



SpaceXAI Adopts NVIDIA Vera CPU to Accelerate Agentic AI at Massive Scale

News Summary:

- SpaceXAI will deploy NVIDIA Vera CPUs to accelerate the work behind its next generation of agentic AI workloads.
- SpaceXAI is expanding its AI infrastructure for Grok with the NVIDIA Vera Rubin platform as it scales toward gigawatts of computing capacity.
- SpaceXAI plans to extend its use of NVIDIA accelerated computing into space, with the first-generation Starmind AI satellite based on an optimized NVIDIA Vera Rubin NVL72 system.

NVIDIA today announced that SpaceXAI will deploy [NVIDIA Vera CPUs](#) to accelerate its next generation of agentic AI applications, bringing the first CPU built for AI agents to one of the world's most ambitious AI deployments.

Agentic AI applications increasingly rely on CPUs to orchestrate tools, execute code, process data and run simulations between model calls. SpaceXAI will use Vera to accelerate these application workloads, helping AI agents act faster while keeping GPUs fed and fully utilized.

SpaceXAI plans to expand its AI infrastructure behind Grok on [NVIDIA Vera Rubin](#), while extending an optimized Vera Rubin NVL72 into space with its first-generation Starmind satellite.

"Agentic AI requires a new kind of computing system -- one built not only to generate answers, but to take action," said Ian Buck, vice president of hyperscale and high-performance computing at NVIDIA. "Vera gives AI agents the CPU performance to act in real time -- executing code, processing data and coordinating complex tasks. SpaceXAI is taking this architecture from massive AI factories to the next frontier of computing in orbit."

"Vera gives us the CPU performance and memory bandwidth to run enormous amounts of orchestration, code and data processing while keeping GPUs doing what they do best," said Mike Nicolls, president of SpaceXAI. "That means higher-performance AI agents and more useful work from every watt of compute."

NVIDIA Vera -- The CPU for Agents

NVIDIA Vera is the first CPU built for AI agents, designed to accelerate the CPU-intensive work that surrounds model inference -- from tool use and code execution to data processing, orchestration and simulation.

Vera features 88 NVIDIA-designed Olympus cores, NVIDIA Spatial Multithreading technology and high-bandwidth LPDDR5X memory, delivering up to 1.2TB/s of bandwidth. Vera enables up to 1.8x faster task completion compared with x86 CPUs across workloads including agentic AI, reinforcement learning and data processing.

SpaceXAI Scales AI Infrastructure With NVIDIA Vera Rubin

At massive scale, AI infrastructure must support demanding training, reasoning and inference workloads while maximizing performance, power efficiency and utilization. NVIDIA Vera Rubin is codesigned across compute, networking and software to optimize the AI factory as a whole.

The platform brings together NVIDIA accelerated computing, NVIDIA NVLink™ interconnect technology, NVIDIA Spectrum-X™ Ethernet networking, NVIDIA BlueField® data processing and NVIDIA software in an integrated architecture designed to deliver high performance and energy efficiency at scale while driving down cost per token.

As SpaceXAI expands toward gigawatts of computing capacity, Vera Rubin provides a common architecture to efficiently scale its next generation of AI factories -- a foundation SpaceXAI plans to take beyond terrestrial data centers and into orbital computing.

NVIDIA Accelerated Computing, From Earth to Orbit

SpaceXAI's work with NVIDIA is moving beyond AI factories on Earth.

SpaceXAI is developing AI computing infrastructure for orbit, where power, thermal management, bandwidth, reliability and physical integration impose dramatically different constraints from conventional data centers.

SpaceXAI's planned first-generation Starmind AI satellite will be based on the optimized NVIDIA Vera Rubin NVL72 rack-scale system, extending the same accelerated computing architecture powering next-generation AI factories on Earth into space.

NVIDIA and SpaceXAI are working to adapt that foundation to the requirements of orbital computing while preserving a common NVIDIA architecture and software ecosystem.

This delivers one computing foundation across a wide range of environments: Vera CPUs accelerating increasingly sophisticated AI agents, Vera Rubin powering the AI infrastructure behind Grok and gigawatt-scale AI factories on Earth, and NVIDIA accelerated computing extending into orbital AI infrastructure.

Learn more about the [NVIDIA Vera Rubin platform](#) and [NVIDIA Vera CPU](#).

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: SpaceXAI deploying Vera for these new application workloads while building its AI infrastructure on Vera Rubin — from AI factories on Earth to computing in orbit; expectations with respect to growth, performance, availability, and benefits of NVIDIA's products, services and technologies, and related trends and drivers; expectations with respect to NVIDIA's third party arrangements, including with its collaborators and partners; expectations with respect to third parties' business plans and third parties' adoption of NVIDIA technology; expectations with respect to technology developments, and related trends and drivers; projected market growth and trends; expectations with respect to AI and related industries; 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 products 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; NVIDIA's ability to realize the potential benefits of business investments or acquisitions; 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 NVIDIA'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, BlueField, 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.

Kacie Thomas
Corporate Communications
NVIDIA Corporation
press@nvidia.com