



NVIDIA ConnectX-7

Also known as: CX7 · MCX755106AS-HEAT · MCX75310AAS-NEAT

ConnectX-7 is the node-level network adapter of the H100/H200 generation: a 400 Gb/s PCIe Gen5 card that runs InfiniBand NDR or Ethernet from the same silicon (VPI). In a rail-optimized HGX design, one ConnectX-7 sits behind each GPU, PCIe-switch-attached so GPUDirect RDMA traffic never crosses the CPU root complex. The connector and protocol you order lock in the fabric you're building toward, InfiniBand for maximum determinism or Spectrum-X Ethernet for operational uniformity, so it is worth getting that decision right before ordering eight of these per node.

400 Gb/s Max bandwidth · single-port NDR InfiniBand or 400GbE	PCIe 5.0 x16 Host interface · Gen4 x16 compatible	330-370 Mmsg/s Message rate · reported, RDMA
1, 2, or 4 Ports · OSFP, QSFP112, or SFP56 depending on SKU	~25 W Typical power · single 400G port, passive cables	

INTERFACES		01/05
Host interface	PCIe Gen 5.0, x16 · 16/32 GT/s, Gen 3.0 compatible	
InfiniBand speeds	NDR 400, HDR 200, EDR 100 Gb/s	
Ethernet speeds	400/200/100/50/25/10 GbE	
Connectors	OSFP, QSFP112, or SFP56 · SKU-dependent, not interchangeable	
Interface technologies	NRZ (10/25G), PAM4 (50/100G)	

PERFORMANCE & OFFLOADS		02/05
RDMA	RoCEv2, zero-touch RoCE	
GPU offload	GPUDirect RDMA, GPUDirect Storage	
Message rate	330-370 Mmsg/s · reported	
Collective offload	SHARP · in-network reduction	
SDN offload	ASAP2 · OVS/VXLAN/GENEVE acceleration	
Security offload	IPsec, TLS, MACsec · 128/256-bit, in-line	

FORM FACTORS		03/05
PCIe cards	HHHL, FHHL	
OCP cards	OCP 3.0 TSFF, SFF	
Multi-Host	Up to 4 hosts · via PCIe bifurcation, 32-lane variants	
POWER		04/05
Single-port 400G, typical	24.9-25.9 W · OSFP or QSFP112, passive cables	
Dual-port variants, typical	24.9-25.9 W · MCX755106 series	
Operating temperature	0-55°C	
Storage temperature	-40-70°C	
PLATFORM SUPPORT		05/05
Reference deployment	1 NIC per GPU · rail-optimized HGX H100/H200 nodes	
OS support	Linux (inbox/MLNX_OFED), Windows, VMware ESXi	
Fabric compatibility	InfiniBand NDR fabrics or Spectrum-X Ethernet	
PROCUREMENT		
Typical lead time Not published by NVIDIA as a standalone figure. Catalog resellers (Newegg, CDW) show single cards in stock same-week; volume orders bundled into OEM server builds (Supermicro, Dell, Lenovo) are quoted separately and run longer. Pricing Reported \$1,780-\$1,820 per single-port 400GbE QSFP112 card at catalog resellers (Newegg, Compute4Less); some marketplace listings for the same part number run as high as \$3,720. Price varies by connector (OSFP vs. QSFP112) and port count. - Price is per card. A rail-optimized 8-GPU H100/H200 node needs 8 cards, roughly \$14,000-\$30,000 in NICs alone at reported street pricing, before switch ports, optics, or cabling. - OSFP and QSFP112 SKUs are not interchangeable; match the connector to the switch-side optics before ordering. - Most volume SKUs are firmware/part-number locked to either InfiniBand or Ethernet at the factory, even though the underlying VPI silicon supports both.		

FIELD NOTES

- NIC-per-GPU cost math is easy to leave out of a cluster estimate: 8× ConnectX-7 at \$1,800-\$3,700 street adds \$14K-\$30K per 8-GPU node before switch ports, optics, and cabling are counted.
- InfiniBand vs. Spectrum-X Ethernet is decided at the fabric level, not the NIC: ConnectX-7 (VPI) runs either. If the far end is an SN5600 Spectrum-X switch, the card's Ethernet firmware and RoCE congestion-control behavior matter more than raw port speed.
- GPUDirect RDMA needs a PCIe-switch topology, NIC and GPU sharing a PCIe switch downstream of the CPU, to avoid a CPU root-complex hop. That is a server board design decision the NIC alone cannot fix.
- Connector choice locks in a cabling ecosystem early: single-port OSFP variants target InfiniBand NDR fabrics, dual-port QSFP112 variants target Ethernet-centric or mixed deployments. Changing later means re-cabling, not re-flashing.
- At roughly 25 W typical per card, power draw is a rounding error next to the GPU it feeds, but 8 cards per node across a few hundred nodes adds a real line item to PDU and airflow planning.

QUESTIONS WE GET ON THIS PART

What is ConnectX-7 used for?

It's the per-GPU network adapter in H100/H200-class clusters, typically one card per GPU in a rail-optimized design, running either InfiniBand NDR (400 Gb/s) or 400GbE.

How much does a ConnectX-7 cost?

Reported street pricing is roughly \$1,800 to \$3,700 per card depending on reseller and configuration; single-port 400GbE QSFP112 cards run near the low end of that range at catalog resellers.

Should I choose InfiniBand or Ethernet for ConnectX-7?

The card supports both from the same silicon (VPI); the decision is made at the fabric level. InfiniBand stays more deterministic for tight GPU collective synchronization. Spectrum-X Ethernet (ConnectX-7 or BlueField-3 paired with SN5600 switches) trades a little of that determinism for standard Ethernet tooling and multi-vendor optics.

Does ConnectX-7 need BlueField-3 to work?

No. ConnectX-7 is the standard host NIC used across both InfiniBand and Spectrum-X deployments. BlueField-3 is a separate DPU-class SuperNIC that NVIDIA's Spectrum-X reference designs use specifically to unlock full adaptive routing and congestion control; ConnectX-7 works with or without it elsewhere in the fabric.

REFERENCES

- [1] <https://www.nvidia.com/content/dam/en-zz/Solutions/networking/infiniband/connectx-7-datasheet.pdf>
- [2] <https://networking-docs.nvidia.com/connectx7hw/specifications>
- [3] <https://www.newegg.com/nvidia-connectx-7/p/N82E16833942008>
- [4] <https://www.servethehome.com/nvidia-connectx-7-400gbe-and-ndr-infiniband-adapter-review-from-pny-supermicro-intel-sapphire-rapids/>
- [5] <https://www.fibermall.com/blog/how-to-build-cluster-with-128-dgx-h100.htm>