

A joint research group from Japan and South Korea has developed Ge/SiGe-AIN core-shell crystals that emit light in the E-band (1360-1460 nm). The team proposed a "shell-first" growth model driven by Plateau-Rayleigh instability to explain the formation of these structures.

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 2/5 | DECISION LOCUS: TECH / PROCESS

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

P09

Photonic Integrated Blue Laser: 461 nm Specs and Tuning Capabilities

A04 | Network system | Pilot

Date not stated

FACT

Researchers have developed a photonic integrated blue laser operating at 461 nm that combines a sub-30-kHz linewidth with high-speed frequency agility. By integrating a GaN-based laser diode with a low-loss silicon nitride microresonator and AlN piezoelectric actuators, the system enables mode-hop-free linear frequency chirps up to 900 MHz. This advancement provides a compact, robust light source for underwater communication, coherent aerosol sensing, and quantum computing.

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 2/5 | DECISION LOCUS: SYSTEM / ADOPTION

P10

Marvell PFMA: Optical Memory Fabric for AI KV Cache Bottlenecks

A11 | Operations | Research

Date not stated

FACT

Marvell has introduced the Photonic Fabric Memory Appliance (PFMA), a rack-mountable system designed to share 16-32 TB of DDR5 memory across up to 16 servers via optical links. By treating memory as a cluster-wide resource, the architecture aims to increase concurrent conversations by approximately six times compared to traditional configurations.

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	NVIDIA OOI Packaging: CPO Roadmap and 2028 Feynman Platform	https://finance.technews.tw/2026/10/20/nvidia-unveils-prototype-silicon-interposer-optical-packaging-chip/
A02	Celestica DS6005: 102.4Tbps CPO Switch Specs and Design	https://www.storagereview.com/news/celestica-ds6005-102-4tbps-co-packaged-optics-1-5ou-liquid-cooled-orv3-switch
A03	Silicon Photonics Interconnect: 3.2 Tbps AI Data Traffic Chip	https://bioengineer.org/silicon-photonics-chip-hits-3-2-terabits-per-second-for-ai-data-traffic/
A04	Photonic Integrated Blue Laser: 461 nm Specs and Tuning Capabilities	https://www.nature.com/articles/s41566-026-02040-9
A05	TDK Photonic Chip: Voltage-Led Control for Smart-Glass Displays	https://optics.org/news/tdk-develops-voltage-led-photonic-chip-to-make-smart-glasses-more-efficient
A06	Ge/SiGe-AlN Core-Shell Crystals: E-band Emission for Optical Comms	https://eetimes.itmedia.co.jp/ee/articles/2610/21/news028.html
A07	GaN MicroLEDs: 2.75 GHz Bandwidth Record and Scaling Roadmap	https://compoundsemiconductor.net/article/125685/275GHz_record_bandwidth_for_GaN_microLED
A08	USFP Optics: NVIDIA and Cisco's 12.8 Tb/s AI Factory Specs	https://www.eenewseurope.com/en/nvidia-and-cisco-unveil-usfp-optics-for-ai-factories/
A09	NVIDIA Rubin Ultra NVL576: CPO-based 576-GPU Scale-up Architecture	https://xenospectrum.com/nvidia-rubin-ultra-nvl576-cpo/
A10	Linktel 1.6T LPO: Performance Specs and AI Network Roadmap	https://www.thailand-business-news.com/pr-news/linktel-technologies-highlights-breakthrough-1-6t-lpo-performance-strong-market-pull-at-ocp-2026
A11	Marvell PFMA: Optical Memory Fabric for AI KV Cache Bottlenecks	https://convergedigest.com/pecc-summit-marvell-optical-memory-kv-cache/
A12	Aegiq Artemis: NVIDIA Ising AI for Quantum Calibration Specs	https://quantumzeitgeist.com/aegiqs-quantum-computer-reduces-engineering/

