NVIDIA Expands Vera CPU Platform to Accelerate Scientific Discovery

Vera Delivers Higher Performance for Agentic Science and 2x Greater Energy Efficiency Than x86 CPUs

News Summary:

- The NVIDIA Vera CPU rack-scale platform combines NVIDIA processors, NVIDIA Quantum-X800 InfiniBand networking and breakthrough scientific software to boost weather, climate and molecular simulations.
- DKRZ has selected Vera CPUs for its HLRE-5 supercomputer for climate science.
- The full Vera CPU platform is coming from Dell Technologies, HPE, Lenovo and Supermicro.

NVIDIA today announced the NVIDIA Vera CPU rack-scale platform for high-performance computing -- combining NVIDIA's first fully custom data center CPU with the gold standard for HPC networking to drive extreme efficiency and throughput for state-of-the-art scientific modeling, simulation and agentic science.

The rack-scale platform brings together 192 Vera CPUs interconnected by NVIDIA Quantum-X800 InfiniBand networking to meet the massive performance demands of climate modeling, molecular dynamics and other unaccelerated, next-generation scientific workloads.

The liquid-cooled system delivers more than 200TB/s of aggregate peak memory bandwidth to support scientific simulation and AI workflows.

“Scientific simulations need CPUs that can keep large datasets moving and sustain performance as applications scale," said Ian Buck, vice president of hyperscale and high-performance computing at NVIDIA. “Vera brings custom CPU cores, high-bandwidth memory and InfiniBand networking together to address those bottlenecks. With compilers and libraries optimized for Vera, researchers can run demanding simulations more efficiently and explore more complex scientific problems."

Engineered for Memory-Intensive HPC
Each Vera CPU features 88 custom NVIDIA Olympus cores, 176 threads and compact, serviceable SOCAMM2 LPDDR5X memory delivering up to 1.2TB/s of bandwidth to accelerate data-intensive scientific applications.

The rack combines 192 Vera CPUs with approximately 300TB of memory and 1 petabyte of storage in a fully liquid-cooled NVIDIA MGX™ design. Each 1U tray integrates eight Vera CPUs, and NVMe storage and management.

The platform reuses NVL72 rack components, including chassis, power distribution and cooling infrastructure, to simplify deployment and service. NVIDIA Quantum-X800 InfiniBand networking connects the system, and NVIDIA SHARP™ In-Network Computing helps scientists finish complex simulations faster by reducing the time processors spend coordinating their work.

In NVIDIA's tests, Vera delivered 1.9x the performance and 2.4x the performance per watt of x86 systems on OpenFOAM's motorBike Large airflow simulation.

Advancing Climate Science Research
NVIDIA has benchmarked the open source Weather Research and Forecasting model on Vera CPUs, with projected performance gains of 1.6x and energy-efficiency gains of 2x compared with x86.

DKRZ, the German Climate Computing Center, has selected Vera CPUs for its HLRE-5 supercomputer, alongside NVIDIA GB200 Grace Blackwell systems and NVIDIA Quantum-2 InfiniBand networking. Built by Bull and expected to be available in early 2028, the system will support climate simulation and AI research, enabling global kilometer-scale climate simulations with lower energy consumption than on the current system.

Software and Availability
To help researchers speed scientific applications, NVIDIA HPC SDK 26.11 is available now, providing compilers and libraries optimized for Vera.

NVIDIA is also introducing a public beta of its next-generation Fortran compiler, which was built in collaboration with leading HPC centers and national labs. The new compiler gives developers a way to run their existing programs built on Fortran -- the language that scientific computing was built on -- with greater speed on NVIDIA platforms, without needing to rewrite their code.

Vera CPU racks are expected to be available from Dell Technologies, HPE, Lenovo and Supermicro beginning in early

2027.

Learn more about the NVIDIA Vera CPU. For detailed benchmark results and test configurations, read the NVIDIA Vera CPU HPC performance brief.

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: the expected capabilities, benefits and availability of the NVIDIA Vera CPU rack and related products and software; the planned Fortran compiler public beta; and anticipated scientific computing uses and future systems 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.

Many of the products and features described herein remain in various stages and will be offered on a when-and-if-available basis. The statements above are not intended to be, and should not be interpreted as a commitment, promise, or legal obligation, and the development, release, and timing of any features or functionalities described for our products is subject to change and remains at the sole discretion of NVIDIA. NVIDIA will have no liability for failure to deliver or delay in the delivery of any of the products, features or functions set forth herein.

© 2026 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, and other NVIDIA product and service names are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries and regions. Other company, product, and service names may be trademarks of the respective companies with which they are associated.

Alex Shapiro
Corporate Communications
NVIDIA Corporation
press@nvidia.com