




ConnectX-9 SuperNIC

Highest-performance networking designed for gigascale AI factories.



The NVIDIA® ConnectX®-9 SuperNIC™ supercharges gigascale AI computing workloads with support for per-port speeds of 800 gigabits per second (Gb/s) over InfiniBand and Ethernet. It delivers extremely fast, efficient network connectivity that significantly enhances system performance for AI factories and cloud platforms.

Powerful Networking for AI

ConnectX-9 SuperNICs fuel the next wave of innovation in accelerated, gigascale AI compute fabrics. Seamlessly integrated with next-generation NVIDIA Spectrum-X™ Ethernet and NVIDIA Quantum-X800 networking platforms, ConnectX-9 SuperNICs deliver up to 1.6 terabits per second (Tb/s) of throughput to NVIDIA Rubin GPUs. These platforms offer the robustness, feature sets, and scalability required for trillion-parameter GPU computing, disaggregated serving architectures, and agentic AI applications. With enhanced power efficiency, ConnectX-9 SuperNICs support the creation of sustainable AI data centers operating hundreds of thousands of GPUs, ensuring a future-ready infrastructure for AI advancements.

ConnectX-9 SuperNICs provide enhanced programmable input/output (IO) and telemetry-based congestion control, achieving industry-leading network performance and peak AI workload efficiency. Additionally, ConnectX-9 InfiniBand SuperNICs extend the capabilities of NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™ to boost NVIDIA In-Network Computing in high-performance computing environments—further enhancing overall efficiency and performance for scientific computing workloads, large-scale training, and inference.

Product Specifications

Supported Network Protocols

- > Ethernet
- > InfiniBand

Total Bandwidth

- > 800 Gb/s

Ethernet Port Speeds

- > 800/400/200/100/50/25 Gb/s

InfiniBand Port Speeds

- > 800/400/200/100 Gb/s

Host Interface

- > PCIe Gen6: up to 48 lanes

Portfolio

- > PCIe HHHL form factor, 1P x OSFP224 and 2P x QSFP112
- > OCP3 TSFF 1P x OSFP224 and 2P x QSFP112
- > Quad ConnectX-9 IO card for NVIDIA Vera Rubin NVL144 systems

Key Features

Network Interface	➤ Max. bandwidth: 800 Gb/s ➤ 200/100/50 Gb/s of PAM4 and 25/10 Gb/s NRZ ➤ Port configuration: • One port supporting up to 800 Gb/s • Two ports supporting up to 400 Gb/s each • Supports up to eight split ports
Host Interface	➤ PCIe Gen6 (up to 48 lanes) ➤ NVIDIA Multi-Host™ (up to 4 hosts) ➤ PCIe switch downstream port containment (DPC) ➤ Message-signaled interrupts (MSI) and MSI-X
AI Networking	➤ Remote direct-memory access (RDMA) and RDMA over Converged Ethernet v2 (RoCEv2) accelerations ➤ NVIDIA Spectrum-X Ethernet: • Spectrum-X multi-plane • Spectrum-X net plug-in ➤ Programmable RDMA transport ➤ Advanced, programmable congestion control ➤ NVIDIA Collective Communications Library (NCCL) ➤ NVIDIA Inference Transfer Library (NIXL) ➤ NVIDIA GPUDirect® RDMA ➤ GPUDirect Storage ➤ In-Network Computing ➤ Advanced timing and synchronization ➤ Message Passing Interface (MPI) accelerations
Cloud Networking	➤ Stateless Transmission Control Protocol (TCP) offloads ➤ Single-root IO virtualization (SR-IOV) ➤ NVIDIA Accelerated Switching and Packet Processing™ (ASAP2) Ethernet for software-defined networking (SDN) and virtual network functions (VNF): • Open vSwitch (OVS) acceleration • Overlay network accelerations • Connection tracking (L4 firewall) and network address translation (NAT)
Security	➤ AES-GCM 128/256-bit key data-in-motion accelerations: Internet Protocol Security (IPsec), Transport Layer Security (TLS), Protocol for Secure Packet (PSP) ➤ AES-XTS 256/512-bit key data-at-rest acceleration ➤ Platform security: • Secure boot with hardware root of trust • Secure firmware update • Flash encryption • Device attestation (SPDM 1.1 and SPDM 1.2) • Commercial National Security Algorithm Suite (CNSA) 2.0 post-quantum cryptography (PQC) Secure Boot and Update support

Key Features

Management and Control

- Network Control Sideband Interface (NC-SI)
- Management Component Transport Protocol (MCTP) over System Management Bus (SMBus) and PCIe Platform Level Data Model (PLDM) for:
 - Monitor and Control DSP0248
 - Firmware Update DSP0267
 - Redfish Device Enablement DSP0218
 - Field-Replaceable Unit (FRU) DSP0257
- Security Protocols and Data Models (SPDM) DSP0274
- Serial Peripheral Interface (SPI) to flash

Network Boot

- InfiniBand or Ethernet
- Preboot eXecution Environment (PXE) boot
- Unified Extensible Firmware Interface (UEFI)

*This document describes hardware features and capabilities. For feature availability, refer to the [firmware release notes](#) and [NVIDIA DOCA™ release notes](#).

Ready to Get Started?

To learn more, contact an NVIDIA sales representative:

nvidia.com/networking-contact-sales