Photonics — Selected Articles

Date: 2026-10-25


Articles: 12

Table of Contents

- #01 NVIDIA OOI Packaging: CPO Roadmap and 2028 Feynman Platform
- #02 Celestica DS6005: 102.4Tbps CPO Switch Specs and Design
- #03 Silicon Photonics Interconnect: 3.2 Tbps AI Data Traffic Chip
- #04 Photonic Integrated Blue Laser: 461 nm Specs and Tuning Capabilities
- #05 TDK Photonic Chip: Voltage-Led Control for Smart-Glass Displays
- #06 Ge/SiGe-AlN Core-Shell Crystals: E-band Emission for Optical Comms
- #07 GaN MicroLEDs: 2.75 GHz Bandwidth Record and Scaling Roadmap
- #08 USFP Optics: NVIDIA and Cisco's 12.8 Tb/s AI Factory Specs
- #09 NVIDIA Rubin Ultra NVL576: CPO-based 576-GPU Scale-up Architecture
- #10 Linktel 1.6T LPO: Performance Specs and AI Network Roadmap
- #11 Marvell PFMA: Optical Memory Fabric for AI KV Cache Bottlenecks
- #12 Aegiq Artemis: NVIDIA Ising AI for Quantum Calibration Specs

#01 NVIDIA OOI Packaging: CPO Roadmap and 2028 Feynman Platform

Published October 20, 2026 TechNews 科技新報 Taiwan



OVERVIEW

NVIDIA has developed a prototype "Optics-on-Interposer" (OOI) chip that integrates optical engines directly into the silicon interposer to eliminate bandwidth bottlenecks in AI computing. This technology significantly reduces component footprint from 100 mm to 8 mm, paving the way for the 2028 Rosa Feynman platform.

Key Findings

At SEMICON West 2026, NVIDIA presented a prototype chip utilizing Optics-on-Interposer (OOI) technology. This approach moves the Optical Engine (OE) from traditional organic substrates to the silicon interposer, aiming to increase bandwidth density and reduce power consumption for large-scale AI calculations.

Technical Details

The OOI architecture restricts electrical signal transmission within the silicon interposer, utilizing shorter and denser interconnects. This design mirrors the proximity used in current GPU and High Bandwidth Memory (HBM) packaging. NVIDIA provided a comparison using a 200T ASIC package: traditional organic substrate solutions require a component range of 100 mm, whereas OOI concentrates critical components within an 8 mm silicon interposer structure.

Additionally, NVIDIA demonstrated research into dense wavelength-division multiplexing (DWDM). The demonstration utilized 8 optical wavelengths, each operating at 32 Gb/s, achieving a total transmission bandwidth of 256 Gb/s per single optical fiber path.

To mitigate transmission bottlenecks during AI system expansion, NVIDIA integrated microbump technology and low-capacitance interfaces between GPUs and switches.

Background and Context

The OOI technology is the physical realization of theories proposed by NVIDIA Chief Research Scientist Nikola Nedovic in a 2023 paper titled "Beyond CPO: Motivations and Methods for Bringing Optics into the Silicon Interposer." NVIDIA is currently advancing its CPO-related product lineup, which includes BlueField-5, NVLink 8 CPO, Spectrum 7 204T CPO, and CX10.

Significance and Outlook

NVIDIA plans to implement these advancements in the Rosa Feynman solution scheduled for 2028. The Feynman architecture GPU is expected to further integrate CPO optical technology, surpassing the current Rubin Ultra architecture to provide the high-speed transmission foundation required for giant AI computing centers after 2028.

Source: <https://finance.technews.tw/2026/10/20/nvidia-unveils-prototype-silicon-interposer-optical-packaging-chip/>

Collected: October 24, 2026 | Automated Research System (AI)

#02 Celestica DS6005: 102.4Tbps CPO Switch Specs and Design

Published October 15, 2026 StorageReview USA



OVERVIEW

Celestica has introduced the DS6005, a liquid-cooled data center switch featuring 102.4Tbps switching capacity and 64 1.6TbE ports. By utilizing co-packaged optics (CPO) and an Open Rack v3 chassis, the system targets power-efficient AI/ML scale-up and scale-out fabrics.

Key Findings

The Celestica DS6005 is a 1.5OU, fully liquid-cooled switch designed for the OCP high-power rack (ORv3-HPR). It delivers 102.4Tbps of switching capacity via 64 optically enabled 1.6TbE ports. The system is built on the Broadcom Tomahawk 6 Davisson, the CPO variant of the 102.4T Tomahawk 6 ASIC.

