

NVIDIA Announces Spectrum-X Photonics, Co-Packaged Optics Networking Switches to Scale AI Factories to Millions of GPUs

- **1.6 Terabits Per Second Per Port Switches to Deliver 3.5x Energy Savings and 10x Resilience in AI Factories**
- **Joint Inventions and Collaborations With TSMC, Coherent, Corning Incorporated, Foxconn, Lumentum and SENKO to Create Integrated Silicon, Optics Process and Supply Chain**

GTC—NVIDIA today unveiled NVIDIA Spectrum-X™ and NVIDIA Quantum-X silicon photonics networking switches, which enable AI factories to connect millions of GPUs across sites while drastically reducing energy consumption and operational costs. NVIDIA has achieved the fusion of electronic circuits and optical communications at massive scale.

As AI factories grow to unprecedented sizes, networks must evolve to keep pace. NVIDIA photonics switches are the world's most advanced networking solution. They integrate optics innovations with 4x fewer lasers to deliver 3.5x more power efficiency, 63x greater signal integrity, 10x better network resiliency at scale and 1.3x faster deployment compared with traditional methods.

"AI factories are a new class of data centers with extreme scale, and networking infrastructure must be reinvented to keep pace," said Jensen Huang, founder and CEO of NVIDIA. "By integrating silicon photonics directly into switches, NVIDIA is shattering the old limitations of hyperscale and enterprise networks and opening the gate to million-GPU AI factories."

NVIDIA silicon photonics networking switches are available as part of the NVIDIA Spectrum-X Photonics Ethernet and NVIDIA Quantum-X Photonics InfiniBand platforms.

The Spectrum-X Ethernet networking platform delivers superior performance and 1.6x bandwidth density compared with traditional Ethernet for multi-tenant, hyperscale AI factories, including [the largest supercomputer in the world](#).

NVIDIA Spectrum-X Photonics switches include multiple configurations, including 128 ports of 800Gb/s or 512 ports of 200Gb/s, delivering 100Tb/s total bandwidth, as well as 512 ports of 800Gb/s or 2,048 ports of 200Gb/s, for a total throughput of 400Tb/s.

NVIDIA Quantum-X Photonics switches provide 144 ports of 800Gb/s InfiniBand based on 200Gb/s SerDes and use a liquid-cooled design to efficiently cool the onboard silicon photonics. NVIDIA Quantum-X Photonics switches offer 2x faster speeds and 5x higher scalability for AI compute fabrics compared with the previous generation.

A Networked Ecosystem

NVIDIA's silicon photonics ecosystem includes TSMC, Browave, Coherent, [Corning Incorporated](#), Fabrinet, Foxconn, Lumentum, SENKO, SPIL, Sumitomo Electric Industries and TFC Communication.

"A new wave of AI factories requires efficiency and minimal maintenance to achieve the scale required for next-generation workloads," said C. C. Wei, chairman and CEO of TSMC. "TSMC's silicon photonics solution combines our strengths in both cutting-edge chip manufacturing and TSMC-SoIC 3D chip stacking to help NVIDIA unlock an AI factory's ability to scale to a million GPUs and beyond, pushing the boundaries of AI."

NVIDIA photonics will drive massive growth for a new wave of state-of-the-art AI factories, alongside pluggable optical transceiver technologies supported by industry leaders including Coherent, Eoptolink, Fabrinet and Innolight.

Availability

NVIDIA Quantum-X Photonics InfiniBand switches are expected to be available later this year, with NVIDIA Spectrum-X Photonics Ethernet switches coming in 2026 from leading infrastructure and system vendors.

Learn more by watching the [NVIDIA GTC keynote](#) and [register for sessions](#) from NVIDIA and industry leaders at the show, which runs through March 21.

About NVIDIA

[NVIDIA](#) (NASDAQ: NVDA) is the world leader in accelerated computing.

Certain statements in this press release including, but not limited to, statements as to: the benefits, impact, availability, and performance of NVIDIA's products, services, and technologies; third parties adopting or offering NVIDIA's products and technologies; and by integrating silicon photonics directly into switches, NVIDIA is shattering the old limitations of hyperscale

and enterprise networks and opening the gate to million-GPU AI factories are forward-looking statements that are subject to risks and uncertainties that could cause results to be materially different than expectations. Important factors that could cause actual results to differ materially include: global economic conditions; our reliance on third parties to manufacture, assemble, package and test our products; the impact of technological development and competition; development of new products and technologies or enhancements to our existing product and technologies; market acceptance of our products or our partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of our products or technologies when integrated into systems; as well as other factors detailed from time to time in the most recent reports NVIDIA files with the Securities and Exchange Commission, or SEC, including, but not limited to, its annual report on Form 10-K and quarterly reports on Form 10-Q. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

© 2025 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo and NVIDIA Spectrum-X are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated. Features, pricing, availability and specifications are subject to change without notice.

Alex Shapiro
Enterprise Networking
1-415-608-5044
ashapiro@nvidia.com