Rillor

EXPLAINER

DGX, HGX and MGX explained for GPU system buyers

DGX is NVIDIA's complete system, HGX is the GPU baseboard server makers build around, and MGX is a modular reference architecture, so each name changes what a price covers.



Rillor

Contents

01 Summary

02 Background

03 Method

04 Findings

05 Implications

06 Limits of this analysis

07 Sources

Summary

- DGX, HGX and MGX name three kinds of NVIDIA product. HGX is a GPU baseboard platform that server makers build around [2, 4]. DGX is NVIDIA's own complete system [7, 10, 17]. MGX is a modular reference architecture for chassis and racks [18, 19]. The DGX B300 involves all three [3, 11].
- From A100 to B300, NVIDIA's Fabric Manager guide describes each DGX system and the matching eight-GPU HGX system with one GPU baseboard topology [3]. The differences between a DGX quote and an HGX quote sit in the host CPUs, memory, networking, storage, power, software and support [7, 12, 28, 29, 30].
- Software inclusion depends on the GPU generation. NVIDIA's licensing guide, last updated 2 September 2026, says NVIDIA AI Enterprise is included in the DGX software bundle for Hopper DGX systems and purchased separately for Blackwell DGX systems [24]. The DGX B200 and B300 web pages list DGX OS as included and AI Enterprise and Mission Control as recommended [7, 10].
- Support differs in scope. DGX datasheets give three years of support for hardware and software [8, 11]. Dell's HGX server carries a three-year hardware warranty without software support [29]. NVIDIA support entitlements may end when a product changes hands [27].
- "NVIDIA-Certified" means a server maker's system passed NVIDIA's functional, security and performance tests after a separate qualification step [21]. The three OEM HGX servers reviewed are on NVIDIA's list updated 30 October 2026 [22].
- NVIDIA states B300 eight-GPU memory as 2.1 TB (HGX page for HGX B300, DGX B300 page and datasheet) and as 2.3 TB (DGX B300 user guide), and DGX B300 power as about 14 kW, 14.5 kW and 15.1 kW [2, 10, 11, 13].

Background

This report is for a buyer who holds a DGX quote next to quotes for servers from Dell, Lenovo, Supermicro or other makers built on an HGX board. Both carry the same GPUs [3]. The prices differ for reasons that the GPU name does not describe.

NVIDIA's annual report describes its data center platform as "typically delivered to customers as rack-scale systems, subsystems, or modules, along with software and services" [1]. It lists among its direct customers "original equipment manufacturers, or OEMs, original device manufacturers, or ODMs, system integrators and distributors which we partner with to help bring our products to market" [1]. Read against that description, a DGX is a system NVIDIA delivers, an HGX board is a subsystem that other makers build into their servers, and MGX is a set of designs those makers build to.

The same filing describes NVIDIA's business model as having "transitioned from offering Hopper HGX systems to Blackwell full-scale datacenter solutions" in fiscal year 2026 [1].

The names also change over time. HGX began in 2017 as a reference architecture shared with original design manufacturers [4]. DGX now covers desktops and rack-scale systems as well as eight-GPU servers [15, 17]. This report sets out what each name covers in NVIDIA's current documents, so that a DGX quote and an HGX quote can be read line by line.

Method

What was examined. For each name: what NVIDIA says it is, what a system built on it contains, and what software, support and certification come with it. Five eight-GPU systems are compared part by part. Two are NVIDIA's: the DGX B200 and the DGX B300 [7, 8, 9, 10, 11, 12, 13]. Three are OEM HGX servers on NVIDIA's certified list: the Supermicro SYS-A21GE-NBRT (HGX B200), the Dell PowerEdge XE9780 (HGX B200 or HGX B300) and the Lenovo ThinkSystem SR680a V4 (HGX B300) [22, 28, 29, 30]. The DGX H100/H200, DGX GB300 and DGX Rubin NVL8 serve as references across generations [14, 15, 16].