Technical Details

The DS6005 integrates optical engines directly onto the switch package, eliminating discrete digital signal processors and long copper traces between the ASIC and pluggable modules to reduce latency and link-level power. Key hardware specifications include:

- **Optical Interface:** 64 1.6TbE ports terminating in 16 MMC-16 fiber connectors.
- **Laser Modules:** 16 field-replaceable ELSFP external laser modules located on the front panel to ensure serviceability.
- **Power and Cooling:** Full liquid cooling with rear-panel quick-disconnects and a DC input labeled for 46 to 56V at 76A.
- **Compute and Security:** An AMD Ryzen x86 host CPU with BMC, TPM 2.0, hardware root of trust, and post-quantum cryptography support.
- **Timing and Software:** IEEE 1588 and SyncE timing hardware; the system runs SONiC with ONIE for installation.

Background and Context

The DS6005 was introduced at the OCP Global Summit alongside other Tomahawk 6 CPO systems, including a 1RU box from Micas Networks and a 2OU switch from Supermicro. Unlike these alternatives, the DS6005 was designed specifically for the OCP high-power rack from the start, offering liquid cooling as the sole cooling option to prevent thermal throttling and enable denser rack configurations.

Significance and Outlook

The shift toward CPO is driven by unprecedented bandwidth demands in AI infrastructure where power and thermal management are critical constraints. While Celestica has not yet published specific power or latency figures, the integration of CPO-aware software management and field-replaceable lasers aims to provide a scalable, power-efficient fabric. According to DellOro Group, pluggable, near-package, and co-packaged optics are expected to coexist as the ecosystem matures.

Source: <https://www.storagereview.com/news/celestica-ds6005-102-4tbps-co-packaged-optics-1-5ou-liquid-cooled-orv3-switch>

Collected: October 24, 2026 | Automated Research System (AI)

#03 Silicon Photonics Interconnect: 3.2 Tbps AI Data Traffic Chip

Published October 18, 2026 Bioengineer.org (論文紹介サイト。原論文はPhotoniX Synergy掲載とされる) China



OVERVIEW

Researchers at Zhejiang University have developed a scalable silicon photonic interconnect architecture achieving an aggregate bandwidth of 3.2 terabits per second in the O-band. By combining hybrid wavelength-port multiplexing with adaptive stabilization, the system addresses the energy and bandwidth bottlenecks of copper-based electrical interconnects in AI clusters.

Key Findings

Researchers at Zhejiang University have reported a silicon photonic interconnect architecture capable of an aggregate bandwidth of 3.2 terabits per second. Published in *PhotoniX Synergy*, the work details a complete transceiver circuit designed for the O-band wavelength window. The system utilizes a combination of hybrid wavelength-port multiplexing, low-crosstalk multiplexing/demultiplexing, and adaptive wavelength stabilization to enable terabit-scale optical I/O for next-generation AI computing.

Technical Details

The architecture employs a heterogeneous design to balance bandwidth and physical footprint:

- **Transmitter:** Uses broadband lattice-filter-based multiplexers consisting of cascaded interferometric stages to ensure wide operating tolerance and controlled passbands.
- **Receiver:** Utilizes compact Vernier dual-microring demultiplexers, which use mismatched free spectral ranges of paired ring resonators to reduce chip area compared to conventional filter banks.
- **Hybrid Multiplexing:** The system scales capacity across two axes: wavelength multiplexing (stacking channels on different colors of light) and port multiplexing (duplicating this across multiple physical waveguide ports).

This dual-axis approach allows aggregate bandwidth to scale as the product of the number of wavelengths per port and the number of ports, preventing excessive channel density on a single waveguide.

Background and Context

Modern AI workloads generate massive data traffic between processors, memory banks, and servers. Traditional electrical interconnects using copper links face severe energy and bandwidth constraints as AI models grow and scale across thousands of accelerators. Silicon photonics mitigates these bottlenecks by encoding data onto light and guiding it through silicon waveguides, which offer lower loss and higher bandwidth density over distance.

Significance and Outlook


According to corresponding author Professor Jianyi Yang, the goal was to create a robust and scalable terabit-scale interconnect. By integrating complementary multiplexer designs and hybrid multiplexing, the team aims to increase capacity without requiring exotic packaging or extreme thermal control. This architecture provides a pathway to scale hardware proportionately with capacity increments, addressing the critical data-movement bottleneck in large-scale AI infrastructure.

