



NVIDIA Spectrum-X Ethernet Switches Speed Up Networks for Meta and Oracle

Hyperscalers Broaden Adoption of NVIDIA Networking Solutions to Drive Giga-Scale AI Data Center Performance

News Summary:

- Meta to introduce switches built on NVIDIA Spectrum Ethernet for its Facebook Open Switching System platform.
- Oracle to build giga-scale AI supercomputers with Spectrum-X Ethernet switches.

OCP—NVIDIA today announced that Meta and Oracle will boost their AI data center networks with [NVIDIA Spectrum-X™](#) Ethernet networking switches.

Meta and Oracle are utilizing Spectrum-X Ethernet switches within their ecosystem as an open, accelerated networking architecture that speeds deployment at scale, unlocks exponential gains in AI training efficiency and shortens time to insights.

“Trillion-parameter models are transforming data centers into giga-scale AI factories, and industry leaders like Meta and Oracle are standardizing on Spectrum-X Ethernet to drive this industrial revolution,” said Jensen Huang, founder and CEO of NVIDIA. “Spectrum-X is not just faster Ethernet — it’s the nervous system of the AI factory, enabling hyperscalers to connect millions of GPUs into a single giant computer to train the largest models ever built.”

Oracle will build giga-scale AI factories accelerated by the NVIDIA Vera Rubin architecture and interconnected by Spectrum-X Ethernet.

“Oracle Cloud Infrastructure is designed from the ground up for AI workloads, and our partnership with NVIDIA extends that AI leadership,” said Mahesh Thiagarajan, executive vice president of Oracle Cloud Infrastructure. “By adopting Spectrum-X Ethernet, we can interconnect millions of GPUs with breakthrough efficiency so our customers can more quickly train, deploy and benefit from the next wave of generative and reasoning AI.”

Meta will integrate Spectrum Ethernet switches into its networking infrastructure for the Facebook Open Switching System (“FBOSS”), a software platform developed to manage and control network switches at massive scale. This integration will speed deployment at scale to unlock gains in AI training efficiency and shorten time to insights.

“Meta’s next-generation AI infrastructure requires open and efficient networking at a scale the industry has never seen before,” said Gaya Nagarajan, vice president of networking engineering at Meta. “By integrating NVIDIA Spectrum Ethernet into the Minipack3N switch and FBOSS, we can extend our open networking approach while unlocking the efficiency and predictability needed to train ever-larger models and bring generative AI applications to billions of people.”

NVIDIA Spectrum-X Ethernet Platform

Designed for the trillion-parameter model era, the NVIDIA Spectrum-X Ethernet platform, consisting of Spectrum-X Ethernet switches and Spectrum-X Ethernet SuperNICs, is the first Ethernet platform purpose-built for AI, enabling hyperscalers to interconnect millions of GPUs with unprecedented efficiency and scale.

Trillion-parameter models and generative AI are redefining the scale of data centers. Spectrum-X Ethernet enables AI at scale, delivering the performance and scalability needed to build the world’s most advanced AI infrastructure.

Spectrum-X Ethernet has already demonstrated record-setting efficiency, enabling [the world’s largest AI supercomputer](#) to achieve 95% data throughput with its congestion-control technology. By contrast, off-the-shelf Ethernet at scale suffers from thousands of flow collisions, limiting throughput to roughly 60%.

This leap in efficiency marks a breakthrough in the economics and performance of AI-scale networking. NVIDIA Spectrum-XGS Ethernet technology, part of the Spectrum-X Ethernet networking platform, enables scale-across capabilities to link data centers across cities, nations and continents into vast, giga-scale AI super-factories.

Spectrum-X builds on NVIDIA’s full-stack platform — including GPUs, CPUs, NVIDIA NVLink™ and software — to deliver seamless performance from compute to network. Its advanced congestion control, adaptive routing and AI-driven telemetry capabilities ensure efficiency and predictability for massive AI training and inference clusters.

About NVIDIA

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

Certain statements in this press release including, but not limited to, statements as to: AI being the defining technology of our time, and the world's largest companies building AI factories to drive this industrial revolution; 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, NVIDIA Spectrum-X and NVLink 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