Sources. NVIDIA's product pages, datasheets, user guides, a technical brief, documentation for Fabric Manager, licensing, support and certification, two press releases and a technical blog [4, 19, 20]. The manufacturers' specifications for the three OEM servers [28, 29, 30]. NVIDIA's Form 10-K [1]. Every source was opened on 24 October 2026, the report's dated statements were checked against their sources again on 7 November 2026, and the Sources list gives each document's own date.

How. Each specification is recorded as the document states it, with the document's date. Where two NVIDIA documents disagree, both figures are given with their sources. FLOPS figures are given as NVIDIA prints them, sparse and dense. Prices are left out. NVIDIA's statements about its own program are reported as its statements.

Findings

1. Three names, three kinds of product

Name	What NVIDIA says it is	Who builds the system	What the name tells a buyer
HGX	"a single baseboard with eight NVIDIA Rubin, NVIDIA Blackwell, or NVIDIA Blackwell Ultra SXMs" [2]	Server makers, which choose the host CPUs and memory [2, 22]	The GPU board and its NVLink switches
DGX	NVIDIA's own systems, from desktops to eight-GPU servers and rack-scale systems [15, 17]	NVIDIA	A configuration NVIDIA sets, with NVIDIA's software list and support term [8, 11, 12]
MGX	"an open modular reference architecture that enables OEMs, ODMs, and ecosystem partners to build accelerated systems faster" [18]	Server makers, and NVIDIA for the DGX B300 chassis and DGX GB300 racks [11, 15]	A chassis or rack design, with no fixed GPU board

Sources: [2, 8, 11, 12, 15, 17, 18, 22].

The names overlap. The DGX B300 is a DGX system. NVIDIA's Fabric Manager guide groups it with HGX B300 under one GPU baseboard topology [3], and its datasheet calls its chassis "NVIDIA MGX-compliant" [11]. A single machine can carry all three names in different senses.

2. HGX: the GPU baseboard

Origin. In May 2017 Foxconn, Inventec, Quanta and Wistron received "early access to the NVIDIA HGX reference architecture" [4]. NVIDIA described HGX as "the same data center design used in Microsoft's Project Olympus initiative, Facebook's Big Basin systems and NVIDIA DGX-1", with eight GPUs connected by NVLink [4]. From the start, HGX and DGX shared a GPU design.

DGX and HGX pairs by generation. NVIDIA's Fabric Manager User Guide, last updated 7 October 2026, describes each generation's DGX and HGX systems together [3].

GPU	Systems the guide names together	GPU baseboards	GPUs per baseboard	NVSwitch chips per baseboard
V100	DGX-2, HGX-2	2	8	6

GPU	Systems the guide names together	GPU baseboards	GPUs per baseboard	NVSwitch chips per baseboard
A100	DGX A100, HGX A100 8-GPU	1	8	6
H100	DGX H100, HGX H100 8-GPU	1	8	4
B100, B200, B300	DGX B200/B300, HGX B200/B300 8-GPU, HGX B100 8-GPU	1	8	2

Source: [3].

For the Blackwell systems the guide says: "The corresponding GPU baseboard NVLink topology has two NVSwitch ASICs and eight B300/B200/B100 GPUs with nine NVLinks from each GPU connected to an NVSwitch" [3]. No source reviewed gives a part number for the baseboard inside a DGX. What NVIDIA documents is one topology for both.

The current HGX table. NVIDIA's HGX page lists HGX B300 and HGX B200 side by side and notes that both are "shipping now" [2].

Figure	HGX B300	HGX B200
Form factor	8x Blackwell Ultra SXM	8x Blackwell SXM
FP4 Tensor Core, sparse and dense	144 and 108 PFLOPS	144 and 72 PFLOPS
FP64/FP64 Tensor Core	10 TFLOPS	296 TFLOPS
Total memory	2.1 TB	1.4 TB
Total NVLink bandwidth	14.4 TB/s	14.4 TB/s
Networking bandwidth	1.6 TB/s	0.8 TB/s

Source: [2].

For Rubin, NVIDIA offers HGX Rubin NVL8 with an NVIDIA Vera CPU or with "x86-based CPU baseboards", and for the x86 version states that "CPU and memory specs are defined by OEM offerings" [2].