Source: <https://bioengineer.org/silicon-photonics-chip-hits-3-2-terabits-per-second-for-ai-data-traffic/>

Collected: October 24, 2026 | Automated Research System (AI)

#04 Photonic Integrated Blue Laser: 461 nm Specs and Tuning Capabilities

Published October 22, 2026 Nature Photonics



OVERVIEW

Researchers have developed a photonic integrated blue laser operating at 461 nm that combines a sub-30-kHz linewidth with high-speed frequency agility. By integrating a GaN-based laser diode with a low-loss silicon nitride microresonator and AlN piezoelectric actuators, the system enables mode-hop-free linear frequency chirps up to 900 MHz. This advancement provides a compact, robust light source for underwater communication, coherent aerosol sensing, and quantum computing.

Key Findings

The reported photonic integrated blue laser operates at 461 nm and achieves a sub-30-kHz linewidth while delivering over 1 mW of optical output power. A primary achievement of this architecture is its frequency agility; the system supports mode-hop-free linear frequency chirps up to 900 MHz at repetition rates up to 1 MHz, maintaining tuning nonlinearity below 2%.

Technical Details

The device utilizes a hybrid integration approach combining a gallium-nitride-based laser diode with a high-Q microresonator. The microresonator is fabricated on a silicon nitride (Si₃N₄) photonic platform using 25-nm-thick waveguides. This specific thickness reduces optical confinement and scattering losses compared to 50-nm-thick waveguides, resulting in a propagation loss of ~ 0.4 dB cm⁻¹ and intrinsic quality factors (Q) exceeding 2.5×10^6 . To enable high-speed modulation of the refractive index and laser frequency tuning, aluminium nitride (AlN) piezoelectric actuators are monolithically integrated onto the photonic circuitry.

Background and Context

Frequency-agile lasers in the ultraviolet-blue spectral range (360–480 nm) are critical for technologies such as optical atomic clocks, Rydberg-atom-based quantum computing, and underwater optical communications. While semiconductor diode lasers (specifically Fabry–Pérot designs) are commonly used, they typically operate in multimode regimes and require complex external free-space cavities or gratings for stabilization. These external components limit portability and environmental robustness. Previous photonic integrated implementations have lacked the fast wavelength-tuning mechanisms necessary for active frequency stabilization via the Pound–Drever–Hall technique and dynamic coherent sensing.

Significance and Outlook

The integration of low-noise performance with fast chirping capabilities addresses a distinct challenge in visible-light photonics. The ability to operate at 461 nm is particularly relevant for the laser cooling and state preparation of strontium (Sr). The researchers have already demonstrated the potential of this integrated laser through experiments in coherent aerosol sensing and underwater communication, where shorter wavelengths offer lower attenuation in water and stronger scattering in air.

Source: <https://www.nature.com/articles/s41566-026-02040-9>

Collected: October 24, 2026 | Automated Research System (AI)

#05 TDK Photonic Chip: Voltage-Led Control for Smart-Glass Displays

Published October 21, 2026 optics.org USA



OVERVIEW

TDK Corporation has developed a voltage-led photonic chip using thin-film lithium niobate to improve the efficiency and resolution of laser beam scanning displays. The system achieves a 10 GHz modulation bandwidth and reduces power consumption by over 75% compared to conventional current-based methods. This advancement enables high-resolution 4K imaging within the compact form factor required for wearable smart glasses.

Key Findings

TDK Corporation has developed a photonic chip that controls laser light intensity via voltage rather than electrical current, significantly reducing power consumption and increasing switching speeds. The researchers report that the system can cut overall power consumption by more than 75% compared to conventional approaches. Furthermore, the chip achieved a modulation bandwidth of 10 gigahertz, which is sufficient to support 4K resolution and higher.

Technical Details

The device is fabricated using a thin film of lithium niobate, a material with electrically controllable optical properties. The architecture consists of the following components:

- Three Mach–Zehnder interferometers: One for each primary color (red, green, and blue), where applied voltage regulates the intensity of light passing through each channel.
- Multiplexer: Formed from lithium niobate waveguides to combine the three colors into a single output beam.
- MEMS scanning mirror: Integrated with the device to demonstrate the production of full-color images.

By utilizing a photonic circuit for intensity control, the design eliminates the need for neutral-density filters typically used to reduce excess laser output.

Background and Context

Conventional laser beam scanning displays produce images by steering a laser beam across a field of view. These systems typically control color intensity through direct current modulation of laser diodes. However, direct modulation has limited bandwidth compared to external optical modulation, restricting imaging resolution. Additionally, laser diodes must operate above a threshold current to maintain intensity, which increases power draw.

