




# NVIDIA BlueField-4 STX Storage Processor

Powering a new class of AI-native data platforms for the agentic AI factory.



The NVIDIA® BlueField®-4 STX storage processor is purpose-built to power AI-native data platforms with accelerated data processing, high-performance NVMe storage, and in-silicon security. Integrating the NVIDIA Vera CPU and NVIDIA® ConnectX-9 SuperNIC™ in a self-hosted architecture, the BlueField-4 STX storage processor functions as both the CPU and data-path engine, enabling high-throughput data movement, efficient storage access, and secure data services for the most demanding AI storage workloads.

As part of the NVIDIA STX reference architecture, BlueField-4 STX storage is designed for a broad range of AI storage workloads, including NVIDIA CMX™ context memory storage, AI data platforms, and high-performance storage infrastructure for AI factories.

## Powering Agentic AI Storage

The NVIDIA BlueField-4 STX storage processor combines advanced compute, networking, programmability, and security to enable AI-native storage infrastructure. Equipped with the NVIDIA Vera CPU, the BlueField-4 STX storage processor delivers power-efficient compute optimized for AI storage services and data-intensive workflows. Purpose-built for agentic AI infrastructure, it orchestrates data access, metadata processing, and storage services directly in the I/O path, reducing complexity and improving efficiency.

Co-designed as part of the NVIDIA Vera Rubin platform, the BlueField-4 STX storage processor works seamlessly with NVIDIA accelerated computing infrastructure as the storage-side CPU and data-path engine. It accelerates secure NVMe storage access and data movement across compute, networking, and storage.

In this architecture, NVIDIA BlueField-4 STX storage processors seamlessly connect AI storage to AI compute over NVIDIA Spectrum-X™ Ethernet networking, creating a high-performance fabric with up to 800 Gb/s per port for ultra-fast, scalable data access.

Through the fully programmable NVIDIA DOCA™ software platform, developers can build optimized AI storage platforms. These support NVIDIA STX high-performance and context memory storage, with accelerated data movement, NVMe access, data protection, encryption, and secure isolation on the NVIDIA BlueField-4

storage processor. In NVIDIA CMX deployments, NVIDIA DOCA Memos™ enables the BlueField-4 STX storage processor to intelligently manage KV cache across AI memory and storage tiers, optimizing data placement and movement for inference workloads. The platform supports file-based data access security enforced in BlueField-4 silicon, helping protect sensitive information and enable secure multi-tenant storage architectures. Trusted by storage industry leaders, BlueField-4 is helping shape a new class of AI storage platforms built to manage the scale and complexity of modern AI infrastructure.

Key Features

Dual BlueField-4 STX Storage Processor

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Network Interface  | <div>&gt; Bandwidth: 3.2 Tb/s</div> <div>&gt; 200/100 Gb/s SerDes PAM4</div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| PCIe Interface     | <div>&gt; PCIe Gen 6 x96 lanes</div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Compute and Memory | <div>&gt; Dual Vera CPU, each with up to 84 Neoverse Arm® V2 cores</div> <div>&gt; 384 GB on-board LPDDR5x</div> <div>&gt; 2 TB on-board M.2 SSD</div> <div>&gt; Redundant M.2 socket</div> <div>&gt; Data path accelerator<ul style="list-style-type: none"><li>Heavy multi-threading (64 cores, 1K HW threads) apps acceleration</li><li>DOCA-programmable</li></ul></div>                                                                                                                                                                                                                                                                          |
| AI Networking      | <div>&gt; DOCA-Accelerated Programmable Steering and Packet Processing</div> <div>&gt; Flexible overlay encapsulation (VxLAN, NVGRE, Geneve)</div> <div>&gt; Built-in support for inline encryption/decryption</div> <div>&gt; Programmable flexible parser: user-defined and advanced protocols</div> <div>&gt; RDMA and RoCEv2</div> <div>&gt; NVIDIA Spectrum-X Ethernet<ul style="list-style-type: none"><li>Spectrum-X multi-plane</li><li>Spectrum-X net plug-in</li></ul></div> <div>&gt; Programmable RDMA transport</div> <div>&gt; Programmable congestion control</div> <div>&gt; NVIDIA GPUDirect RDMA</div> <div>&gt; TCP offloads</div> |
| Storage            | <div>&gt; DOCA Memos for intelligent KV cache storage</div> <div>&gt; NVMe-oF/NVMe-TCP target accelerations</div> <div>&gt; S3 over RDMA</div> <div>&gt; T10-Diff</div> <div>&gt; AES-XTS 256/512-key data-at-rest acceleration</div> <div>&gt; Real-time data access control</div> <div>&gt; NVIDIA GPUDirect® Storage</div>                                                                                                                                                                                                                                                                                                                         |

**Cybersecurity**

- > Platform security:
  - Secure boot with PQC + classical hardware root of trust
  - PQC + classical secure firmware update
  - Flash encryption
  - Device attestation (SPDM 1.1/1.2)
- > Data protection:
  - Real-time file-based access control
  - Multi-tenant isolation
- > Connection tracking for stateful firewall
- > AES-GCM 128/256-key inline data-in-motion accelerations:
  - IPsec
  - TLS
  - PSP

---

**Management and Control**

- > System and SMM-level out-of-band management
  - > Physical interfaces: USB, I2C, and PCIe
  - > NSM and PLDM over MCTP protocols
  - > Monitoring, reset, recovery and secure upgrade
- 

## Ready to Get Started?

To learn more visit the [NVIDIA STX](#) page, or contact an NVIDIA sales representative: [nvidia.com/networking-contact-sales](https://nvidia.com/networking-contact-sales)