Datasheet

ConnectX-8 SuperNIC

Highest-performance 800 G networking designed for massive-scale AI.

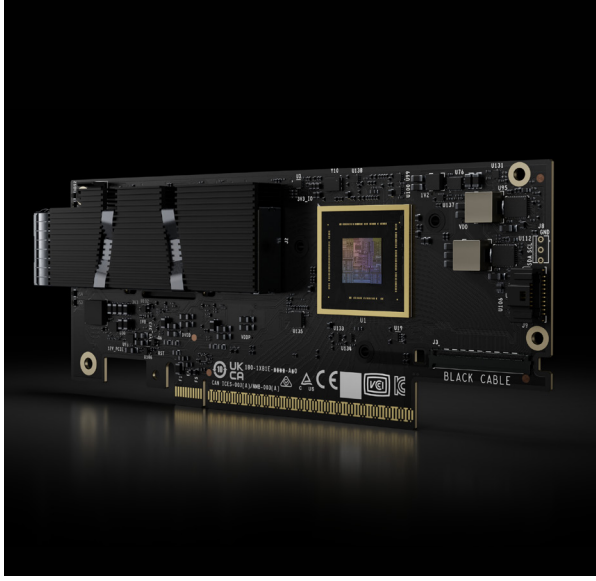

The NVIDIA® ConnectX®-8 SuperNIC™ is optimized to supercharge hyperscale AI computing workloads. With support for both InfiniBand and Ethernet networking at up to 800 gigabits per second (Gb/s), ConnectX-8 SuperNIC delivers extremely fast, efficient network connectivity, significantly enhancing system performance for AI factories and cloud data center environments.

Powerful Networking for the Future of AI

Central to NVIDIA's AI networking portfolio, ConnectX-8 SuperNICs fuel the next wave of innovation in forming accelerated, massive-scale AI fabrics. They seamlessly integrate with next-gen NVIDIA networking platforms, providing up to 800 Gb/s of end-to-end connectivity. These platforms offer the robustness, feature sets, and scalability required for trillion-parameter GPU computing, AI data platforms, and agentic AI applications. With enhanced power efficiency, ConnectX-8 SuperNICs support the creation of increasingly sustainable AI data centers operating hundreds of thousands of GPUs, ensuring a future-ready infrastructure for AI advancements.

ConnectX-8 SuperNICs enable advanced routing and telemetry-based congestion control capabilities, achieving the highest network performance and peak AI workload efficiency. Additionally, ConnectX-8 InfiniBand SuperNICs extend the capabilities of NVIDIA® Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™ to boost in-network computing in high-performance computing environments, further enhancing overall efficiency and performance for both training and inferencing at scale.

Specifications

Supported network protocols	> InfiniBand > Ethernet
Maximum total bandwidth	800 Gb/s
InfiniBand speeds	800/400/200/100 Gb/s
Ethernet speeds	400/200/100/50/25 Gb/s
Host interface	PCIe Gen6: up to 48 lanes
Portfolio	> PCIe HHHL 1P x OSFP > PCIe HHHL 2P x QSFP112 > Dual ConnectX-8 Mezzanine > OCP 3.0 TSFF 1P x OSFP

Key features

Network Interface	InfiniBand	Ethernet
	> Supports 200/100/50 G PAM4	> Supports 100/50 G PAM4 and 25/10 G NRZ
	> Speeds: • 1 port x 800/400/200/100 Gb/s • 2 ports x 400/200/100 Gb/s	> Speeds: • 1 port x 400/200/100 Gb/s • 2 ports x 400/200/100/50/25 Gb/s
	> Max. bandwidth: 800 Gb/s	> Supports up to 8 split ports
	> IBTA v1.7-compliant	> Max. bandwidth: 800 Gb/s
	> 16 million I/O channels	
	> 256- to 4,096-byte MTU, 2 GB messages	
Host Interface	> PCIe Gen6 (up to 48 lanes)	
	> NVIDIA Multi-Host™ (up to four hosts)	
	> PCIe switch downstream port containment (DPC)	
	> MSI/MSI-X	
Optimized Cloud Networking	> Stateless TCP offloads: IP/TCP/UDP checksum	
	> LSO, LRO, GRO, TSS, RSS	
	> SR-IOV	
	> Ethernet Accelerated Switching and Packet Processing™ (ASAP²) for SDN and VNF: • OVS acceleration • Overlay network accelerations: VXLAN, GENEVE, NVGRE • Connection tracking (L4 firewall) and NAT • Hierarchical QoS, header rewrite, flow mirroring, flow-based statistics, flow aging	
Advanced AI/ HPC Networking	> RDMA and RoCEv2 accelerations	
	> Advanced, programmable congestion control	
	> NVIDIA® GPUDirect® RDMA	
	> GPUDirect Storage	
	> In-network computing	
	> High-speed packet reordering	
	> MPI accelerations • Burst-buffer offloads • Collective operations • Enhanced atomic operations • Rendezvous protocol offloads	

Key features	
AI/HPC Software	> NCCL > HPC-X > DOCA UCC/UCX > Open MPI > MVAPICH2
Cybersecurity	> Platform security • Secure boot with hardware root of trust • Secure firmware update • Flash encryption • Device attestation (SPDM 1.1) > Inline crypto accelerations: IPsec, MACsec, PSP
Advanced Timing and Synchronization	> Advanced Precision Time Protocol (PTP): IEEE 1588v2 (any profile), G.8273.2 Class C, line-rate hardware timestamp (UTC format) > SyncE: Meets G.8262.1 (eEEC) > Precise Time Measurement (PTM) > Configurable pulse per second (PPS) in and out > Time-triggered scheduling > PTP-based packet pacing
Management and Control	> Network Control Sideband Interface (NC-SI) > MCTP over SMBus and PCIe 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 > Joint Test Action Group (JTAG) IEEE 1149.1 and IEEE 1149.6
Network Boot	> InfiniBand or Ethernet > PXE boot > iSCSI boot > Unified Extensible Firmware Interface (UEFI)

Ready to Get Started?

To learn more, contact an NVIDIA sales representative:
pny.com/en-eu/nvidia-connectx-8-supernic