Significance and Outlook


This voltage-controlled approach addresses the interconnected challenges of switching speed, power consumption, and component count in the miniaturization of laser scanning displays. The 10-GHz bandwidth provides experimental evidence that external voltage control can support ultracompact, high-resolution wearable displays. Future research will focus on reducing optical losses within the lithium niobate waveguides and during the color combination process to further increase overall system efficiency.

Source: <https://optics.org/news/tdk-develops-voltage-led-photonic-chip-to-make-smart-glasses-more-efficient>

Collected: October 24, 2026 | Automated Research System (AI)

#06 Ge/SiGe-AlN Core-Shell Crystals: E-band Emission for Optical Comms

Published October 21, 2026 EE Times Japan Japan



OVERVIEW

A joint research group from Japan and South Korea has developed Ge/SiGe-AlN core-shell crystals that emit light in the E-band (1360–1460 nm). The team proposed a "shell-first" growth model driven by Plateau-Rayleigh instability to explain the formation of these structures.

Key Findings

A joint research team comprising Nagoya University, the Electronics and Telecommunications Research Institute (ETRI), Kookmin University, and Korea National Ocean University has confirmed that Ge/SiGe-AlN core-shell crystals emit light within the E-band (1360–1460 nm), a critical spectrum for optical fiber communications. Room-temperature photoluminescence measurements identified specific emission peaks at 1372 nm, 1377 nm, 1381 nm, and 1392 nm.

Technical Details

The researchers combined Ge (germanium)/SiGe (silicon germanium) semiconductors with AlN (aluminum nitride) to create core-shell crystals, where a central core is surrounded by a different material shell. To analyze the resulting structures, the group utilized X-ray CT and electron microscopy. These observations revealed various morphologies, including structures with internal cavities and those where Ge/SiGe materials were partially incorporated into the interior.

Based on these results, the team proposed a "shell-first" growth model. Unlike conventional processes where the core forms first, this model suggests that the AlN-containing shell structure forms first, subsequently incorporating the Ge/SiGe materials. The researchers attribute this unique growth mechanism to the influence of Plateau-Rayleigh instability.

Background and Context

Core-shell crystals are utilized to combine semiconductors with different properties to enable new optical and electronic devices. However, the specific processes of core and shell formation, as well as their resulting optical characteristics, had not been sufficiently elucidated until this study.

Significance and Outlook


The ability to achieve E-band emission using Ge/SiGe-AlN materials provides a potential new material path for optical communication components. The research was conducted by a diverse team including Choi Sunghoon (ETRI), Park Jun-Seok and Won Joon Bong (Kookmin University), Lee Jae Hak and Kim Kyoung Hwa (LNBS), Ahn Hyung-Soo, Kwon Eunmin, Kim Sohee, Kim Myungjun, and Jeong Donghyeon (Korea National Ocean University), as well as Anna Honda, Yoshio Honda, and Tsuyoshi Kato from Nagoya University's Institute for Materials and Systems.

Source: <https://eetimes.itmedia.co.jp/ee/articles/2610/21/news028.html>

Collected: October 24, 2026 | Automated Research System (AI)

#07 GaN MicroLEDs: 2.75 GHz Bandwidth Record and Scaling Roadmap

Published October 22, 2026 Compound Semiconductor USA



OVERVIEW

Researchers at the University of Illinois at Urbana-Champaign have achieved a world-record modulation bandwidth of 2.75 GHz in a blue GaN microLED. By analyzing devices from 100 μ m down to 2 μ m, the team established a device-level roadmap to replace copper links with high-bandwidth optical interconnects for AI computing.

Key Findings

A research team led by Can Bayram has set a world record for modulation bandwidth in a blue GaN microLED, reaching 2.75 GHz. This performance was achieved using a device with a 5 μ m diameter, which also demonstrated stability and linearity under extreme current densities. The findings, published in IEEE Electron Device Letters, provide a technical roadmap for integrating these emitters into high-bandwidth optical interconnects.

Technical Details

The team fabricated and tested blue microLEDs across nearly two orders of magnitude in size, with diameters ranging from 100 μ m down to 2 μ m. Technical analysis using Smith-chart measurements of RF power delivery revealed that standard, larger microLEDs are primarily limited by junction capacitance. Reducing the device diameter significantly lowered this capacitance, enabling the record-breaking speeds.

