



NVIDIA Introduces Spectrum-XGS Ethernet to Connect Distributed Data Centers Into Giga-Scale AI Super-Factories

CoreWeave to Deploy NVIDIA Spectrum-XGS Ethernet With Scale-Across Technology

Hot Chips—NVIDIA today announced NVIDIA® Spectrum-XGS Ethernet, a scale-across technology for combining distributed data centers into unified, giga-scale AI super-factories.

As AI demand surges, individual data centers are reaching the limits of power and capacity within a single facility. To expand, data centers must scale beyond any one building, which is limited by off-the-shelf Ethernet networking infrastructure with high latency and jitter and unpredictable performance.

Spectrum-XGS Ethernet is a breakthrough addition to the [NVIDIA Spectrum-X™](#) Ethernet platform that removes these boundaries by introducing scale-across infrastructure. It serves as a third pillar of AI computing beyond scale-up and scale-out, designed for extending the extreme performance and scale of Spectrum-X Ethernet to interconnect multiple, distributed data centers to form massive AI super-factories capable of giga-scale intelligence.

“The AI industrial revolution is here, and giant-scale AI factories are the essential infrastructure,” said Jensen Huang, founder and CEO of NVIDIA. “With NVIDIA Spectrum-XGS Ethernet, we add scale-across to scale-up and scale-out capabilities to link data centers across cities, nations and continents into vast, giga-scale AI super-factories.”

Spectrum-XGS Ethernet is fully integrated into the Spectrum-X platform, featuring algorithms that dynamically adapt the network to the distance between data center facilities.

With advanced, auto-adjusted distance congestion control, precision latency management and end-to-end telemetry, Spectrum-XGS Ethernet nearly doubles the performance of the [NVIDIA Collective Communications Library](#), accelerating multi-GPU and multi-node communication to deliver predictable performance across geographically distributed AI clusters. As a result, multiple data centers can operate as a single AI super-factory, fully optimized for long-distance connectivity.

Hyperscale pioneers embracing the new infrastructure include CoreWeave, which will be among the first to connect its data centers with Spectrum-XGS Ethernet.

“CoreWeave’s mission is to deliver the most powerful AI infrastructure to innovators everywhere,” said Peter Salanki, cofounder and chief technology officer of CoreWeave. “With NVIDIA Spectrum-XGS, we can connect our data centers into a single, unified supercomputer, giving our customers access to giga-scale AI that will accelerate breakthroughs across every industry.”

The Spectrum-X Ethernet networking platform provides greater bandwidth density than off-the-shelf Ethernet for multi-tenant, hyperscale AI factories — including the [world’s largest AI supercomputer](#). It comprises NVIDIA Spectrum-X switches and [NVIDIA ConnectX®-8 SuperNICs](#), delivering seamless scalability, ultralow latency and breakthrough performance for enterprises building the future of AI.

Today’s announcement follows a drumbeat of networking innovation announcements from NVIDIA, including [NVIDIA Spectrum-X](#) and [NVIDIA Quantum-X silicon photonics networking switches](#), which enable AI factories to connect millions of GPUs across sites while reducing energy consumption and operational costs.

Availability

NVIDIA Spectrum-XGS Ethernet is available now as part of the NVIDIA Spectrum-X Ethernet platform.

Learn more about Spectrum-XGS Ethernet at [Hot Chips](#).

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: with NVIDIA Spectrum-XGS Ethernet, NVIDIA adding scale-across to scale-up and scale-out capabilities to link data centers across cities, nations and continents into vast, giga-scale AI super-factories; the benefits, impact, performance, and availability of NVIDIA’s products, services, and technologies; expectations with respect to NVIDIA’s third party arrangements, including with its collaborators and partners; expectations with respect to technology developments; and other statements that are not historical facts are forward-looking statements within the meaning of Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended, which are subject to the “safe harbor” created by those sections based on

management's beliefs and assumptions and on information currently available to management and 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 and political conditions; NVIDIA's reliance on third parties to manufacture, assemble, package and test NVIDIA's products; the impact of technological development and competition; development of new products and technologies or enhancements to NVIDIA's existing product and technologies; market acceptance of NVIDIA's products or NVIDIA's partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of NVIDIA's products or technologies when integrated into systems; and changes in applicable laws and regulations, 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, ConnectX and 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