Four-GPU boards. Eight GPUs is the form on the current HGX page [2]. Earlier generations also had four-GPU boards. NVIDIA's H200 page offers "NVIDIA HGX H200 partner and NVIDIA-Certified Systems with 4 or 8 GPUs" [5]. NVIDIA's certified list includes "HGX H200 4-GPU" and "HGX H100 4-GPU" servers from Dell, Lenovo and Supermicro [22]. A quote that says "HGX H100" needs a GPU count.

Networking on the HGX B300 board. NVIDIA's Blackwell Ultra datasheet describes HGX B300's "networking integration with NVIDIA ConnectX-8 SuperNICs" [6]. Dell's XE9780 with B300 GPUs has "8 x OSFP with B300 GPU" as embedded ports [29]. Lenovo's SR680a V4 offers "8x OSFP cages which are connected to the 8x ConnectX-8 switches integrated into the GPU system board" [30].

3. DGX: NVIDIA's complete system

A family of systems. The DGX name covers a desktop machine, DGX Spark, built on the GB10 Superchip with 128 GB of unified memory [17], and a rack-scale system, DGX GB300, with 72 Blackwell Ultra GPUs and 36 Grace CPUs [15]. On 24 October 2026 the former DGX H200 product address redirected to the DGX platform page [17], so Hopper DGX hardware is cited from its user guide [14].

Eight-GPU DGX systems across generations.

Part	DGX H200	DGX B200	DGX B300	DGX Rubin NVL8
GPUs and GPU memory	8 x H200, 1,128 GB	8 x Blackwell, 1,440 GB	8 x Blackwell Ultra, 2.1 TB (web page and datasheet) or 2.3 TB (user guide)	8 x Rubin, 2 TB
Host CPUs	2 x Xeon 8480C, 56 cores each	2 x Xeon Platinum 8570, 56 cores each	2 x Xeon 6776P	2 x Xeon 6776P
Host memory	2 TB, 32 DIMMs	2 TB, configurable to 4 TB	2 TB default, up to 4 TB	Not listed on the page
Cluster network	8 x ConnectX-7 single-port InfiniBand cards	8 x ConnectX-7 single-port, up to 400 Gb/s	8 x ConnectX-8, up to 800 Gb/s	8 x ConnectX-9, up to 800 Gb/s
Storage and management network	2 x ConnectX-7 dual-port Ethernet cards	2 x BlueField-3 DPUs	2 x BlueField-3 DPUs	1 x BlueField-4 DPU

Part	DGX H200	DGX B200	DGX B300	DGX Rubin NVL8
Power	10.2 kW max, six 3.3 kW supplies, 4+2	About 14.3 kW max, six 3.3 kW supplies, 5+1	About 14 kW (web page), 14.5 kW busbar or 15.1 kW PSU (datasheet)	About 24 kW
Height	8U	10U	10U	Not listed on the page

Sources: [7, 9, 10, 11, 13, 14, 16].

FP4 is listed at 144 PFLOPS sparse and 72 dense for DGX B200, and 144 sparse and 108 dense for DGX B300 [7, 10]. DGX Rubin NVL8 lists NVFP4 inference at 400 PFLOPS sparse and NVFP4 training at 270 PFLOPS dense [16].

A fixed configuration. NVIDIA's DGX B300 technical brief says the motherboard "holds two Intel Xeon Platinum 6776P CPUs, and the system is configured with two terabytes of RAM, which can be upgraded by the customer to four terabytes if the customer's workload requires it. The rest of the components in the system are in a fixed configuration" [12]. The user guide adds that the system "Supports up to 4 TB of RAM, configurable on-site with 128 GB DIMMs purchased from NVIDIA" [13]. Apart from choosing the AC or DC power version of DGX B300 [12, 13], a DGX quote leaves one hardware choice to the buyer, the memory size.

What changed in DGX B300. Two changes bear on a quote.

- Network ports moved. "The GPU tray is equipped with eight 800Gbps NVIDIA ConnectX-8 OSFP network ports", which "were previously located in the motherboard tray" [12]. The user guide says "ConnectX-8 networking connects directly to GPUs via PCIe Gen6" [13].
- Two power versions. The system is "Available in both AC/PDU and DC/busbar versions" [13]. The AC version has "12 x 3.2 kW" power supplies and twelve AC inlets [13]. The DC version takes power from external busbars in MGX-compliant racks, which can deliver "54V of DC power to each DGX B300" [12].