However, the study identified a performance ceiling below the 5 μ m threshold. Further miniaturization increases current-spreading and contact resistances, leading to an impedance mismatch with standard 50- Ω radio-frequency systems. The researchers identified p-contact optimization and pad engineering as the necessary next steps to overcome these resistances and further increase performance.

Background and Context

Optical interconnects are being explored to address the energy and bandwidth bottlenecks of conventional electrical copper links in data centers and high-performance computing. MicroLED arrays are considered strong candidates for these applications due to their compact footprint and potential for scalable, massively parallel, and error-free optical data links.

Significance and Outlook

By unifying measurements of impedance, linearity, and modulation across a vast scaling range, the researchers have established a clear pathway for creating dense, low-energy optical channels. This development aims to reduce packaging complexity in future optical networks, potentially breaking the data bottleneck for next-generation AI and high-performance computing through massively parallel optical communication.


Source:

https://compoundsemiconductor.net/article/125685/275GHz_record_bandwidth_for_GaN_microLED

Collected: October 24, 2026 | Automated Research System (AI)

#08 USFP Optics: NVIDIA and Cisco's 12.8 Tb/s AI Factory Specs

Published October 21, 2026 eeNews Europe UK



OVERVIEW

NVIDIA and Cisco have developed Ultra-Dense Small Form Factor Pluggable (USFP) optics to support 400 Gb/s-per-lane connectivity for AI factories. The technology delivers four times the bandwidth density of current 200 Gb/s-per-lane pluggables and improves power efficiency by approximately 1.3 times.

Key Findings

NVIDIA and Cisco have introduced the Ultra-Dense Small Form Factor Pluggable (USFP), a new optical transceiver form factor designed for next-generation AI factories. The USFP architecture supports a total capacity of up to 12.8 Tb/s and provides four times the bandwidth and front-panel bandwidth density compared to existing 200 Gb/s-per-lane pluggable modules.

Technical Details

The initial USFP transceivers integrate 16 full-duplex lanes within an OSFP-sized footprint, compatible with both SuperNICs and switches, with potential future expansions to 32 lanes or more. A primary innovation is the two-dimensional array connector featuring a two-step insertion mechanism; high-speed electrical contacts engage only after the module is fully inserted to reduce wear and improve signal integrity.

- **Performance:** Supports electrical signaling at 400 Gb/s per lane and improves performance at 200 Gb/s.
- **Thermal Management:** Modules can dissipate up to 100W without requiring a dedicated embedded cooling circuit.
- **Efficiency:** Offers approximately 1.3 times better power efficiency than previous generations.

Background and Context

The development targets the reduction of network complexity and power consumption in AI infrastructure. By increasing density, operators can connect more computing systems within the same footprint, avoiding additional switching layers that typically increase latency and costs. The initiative builds upon existing collaborations between Cisco Silicon One-based networks and NVIDIA Spectrum-X Ethernet networking.

Significance and Outlook

To ensure interoperability, NVIDIA and Cisco are partnering with Amphenol and TE Connectivity to develop USFP connectors, cages, and modules. The companies intend to establish an open industry specification via a community-led multi-source project. This open ecosystem is intended to provide European data center operators, network equipment makers, and optical component suppliers a flexible path to higher-bandwidth networking while maintaining the serviceability of pluggable optics for enterprise, sovereign cloud, and neocloud deployments.

Source: <https://www.eenewseurope.com/en/nvidia-and-cisco-unveil-usfp-optics-for-ai-factories/>

Collected: October 24, 2026 | Automated Research System (AI)

#09 NVIDIA Rubin Ultra NVL576: CPO-based 576-GPU Scale-up Architecture

Published October 22, 2026 XenoSpectrum USA



OVERVIEW

NVIDIA presented the Vera Rubin Ultra NVL576 at OCP 2026, a configuration that links 576 GPUs into a single NVLink domain using Co-Packaged Optics (CPO). This architecture targets AI models with over 10 trillion parameters, aiming to improve power efficiency by expanding high-speed connectivity across eight racks.

Key Findings

At the OCP Global Summit held in San Jose starting October 12, 2026, NVIDIA detailed the "Vera Rubin Ultra NVL576 CPO" configuration. The system utilizes an NVLink 7 CPO switch tray to unify 576 GPUs into a single high-speed connectivity domain. This design is specifically optimized for AI models exceeding 10 trillion parameters, focusing on enhancing performance per watt.

