



NVIDIA DGX B200

Also known as: DGX B200 · DGXB-G1440

DGX B200 is NVIDIA's fully integrated, factory-built Blackwell system: eight B200 SXM GPUs, NVLink 5 fabric, dual Xeon hosts and eight ConnectX-7 NICs in one 10U chassis with a single support contract from NVIDIA. It is the reference appliance the DGX SuperPOD architecture is built from. Teams that want a validated, drop-in node with NVIDIA-direct support buy DGX; teams with an existing OEM relationship, in-house cluster ops, and a preference to tune the CPU, storage and networking BOM buy the equivalent HGX B200 platform from Supermicro, Dell, HPE, Lenovo or similar and integrate it themselves. The delta is mostly integration risk and support terms, not raw GPU performance.

8x B200 SXM GPUs · Blackwell, NVLink 5	1,440 GB HBM3e Total GPU memory	~14.3 kW System max power · 200-240V input
10U rackmount Form factor	8x ConnectX-7 Networking · up to 400 Gb/s InfiniBand or Ethernet per port	313.9 lb (142.4 kg) max Weight

COMPUTE		01/05
GPUs	8x NVIDIA B200 SXM5 · Blackwell architecture	
GPU-to-GPU fabric	NVLink 5 · 1.8 TB/s per GPU bidirectional	
CPUs	2x Intel Xeon Platinum 8570 · 56 cores each, 2.1/4 GHz base/max boost	
FP4 dense compute	~14.4 TB/s aggregate GPU-to-GPU bandwidth · 8-GPU node	

MEMORY & STORAGE		02/05
Total GPU memory	1,440 GB HBM3e	
System memory	2 TB · 32 DIMMs, upgradable to 4 TB	
OS storage	2x 1.92 TB NVMe M.2 · RAID 1	
Data cache storage	8x 3.84 TB NVMe U.2 SED · RAID 0, ~30.7 TB raw	

NETWORKING		03/05
Compute fabric NICs	8x ConnectX-7 single-port · up to 400 Gb/s InfiniBand	
Ethernet option	400/200/100/50/40/25/10 GbE · per port, same ConnectX-7 cards	
In-band management/storage	2x BlueField-3 DPU dual-port cards	
Host connectivity	2 network module trays · field-serviceable	
POWER & COOLING		04/05
System max power	~14.3 kW · reported by NVIDIA and resellers	
Power supplies	6x 3.3 kW · 5+1 redundant, 200-240V/16A/50-60Hz each	
Cooling	Air-cooled · front-to-back	
Airflow	1,550 CFM	
Heat output	48,794 BTU/hr	
Operating temperature	10 C to 35 C (50 F to 95 F)	
PHYSICAL		05/05
Form factor	10U rackmount	
Height	17.5 in (444 mm)	
Width	19 in (482.3 mm) max	
Depth	35.3 in (897.1 mm) max	
Weight	313.9 lb (142.4 kg) max	
PROCUREMENT		
Typical lead time		
8 to 16 weeks from order to delivery, reported, varies by region and order volume.		
Pricing		
\$450,000 to \$600,000 per unit, reported and street-sourced (one reseller lists \$600,000 flat; other market surveys cite figures nearer \$515,000). NVIDIA does not publish a list price; every quote runs through a partner.		
- Price is quote-only from NVIDIA and its resellers; nobody publishes a public list price. - The DGX premium buys a single-throat-to-choke support contract, NVIDIA-validated firmware/driver stack, and DGX SuperPOD reference-architecture compatibility, not extra GPU performance over an equivalent HGX B200 OEM box. - Base Command Manager ships free for up to 8 accelerators per system but without support; production clusters typically add NVIDIA Mission Control or AI Enterprise for supported cluster management. - Blackwell allocation has been tight through mid-2026; expect priority given to large committed-volume orders over single-unit purchases.		

FIELD NOTES

- At ~14.3 kW max draw, one DGX B200 alone exceeds what a standard 5 to 8 kW colo rack can deliver. Plan on high-density power (208V three-phase PDUs) or spreading systems across racks, not stacking multiple B200 nodes in one legacy cabinet.
- 1,550 CFM of front-to-back air is a lot of air to move through a 10U chassis; hot-aisle containment stops working as a nice-to-have at this density and becomes mandatory to avoid recirculation and thermal throttling.
- 313.9 lb in 10U concentrated at the front of a rack changes the floor-loading and rack-tip-over math versus a typical 1U/2U server deployment; check raised-floor or slab load ratings before racking multiple units.
- If your team already runs an OEM fleet (Supermicro, Dell, HPE) and has in-house firmware/driver ops, an HGX B200 build usually lands at comparable or lower cost with more BOM flexibility (storage, NIC count, CPU choice). DGX earns its premium mainly on turnkey support and SuperPOD-grade validation, which matters most for teams without a bench of GPU cluster engineers.
- 8 ConnectX-7 NICs means the networking bill of materials (cables, switch ports, optics) is nearly as consequential to total system cost as the GPUs; budget switch fabric alongside the node price, not after it.

QUESTIONS WE GET ON THIS PART

How much does a DGX B200 cost?

NVIDIA does not publish a list price. Reported street pricing runs roughly \$450,000 to \$600,000 per 8-GPU system depending on reseller and configuration, with one reseller listing it flat at \$600,000. Get a quote; volume and region move the number substantially.

DGX B200 vs HGX B200: what's the actual difference?

Same B200 GPUs and NVLink 5 fabric either way. DGX is NVIDIA's own pre-integrated, pre-validated appliance sold and supported directly by NVIDIA; HGX B200 is the reference board OEMs (Dell, Supermicro, HPE, Lenovo) build into their own chassis with their own CPU, storage and support terms. Choose DGX for turnkey support and SuperPOD compatibility, HGX-based OEM builds if you want BOM control or already have OEM support contracts.

What power and cooling does a DGX B200 need?

Up to ~14.3 kW per unit, drawn through six 3.3 kW power supplies in 5+1 redundancy on 200-240V circuits, with roughly 1,550 CFM of front-to-back air cooling. Most legacy colo racks (5 to 8 kW) cannot host one of these alone without power and airflow upgrades.

Is the DGX B200 still current, or has NVIDIA moved on to B300?

DGX B300 (Blackwell Ultra) began shipping in January 2026 and targets denser memory and reasoning workloads, but DGX B200 is still sold and shipping as of mid-2026, and it retains an FP64 advantage over B300 for HPC-adjacent workloads.

REFERENCES

- [1] <https://docs.nvidia.com/dgx/dgxb200-user-guide/introduction-to-dgxb200.html>
- [2] <https://itctshop.com/product/nvidia-dgx-b200/>
- [3] <https://www.exxactcorp.com/NVIDIA-DGXB-G1440-P2CMI36-E161162479>
- [4] <https://newsletter.semianalysis.com/p/nvidias-blackwell-reworked-shipment>
- [5] <https://nvidianews.nvidia.com/news/nvidia-blackwell-dgx-generative-ai-supercomputing>
- [6] <https://www.nvidia.com/en-us/data-center/dgx-b300/>
- [7] <https://www.gpuyard.com/tutorials/howto/setup-nvlink-blackwell-gpu-servers/>
- [8] <https://www.server-parts.eu/post/nvidia-b200-gb200-blackwell-gpu-server-comparison>

Latest version of this datasheet: <https://gpumith.com/hardware/systems/nvidia-dgx-b200>. Specifications compiled from the references above; verify against vendor documentation before purchase.