Two figures stated two ways. The DGX B300 page and datasheet give the system's GPU memory as 2.1 TB [10, 11], the same figure NVIDIA's HGX page gives for HGX B300 [2]. The Blackwell Ultra datasheet gives 270 GB of HBM3E per GPU for HGX B300 [6]. The DGX B300 user guide, last updated 4 November 2026, gives "8 x 288 GB = 2.3 TB total" [13]. Dell and Lenovo give 288 GB per B300 GPU [29, 30]. No document reviewed explains the difference. DGX B300 power appears as "~14 kW" on the web page [10], "14.5 kW (Busbar)" and "15.1 kW (PSU)" in the datasheet [11], and 14.5 kW in the user guide [13].

4. MGX: a modular reference architecture

What NVIDIA says MGX is. The MGX page describes "an open modular reference architecture" for designs "from single-node servers to rack-scale AI factories" [18]. It supports "rack-scale and PCIe GPU solutions", and on the CPU side "both Arm CPUs, such as NVIDIA Grace and Vera, as well as x86 and other Arm-based CPUs" [18]. Its "third-generation MGX rack-scale design" covers servers, cooling, power and connectors, and it supports PCIe, OCP and EIA rack standards [18].

MGX vs HGX in NVIDIA's words. The 28 May 2023 launch release gives the contrast directly. "MGX differs from NVIDIA HGX in that it offers flexible, multi-generational compatibility with NVIDIA products to ensure that system builders can reuse existing designs and easily adopt next-generation products without expensive redesigns. In contrast, HGX is based on an NVLink-connected, multi-GPU baseboard" [19]. With MGX, "manufacturers start with a basic system architecture optimized for accelerated computing for their server chassis, and then select their GPU, DPU and CPU" [19].

MGX counts are NVIDIA statements. The page says "over 100 combinations of modular designs" [18], the 2023 release "more than 100 server variations" [19] and a 2025 blog "hundreds of GPU, DPU, CPU, storage, and networking combinations" [20]. None is a count a buyer can check.

MGX in racks. The 2025 blog says "MGX components in the NVIDIA GB200 NVL72 and GB300 NVL72 systems serve as the foundational infrastructure for managing power density and thermal loads" [20]. DGX GB300 is "a rack-scale, liquid-cooled system designed to fit into modern data centers using NVIDIA MGX racks" [15].

MGX in DGX B300. The DGX B300 web page says the system is "deployable in NVIDIA MGX racks for the first time" and "Compatible with NVIDIA MGX and traditional enterprise racks" [10]. Its datasheet describes a "redesigned, NVIDIA MGX-compliant, air-cooled chassis" [11]. The technical brief points customers who plan higher compute density, and who can deliver 400 to 480 V AC to the rack, to MGX-compliant racks with power shelves and busbars [12].

MGX and HGX for one GPU. The H200 page shows the two paths side by side. H200 SXM is offered in "NVIDIA HGX H200 partner and NVIDIA-Certified Systems with 4 or 8 GPUs", and the PCIe H200 NVL in "NVIDIA MGX H200 NVL partner and NVIDIA-Certified Systems with up to 8 GPUs" [5]. For a buyer, MGX says how a chassis or rack is built, and HGX says which GPU board is inside.

5. Qualified and NVIDIA-Certified

Two levels. NVIDIA's certified-systems page separates qualification from certification. "A server that is qualified has undergone thermal, mechanical, power, and signal integrity qualification to ensure that the accelerated hardware is fully functional in that server design. [...] This testing process is a prerequisite to submitting a system for the NVIDIA-Certified program" [21]. Certified systems then undergo a suite of "functional, security, and performance tests", including "single-GPU, multi-GPU, and multi-node tests" [21].

Testing also covers network and storage offload, the Trusted Platform Module and remote management [21]. The configuration guide adds TPM 2.0 and UEFI boot checks, and testing "within the OEM manufacturer's temperature and airflow specifications" [23].