Technical Details

The NVL576 consists of eight MGX NVL racks, with each rack containing 72 Rubin Ultra GPUs. To create a single NVLink domain, the system employs a two-layer all-to-all connection topology using a combination of copper wiring and direct optical connections. A critical component is the Co-Packaged Optics (CPO) switch tray, which integrates optical engines within the same package as the semiconductor chips to reduce electrical wiring length and mitigate power consumption and signal transmission constraints.

NVIDIA distinguishes this scale-up NVLink architecture from its scale-out Ethernet networking. While the NVL576 connects 576 GPUs in a tight domain, the "Spectrum-X Photonics Multiplane Ethernet" is designed for AI factories scaling up to 512K GPUs. Additionally, the Rubin Ultra design is integrated into the MGX 1.2 architecture, which standardizes liquid cooling and power interfaces to simplify assembly and maintenance.

Background and Context

The 576-GPU concept was previously introduced in a March technical briefing, which mentioned "Polyphe," a prototype using GB200 GPUs. However, the OCP presentation focuses on the Rubin Ultra production path. It is noted that the NVL576 is distinct from the "Kyber NVL144," which houses 144 GPUs in a single rack; the NVL576 requires eight racks to reach its total GPU count.

Significance and Outlook


While the CPO configuration was presented, industry reports suggest NVIDIA may still consider Near-Packaged Optics (NPO), with a potential product appearance in 2027. The real-world utility of the NVL576 will depend on actual performance measurements under unified conditions and the maintainability of the optical components, specifically how the replacement of failed optical engines affects system uptime.

Source: <https://xenospectrum.com/nvidia-rubin-ultra-nvl576-cpo/>

Collected: October 24, 2026 | Automated Research System (AI)

#10 Lintel 1.6T LPO: Performance Specs and AI Network Roadmap

Published October 18, 2026 Thailand Business News (PR Newswire配信の転載) USA



OVERVIEW

Linktel Technologies demonstrated 1.6T Linear Pluggable Optics (LPO) at OCP 2026, featuring power consumption as low as 10W and ultra-low latency. The company has moved these modules into volume production to address bandwidth and thermal constraints in AI data centers.

Key Findings

Linktel Technologies showcased its 1.6T LPO solutions, specifically the 1.6T LPO 2DR4, at the 2026 OCP Global Summit. These modules are designed for next-generation AI computing networks, focusing on reducing power loads and signal processing delays. The company confirmed that volume production for these 1.6T connectivity solutions is already in place.

Technical Details

The 1.6T LPO architecture provides several technical advantages for high-performance AI scale-out fabrics:

- **Power Efficiency:** Operates with power consumption as low as 10W.
- **Latency:** Utilizes a linear optics architecture to minimize signal processing delays.
- **Supply Chain:** Removes the optical DSP to eliminate component choke points.
- **Interoperability:** Validated multi-switch platform interoperability during the 1.6T LPO-MSA Plugfest in early September 2026, supporting the 200G-per-wavelength ecosystem.

Other Architectural Innovations

Beyond LPO, Linktel presented additional high-speed connectivity pillars:

- **12.8T Liquid-Cooled XPO:** Integrated with advanced liquid cooling for high-density AI nodes. As a founding member of the XPO MSA, Linktel is scaling XPO series variants including FRO, LRO, and LPO.
- **6.4T/7.2T NPO Modules:** Supports both ELS and ILS configurations, combining high density with low power and latency.
- **DMD Technology:** Dr. Grace Zhang presented Dispersion Managed Transceivers for OCS SuperPOD architectures, specifically cost-effective 800G (20km) and 1.6T (10km) solutions, and the feasibility of 400G per lambda (3km).

Significance and Outlook


The DMD product line has entered volume production. To accelerate the maturity of next-gen AI fabrics, Linktel has stated it is open to sharing its U.S.-filed patents with ecosystem partners. These developments target the scaling pressures of AI data centers that require higher bandwidth densities within strict thermal and power budgets.

Source: <https://www.thailand-business-news.com/pr-news/linktel-technologies-highlights-breakthrough-1-6t-lpo-performance-strong-market-pull-at-ocp-2026>

Collected: October 24, 2026 | Automated Research System (AI)

#11 Marvell PFMA: Optical Memory Fabric for AI KV Cache Bottlenecks

Published October 22, 2026 Converge Digest USA



OVERVIEW

Marvell has introduced the Photonic Fabric Memory Appliance (PFMA), a rack-mountable system designed to share 16–32 TB of DDR5 memory across up to 16 servers via optical links. By treating memory as a cluster-wide resource, the architecture aims to increase concurrent conversations by approximately six times compared to traditional configurations.

