



NVIDIA Vera BlueField-4 STX Brings Agentic AI Storage Processing With In-Silicon Security

New and Enhanced NVIDIA DOCA Security Capabilities for Vera BlueField-4 STX Help Protect AI Agents, Context Memory and File-Based Data Access Directly in Silicon

News Summary:

- New and enhanced NVIDIA DOCA Vault, DOCA Argus and DOCA Flow security capabilities for NVIDIA BlueField-4 provide zero-trust file access, agent behavior visibility and network-level isolation.
- With DOCA, Vera BlueField-4 STX delivers runtime threat detection up to 1,000x faster than existing agentless runtime solutions, while enabling network and file access enforcement at speeds of up to 800Gb/s.
- Storage, cybersecurity and system partners in the Vera BlueField-4 STX ecosystem are building platforms for enterprise-scale agentic AI.

NVIDIA GTC Taipei--NVIDIA today announced new [NVIDIA DOCA](#)™ security innovations for [NVIDIA Vera BlueField-4 STX](#), defining a new class of secure-by-design storage for agentic AI factories.

As enterprises move from using AI chatbots to harnessing autonomous agents that reason, retrieve and act across business data, storage is becoming a real-time control point for intelligence. Agentic AI systems depend on fast access to proprietary data and context memory. They also create new exposures as they continuously read, write and share information without direct human oversight.

Vera BlueField-4 STX extends NVIDIA's accelerated storage architecture with [security capabilities designed for the agentic AI era](#). Powered by a unified NVIDIA DOCA security stack and enforced in [NVIDIA BlueField®-4](#) silicon, STX-based platforms can inspect and govern interactions between agents, data and context memory inline, helping enterprises enable continuous policy enforcement in the AI data path.

"Agentic AI turns enterprise data into a living, real-time system -- and that system must be protected where data moves, where context is stored and where agents act," said Jensen Huang, founder and CEO of NVIDIA. "With Vera BlueField-4 STX, NVIDIA and its ecosystem are building secure-by-design storage infrastructure that enforces trust in silicon at the speed of AI."

The security innovations enhance the NVIDIA STX-accelerated AI-native storage foundation, which keeps data close and accessible for enterprise systems and long-context reasoning. NVIDIA DOCA enables Vera BlueField-4 STX to provide up to 1,000x faster runtime threat detection than existing agentless runtime solutions, while enforcing network and file access policies at speeds of up to 800Gb/s.

Secure-by-Design Storage for Agentic AI

The expanded Vera BlueField-4 STX platform brings [NVIDIA DOCA security libraries and microservices](#) to the [AI storage](#) layer, helping enterprises protect data, agents and context memory as agentic AI moves into production. They include:

- New NVIDIA DOCA Vault microservices, which help ensure that only authorized AI workloads can access the right files with the right permissions.
- NVIDIA DOCA Argus, which provides visibility into agent behavior and AI workload activity.
- NVIDIA DOCA Flow, which helps isolate network traffic and protect sensitive data across multi-tenant AI environments.

These new and enhanced capabilities run on BlueField-4, where policies can be enforced in silicon while data continues moving at AI factory speeds. Together, they help partners deliver storage platforms that are secure by design, reducing the risk of unauthorized access, data extractions and sensitive context exposure without adding complexity to enterprise AI deployments.

Partners Build the Secure AI Storage Ecosystem

Cybersecurity leaders integrating enterprise security solutions with Vera BlueField-4 STX include [Akamai](#), [Armis from ServiceNow](#), [Check Point](#), Cisco, [CrowdStrike](#), [EQTY](#), [F5](#), Fortinet, [Palo Alto Networks](#), [TrendAI](#), [Xage Security](#) and [Zscaler](#).

Leading storage providers and systems manufacturers are building on STX to deliver AI-native storage infrastructure for inference, training and analytics. Storage providers building STX-based platforms include [Cloudian](#), [DDN](#), Dell Technologies, [Everpure](#), [Hitachi Vantara](#), HPE, IBM, [MiniIO](#), NetApp, Nutanix, [VAST Data](#) and [WEKA](#).

Manufacturing partners AIC, ASUS, Foxconn, Gigabyte, Quanta Cloud Technology (QCT), Supermicro, Wistron and Wiwynn are developing STX-based systems.

Global systems integrators Accenture, Deloitte and Worldwide Technology are working with NVIDIA and its partners to bring secure AI storage solutions to enterprises globally.

Availability

STX-based platforms are expected to be available from partners in the second half of 2026.

Watch Huang's [keynote](#) and learn more at [NVIDIA GTC Taipei](#).

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: 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 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 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, BlueField and NVIDIA DOCA are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and/or 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
Corporate Communications
NVIDIA Corporation
press@nvidia.com