How many. NVIDIA states that "there are more than 400 NVIDIA-Certified Systems available" [21]. This report did not count the list.

The list. NVIDIA's list, last updated 30 October 2026, says "Each NVIDIA-Certified System is tested with supported NVIDIA GPUs to validate the performance and reliability of the combined system" [22]. The three OEM servers in this report appear in its data center table: the Supermicro SYS-A21GE-NBRT with HGX B200, the Dell PowerEdge XE9780 with HGX B300 and HGX B200, and the Lenovo ThinkSystem SR680a V4 with HGX B300 [22]. DGX B200 and DGX B300 do not appear in the list's data center tables [22].

HGX reference configurations. The same list defines "HGX Reference Configurations" as "NVIDIA-Certified scale-up server platforms" whose names "follow the pattern CPU-GPU-Network Adaptor-Average GPU Network Bandwidth" [22]. So 2-8-9-400 means two CPUs, eight GPUs, nine network adaptors and 400 Gb/s [22]. "The 2-8-9-800 and 2-8-10-800 configurations support NVIDIA HGX B200/B300 systems with an average network bandwidth of 800 Gb/s" [22]. The Lenovo SR680a V4 appears under 2-8-9-800 with HGX B300 and ConnectX-8 [22].

What the list leaves to the quote. Each entry names the GPU boards and, for reference configurations, the network devices tested [22]. It gives no CPU model or memory size for each server. Certification of a server model therefore says nothing about a particular quote's CPU and memory choices.

Support for certified systems. "Hardware support for NVIDIA-Certified Systems is available from the system vendors", and NVIDIA's software products carry their own support [21]. The list adds that "NVIDIA AI Enterprise is supported on NVIDIA-Certified Systems across bare metal, virtualized, and cloud deployments" [22].

6. Software in a DGX price, by generation

The licensing rule. NVIDIA's licensing guide, last updated 2 September 2026, states the rule by GPU generation. "NVIDIA AI Enterprise is included in the NVIDIA DGX software bundle for NVIDIA DGX systems using the NVIDIA Hopper architecture" and "NVIDIA AI Enterprise licenses need to be purchased separately for NVIDIA DGX systems using the NVIDIA Blackwell architecture" [24]. It "is licensed on a per-GPU basis", with a license "required for every GPU installed on the server" that hosts it [24]. It is sold as a subscription, by consumption in cloud marketplaces, or as a perpetual license with five years of required support [24].

GPUs that include it. The guide lists GPUs that come with a five-year AI Enterprise subscription, among them the H100 PCIe or NVL and the H200 NVL, and says the software "can be installed on the NVIDIA Certified Systems in which the GPUs are installed, but not on a different server or instance" [24]. NVIDIA's H200 page marks AI Enterprise as "Add-on" for H200 SXM and "Included" for H200 NVL [5].

What each document shows.

Document	Date	DGX OS	NVIDIA AI Enterprise	NVIDIA Mission Control
DGX B200 web page [7]	Viewed 24 October 2026	"Included"	"Recommended"	"Recommended"
DGX B200 datasheet [8]	October 2025	Listed, no label	Listed, no label	Listed, no label
DGX B300 web page [10]	Viewed 24 October 2026	"Included"	"Recommended"	"Recommended"
DGX B300 datasheet [11]	February 2026	Listed, no label	Listed, no label	Listed, no label
DGX GB300 web page [15]	Viewed 24 October 2026	Listed, no label	Listed, no label	Listed, no label
DGX Rubin NVL8 web page [16]	Viewed 24 October 2026	Operating systems listed	Not mentioned	Described in the page text, outside the Software row
DGX Resources [25]	Viewed 7 November 2026	"now included with your system purchase"	"Recommended for all DGX systems and included in the DGX Software Bundle for DGX systems"	Listed, no label
Licensing guide [24]	2 September 2026		In the DGX bundle for Hopper DGX, bought separately for Blackwell DGX	

Sources: [7, 8, 10, 11, 15, 16, 24, 25].