Key Findings

Marvell is targeting the key-value (KV) cache bottleneck in AI inference to prevent redundant computations during long conversations in reasoning models and agentic applications. Simulation-based analysis indicates that a shared-memory configuration can support approximately six times as many concurrent conversations at a comparable time-to-first-token using the same compute resources. These results are based on workload-specific modeling and pre-silicon estimates.

Technical Details

The PFMA is a rack-mountable appliance featuring 128 standard DDR5 DIMM slots and approximately 1 TB of HBM3e, which functions as a cache for the larger DDR5 pool. The system utilizes sixteen Photonic Fabric Memory Modules connected through a passive fiber shuffle to create a unified shared address space. Marvell provides two integration options:

- **PF-NIC card:** Uses PCIe Gen6 x16/CXL 3.x for server connectivity.
- **PF-Chiplet:** Integrated directly into an XPU package.

The practical optical reach for memory sharing across racks is identified as 25–30 meters to balance physical separation against propagation delay. Marvell positions the chiplet-attached PFMA as a G1.5 tier below local GPU HBM, while the server-attached version serves as a G2.5 tier between system DRAM and storage.

Background & Context

The technology, acquired through Celestial AI in February 2026, expands Marvell's memory strategy beyond Bravera storage controllers and Structera CXL expansion. The goal is to keep reusable KV cache closer to compute than SSD-based offloading, allowing inference software to flexibly assign workloads to different processors without repeating prefill work.

Significance & Outlook


This architecture shifts KV cache placement from a local component to a cluster architecture decision. However, realized performance remains dependent on cache hit rates, DDR bandwidth, fiber length, and the ability of serving software to manage compatible cache entries. The PFMA represents a move toward treating memory as a shared cluster resource rather than a standalone server component.

Source: <https://convergedigest.com/pecc-summit-marvell-optical-memory-kv-cache/>

Collected: October 24, 2026 | Automated Research System (AI)

#12 Aegiq Artemis: NVIDIA Ising AI for Quantum Calibration Specs

Published October 22, 2026 Quantum Zeitgeist UK



OVERVIEW

Aegiq has integrated NVIDIA Ising AI models into the calibration workflow of its Artemis photonic quantum computer, reducing weekly engineering time threefold. The system utilizes a Vision-Language Model (VLM) to automate the optimization of single photon source performance, enabling calibration by non-specialists.

Key Findings

Aegiq's Artemis photonic quantum computer, deployed at the UK's National Quantum Computing Centre (NQCC), has achieved a threefold reduction in weekly engineering time by integrating NVIDIA Ising AI models into its calibration workflow. This automation allows the system to maintain high performance while reducing the reliance on specialist expertise for routine calibration tasks.

Technical Details

The system employs the NVIDIA Ising Calibration Vision-Language Model (VLM), which runs locally on an NVIDIA DGX Spark to ensure the low-latency inference required for real-time control. The VLM automatically balances three key performance metrics for the single photon source: brightness, purity, and indistinguishability.

Aegiq utilizes an agent-based architecture to manage the process. Users interact with the system via natural language prompts, such as "Check the performance of the single photon source" or requests to find optimal values for laser power, wavelength, and source bias. The agents autonomously develop and execute calibration plans, iteratively analyzing results to optimize performance. The VLM's pre-existing training across multiple quantum modalities allowed for immediate application to these tasks without additional specialized training.

Background & Context

Previously, calibration was a manual process that demanded significant engineering time and limited overall system uptime. By shifting to an AI-driven approach, Aegiq aims to increase stability and uptime for its customers. The current deployment on the Artemis system has successfully achieved target performance metrics for this first-generation quantum computer.

Significance & Outlook

Aegiq is expanding the use of NVIDIA Ising technology beyond calibration; the company is exploring NVIDIA Ising decoding for quantum error correction within its QGATE architecture. This dual application is part of a broader strategy to redesign calibration routines and workflows around an AI-enhanced paradigm. These developments are intended to prepare the system for the scaling and implementation of future fault-tolerant quantum computers.

Source: <https://quantumzeitgeist.com/aegiqs-quantum-computer-reduces-engineering/>

Collected: October 24, 2026 | Automated Research System (AI)