


HOPPER / HGX & NVL

NVIDIA H200. Choose the right fabric.

The HGX configuration references Supermicro's SYS-821GE-TNHR, an 8U dual-Xeon system supporting eight H200 SXM GPUs. H200 NVL is a separate PCIe option; it uses a compatible PCIe chassis such as the SYS-521GE-TNRT and a different interconnect design.



Official NVIDIA HGX H200 eight-GPU baseboard. This component view is separate from the OEM chassis image.

GPU baseboard	8× H200 SXM with 141 GB HBM3e each; 1,128 GB aggregate. Up to 4.8 TB/s local memory bandwidth and 900 GB/s NVLink per GPU.
GPU fabric	HGX eight-GPU designs use NVLink and four NVSwitch devices. GPU fabric and external scale-out network are separate parts of the architecture.
NVL alternative	H200 NVL: 141 GB HBM3e, up to 600 W, PCIe 5.0 x16. Supported 2- or 4-GPU NVLink bridge groups; it is not the HGX eight-way switched fabric.

A complete system

H200 SXM requires a supported HGX platform. H200 NVL is a PCIe accelerator with supported 2- or 4-way NVLink bridge groups. Neither the GPU count nor aggregate memory alone defines the interconnect topology.

SYSTEM PLATFORM / ENGINEERING SCOPE

From baseboard to rack.

HGX reference platform	Supermicro SYS-821GE-TNHR, 8U, supporting HGX H100 / H200. The NexTec proposal specifies H200 SXM.
Host compute	Dual 4th / 5th Gen Intel Xeon Scalable. CPU models, core count, TDP, and firmware are selected against the supported OEM list.
ECC memory	32 DDR5 DIMM slots; up to 8 TB platform capacity. Actual speed depends on CPU, DIMM type and DIMMs per channel.
Boot / storage	2× M.2 NVMe for boot. Up to 16 hot-swap 2.5-inch NVMe bays: 12 default, 4 optional; SATA configuration also available.
Fabric expansion	Eight PCIe 5.0 x16 low-profile slots, plus two FHHL slots and two optional FHHL slots. Select NIC count, GPU affinity, fabric, optics, and cabling as a complete design.
Power / facility	H200 SXM up to 700 W each: 5.6 kW across eight GPUs before the host and other components. Complete feed, PSU, cooling and rack requirements are confirmed for the ordered build.
System acceptance	GPU and fabric health, PCIe/NVLink enumeration, memory and storage checks, sustained power/thermal behavior, representative training or inference, and documented firmware/software versions.

Technical references

[NVIDIA H200 specifications ↗](#)[Supermicro HGX platform specifications ↗](#)[Request an H200 system ↗](#)