Read alone, a datasheet or the DGX Resources line could suggest AI Enterprise comes with every DGX [8, 11, 25]. The licensing guide limits that to Hopper DGX systems [24].

Mission Control and Base Command Manager. With Mission Control, Base Command Manager "is included as a core component with full cluster management capabilities" [25]. NVIDIA's Mission Control page says its version 2.5 adds early access support for Vera Rubin and continues support for DGX systems built on the Blackwell architecture "with enhancements and bug fixes across the stack" [26]. Neither page gives a price or

a licensing unit [25, 26]. Base Command Manager is "included with NVIDIA AI Enterprise for basic cluster management", and without a purchase it "is available for free for up to eight accelerators per system. This free-to-use offer does not include support" [25].

On the HGX side. The OEM documents name operating systems: RHEL for Supermicro [28], Ubuntu Server for Dell with B200 or B300 and Red Hat Enterprise Linux for Dell with B300 only [29], and Red Hat Enterprise Linux and Ubuntu Server for Lenovo [30]. None of the three mentions NVIDIA AI Enterprise [28, 29, 30].

7. Support and warranty

DGX. The DGX B200 datasheet gives "Three-year Enterprise Business-Standard Support for hardware and software" [8]. The DGX B300 datasheet gives "Three-year business-standard hardware and software support" [11], and the DGX GB300 page gives the same wording [15]. The DGX Rubin NVL8 page lists no support term [16].

What NVIDIA's Business Standard level covers. NVIDIA's support documentation lists, for Business Standard, technical support access during local business hours and an initial response within four business hours for severity 1 cases, against 24/7 access and one hour at the Business Critical level [27]. For DGX hardware, "Advanced RMA With Onsite Engineer Arrival" with next-business-day delivery is included at Business Standard, with a four-hour option [27]. Footnotes limit this to "selected countries", with a "Minimum volume purchase required" [27].

Entitlements and resale. "Excluding the limited standard warranty on NVIDIA hardware, customers are not entitled to support unless they have purchased a support contract and have an active support entitlement" [27]. Support entitlements "are tied to the company and not an individual, are non-transferable, are non-assignable and may be terminated when the product is transferred to another party" [27]. A buyer of a used DGX system cannot assume the seller's NVIDIA support transfers. The page does not say what a second owner can buy.

OEM HGX servers.

System	Stated warranty or support	Software support
Dell PowerEdge XE9780	"a standard three-year hardware warranty covering repair or replacement of defective components"	"does not include software support". Dell points to its ProSupport Infrastructure Suite for hardware and software
Lenovo ThinkSystem SR680a V4	"Three-year customer-replaceable unit and onsite limited warranty with 9x5 next business day (NBD)"	Optional upgrades through Lenovo Services, including 4-hour or 2-hour response and "software support for Lenovo hardware and some third-party applications"

System	Stated warranty or support	Software support
Supermicro SYS-A21GE-NBRT	No warranty term on the product page. The description reads "onsite service required"	Not stated

Sources: [28, 29, 30].

8. Same GPUs, different bills of materials

Each cell gives what the manufacturer's document states.

Part	DGX B200	DGX B300	Supermicro SYS-A21GE-NBRT (HGX B200)	Dell XE9780 (HGX B200 or B300)	Lenovo SR680a V4 (HGX B300)
Host CPUs	2 x Xeon Platinum 8570, 112 cores total	2 x Xeon 6776P	2 x 4th or 5th Gen Xeon, up to 64 cores	2 x 6th Gen Xeon, up to 86 cores each	2 x Xeon 6700P-series, up to 86 cores
Host memory	2 TB in 32 DIMMs, upgradable to 4 TB	2 TB default, up to 4 TB	32 DIMM slots, up to 4 TB (one DIMM per channel) or 8 TB (two)	Up to 32 DIMM slots, 4 TB of RDIMMs, or 8 TB if 256 GB DIMMs are supported	Up to 4 TB with 32 x 128 GB RDIMMs
Cluster network	8 x ConnectX-7, up to 400 Gb/s	8 x ConnectX-8, up to 800 Gb/s, on the GPU tray	8 low-profile and 2 FHHL PCIe 5.0 slots	B200: up to 12 slots and 12 NICs and DPUs. B300: 8 embedded OSFP, plus at most 4 NICs and DPUs	8 x OSFP 800 Gb/s from ConnectX-8 switches on the GPU board
DPUs	2 x BlueField-3	2 x BlueField-3	None listed	Within the NIC and DPU limit	BlueField-3 B3220 among PCIe adapter options
Storage	2 x 1.9 TB M.2, 8 x 3.84 TB U.2	2 x 1.9 TB M.2, 8 x 3.84 TB E1.S	10 front NVMe bays, 2 M.2	Up to 16 E3.S or 10 U.2	8 NVMe drives
Power supplies	6 x 3.3 kW, 5+1	12 x 3.2 kW (AC) or 54 V DC busbar	6 x 5,250 W, 3+3	12 x 3,200 W	8 x 3,200 W (N+1) or 3,800 W (N+N)

