




# 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

<b>Supported network protocols</b>	<ul style="list-style-type: none"> <li>&gt; InfiniBand</li> <li>&gt; Ethernet</li> </ul>
<b>Maximum total bandwidth</b>	800 Gb/s
<b>InfiniBand speeds</b>	800/400/200/100 Gb/s
<b>Ethernet speeds</b>	400/200/100/50/25 Gb/s
<b>Host interface</b>	PCIe Gen6: up to 48 lanes
<b>Portfolio</b>	<ul style="list-style-type: none"> <li>&gt; PCIe HHHH 1P x OSFP</li> <li>&gt; PCIe HHHH 2P x QSFP112</li> <li>&gt; Dual ConnectX-8 Mezzanine</li> <li>&gt; OCP 3.0 TSFF 1P x OSFP</li> </ul>

## Key features

Network Interface	InfiniBand	Ethernet
	<ul style="list-style-type: none"><li>&gt; Supports 200/100/50 G PAM4</li><li>&gt; Speeds:<ul style="list-style-type: none"><li>• 1 port x 800/400/200/100 Gb/s</li><li>• 2 ports x 400/200/100 Gb/s</li></ul></li><li>&gt; Max. bandwidth: 800 Gb/s</li><li>&gt; IBTA v1.7-compliant</li><li>&gt; 16 million I/O channels</li><li>&gt; 256- to 4,096-byte MTU, 2 GB messages</li></ul>	<ul style="list-style-type: none"><li>&gt; Supports 100/50 G PAM4 and 25/10 G NRZ</li><li>&gt; Speeds:<ul style="list-style-type: none"><li>• 1 port x 400/200/100 Gb/s</li><li>• 2 ports x 400/200/100/50/25 Gb/s</li></ul></li><li>&gt; Supports up to 8 split ports</li><li>&gt; Max. bandwidth: 800 Gb/s</li></ul>
Host Interface	<ul style="list-style-type: none"><li>&gt; PCIe Gen6 (up to 48 lanes)</li><li>&gt; NVIDIA Multi-Host™ (up to four hosts)</li><li>&gt; PCIe switch downstream port containment (DPC)</li><li>&gt; MSI/MSI-X</li></ul>	
Optimized Cloud Networking	<ul style="list-style-type: none"><li>&gt; Stateless TCP offloads: IP/TCP/UDP checksum</li><li>&gt; LSO, LRO, GRO, TSS, RSS</li><li>&gt; SR-IOV</li><li>&gt; Ethernet Accelerated Switching and Packet Processing™ (ASAP²) for SDN and VNF:<ul style="list-style-type: none"><li>• OVS acceleration</li><li>• Overlay network accelerations: VXLAN, GENEVE, NVGRE</li><li>• Connection tracking (L4 firewall) and NAT</li><li>• Hierarchical QoS, header rewrite, flow mirroring, flow-based statistics, flow aging</li></ul></li></ul>	
Advanced AI/ HPC Networking	<ul style="list-style-type: none"><li>&gt; RDMA and RoCEv2 accelerations</li><li>&gt; Advanced, programmable congestion control</li><li>&gt; NVIDIA® GPUDirect® RDMA</li><li>&gt; GPUDirect Storage</li><li>&gt; In-network computing</li><li>&gt; High-speed packet reordering</li><li>&gt; MPI accelerations<ul style="list-style-type: none"><li>• Burst-buffer offloads</li><li>• Collective operations</li><li>• Enhanced atomic operations</li><li>• Rendezvous protocol offloads</li></ul></li></ul>	

## 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 (eEEEC)
- > 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:  
[nvidia.com/en-us/contact/sales](https://www.nvidia.com/en-us/contact/sales)