Part	DGX B200	DGX B300	Supermicro SYS-A21GE-NBRT (HGX B200)	Dell XE9780 (HGX B200 or B300)	Lenovo SR680a V4 (HGX B300)
Height	10U	10U	10U	10U	8U
Cooling	Air	Air	Air	Air	Air

Sources: [7, 9, 10, 11, 12, 13, 28, 29, 30].

Host memory floors. Dell's guide relays a rule from NVIDIA's design guide: "the minimum total memory capacity for Server with 8-way GPU" must be greater than 1 TB per system for B200 and greater than 2 TB per system for B300 [29]. The DGX B300 ships with 2 TB by default [13]. No source explains how the DGX default relates to the relayed floor.

Network cards. With B300 GPUs, Dell's server has eight embedded OSFP ports and "supports a combined maximum of 4 NICs and DPUs per unit" in its PCIe slots. With B200 GPUs the maximum is 12 [29]. Dell's guide states that "For AI and HPC deployments, a 1:1 NIC-to-GPU ratio is recommended", and warns that "Configuring fewer than 8 NICs on the Dell PowerEdge XE9780 may not meet the performance demands of these workloads" [29]. The DGX B200 includes eight cluster cards [7]. A B200 server quoted without cluster cards covers less than a DGX B200 does.

Cooling. All five are air-cooled in their documents [7, 11, 28, 29, 30].

CPU cores. The DGX B300 page gives no core count for its Xeon 6776P [10]. Lenovo's processor table lists that model at 64 cores [30].

9. Checklist: a DGX quote beside an HGX quote

1. **GPU board.** Generation, GPU count and GPU memory, with the document each figure comes from [2, 5, 13, 22].
2. **Host CPUs.** Model, count and cores [7, 12, 28, 29, 30].
3. **Host memory.** Capacity, DIMM count and room to expand, against the minimums Dell relays [13, 29].
4. **Cluster networking.** Card model, count, speed and where the cards attach [12, 29, 30].
5. **DPUs, management ports and storage** [7, 10, 28, 29, 30].
6. **Software.** Operating system, AI Enterprise, Mission Control and Base Command Manager, with license counts and terms [24, 25, 26].
7. **Support.** Term, scope, response times, country limits and transfer on resale [8, 11, 27, 29, 30].
8. **Power and rack.** AC supplies or DC busbar, supply count and rating, and rack type [11, 12, 13, 28, 29, 30].

9. **Certification.** The list entry and any reference configuration [21, 22].

Implications

For procurement and infrastructure buyers. A DGX quote and an HGX quote become comparable once both are written out at the level of the checklist. The DGX fixes every part except memory size and, for DGX B300, the AC or DC power version [12, 13]. A B200 server quote that leaves out the eight cluster cards, or carries less than 2 TB of host memory, covers less hardware than a DGX B200 quote does [7, 29]. A Blackwell DGX quote carries three years of hardware and software support [8, 11], and the AI Enterprise license is a separate purchase for it [24].

For finance teams, lenders and valuers. The asset description should name the configuration, the software licenses and their terms, and the support entitlement. NVIDIA's support entitlements may end on transfer [27], and AI Enterprise subscriptions that come with certain GPUs may be installed only on the system that holds those GPUs [24]. Both points bear on what passes to a later owner.

For analysts and data teams. Configurations need normalizing before prices can be compared. Each specification should carry its document and date, because NVIDIA's own documents differ on B300 memory and power [2, 10, 11, 13]. Rillor's GPU systems pricing method models configurations by GPU family, GPU count, memory, CPU, interconnect, cooling, condition, location and seller type, and each observation in its schema keeps its evidence class.

Limits of this analysis

- The report covers five B200 and B300 systems. Many other certified HGX servers exist [22], and their parts can differ.
- No source reviewed gives part numbers for the GPU baseboard in a DGX system. The report relies on NVIDIA's statement of one baseboard topology [3].
- NVIDIA's HGX design guide was not among the sources reviewed. Host-memory minimums come only through Dell's guide [29].
- Where NVIDIA's documents disagree, on B300 GPU memory and DGX B300 power, the report gives each figure with its source and does not choose between them [2, 10, 11, 13].
- Software and support terms come from NVIDIA's and the manufacturers' published documents as of 24 October 2026, and from NVIDIA's DGX Resources and Mission Control pages as of 7 November 2026. Licensing guides and product pages change, and a signed contract governs any actual purchase. This report is not legal, tax or procurement advice.
- The report states no prices. It describes what a price covers.
- NVIDIA, DGX, HGX, MGX, ConnectX, BlueField and related names are trademarks of NVIDIA Corporation. Dell, Lenovo, Supermicro and their product names are trademarks of their owners. They are used only to identify products. No affiliation or endorsement is implied.

Sources

1. NVIDIA Corporation Annual Report on Form 10-K for the fiscal year ended 25 January 2026. U.S. Securities and Exchange Commission, filed 25 February 2026. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm

2. NVIDIA HGX Platform. NVIDIA, viewed 24 October 2026. https://www.nvidia.com/en-us/data-center/hgx/

3. NVIDIA Fabric Manager (user guide). NVIDIA Docs, last updated 7 October 2026. https://docs.nvidia.com/datacenter/tesla/fabric-manager-user-guide/index.html

4. NVIDIA Partners with World's Top Server Manufacturers to Advance AI Cloud Computing. NVIDIA Newsroom, 29 May 2017. https://nvidianews.nvidia.com/news/nvidia-partners-with-world-s-top-server-manufacturers-to-advance-ai-cloud-computing

5. NVIDIA H200 GPU. NVIDIA, viewed 24 October 2026. https://www.nvidia.com/en-us/data-center/h200/

6. NVIDIA Blackwell Ultra Datasheet. NVIDIA, October 2025. https://dam-cdn.nvd.orangelogic.com/AssetLink/1k0p832eq8r5ca0u5383ie5o4tp3bst1.pdf

7. NVIDIA DGX B200. NVIDIA, viewed 24 October 2026. https://www.nvidia.com/en-us/data-center/dgx-b200/

8. NVIDIA DGX B200 Datasheet. NVIDIA, October 2025. https://dam-cdn.nvd.orangelogic.com/AssetLink/h226p7256j0xwmhmx4qqc1t3dv4jqg2m.pdf

9. NVIDIA DGX B200 User Guide: Introduction to NVIDIA DGX B200 Systems. NVIDIA, last updated 17 October 2026. https://docs.nvidia.com/dgx/dgxb200-user-guide/introduction-to-dgxb200.html

10. NVIDIA DGX B300. NVIDIA, viewed 24 October 2026. https://www.nvidia.com/en-us/data-center/dgx-b300/

11. NVIDIA DGX B300 Datasheet. NVIDIA, February 2026. https://dam-cdn.nvd.orangelogic.com/AssetLink/625w0j7hnw6f07ui7fo63211mti2v08o.pdf

12. NVIDIA DGX B300 Technical Brief: Ready for Deployment, a Deep Dive into the New Architecture of NVIDIA DGX B300 Systems. NVIDIA, October 2025. https://dam-cdn.nvd.orangelogic.com/AssetLink/reof7qjn8r1ht44ud4qc1gt7twrs6x1w.pdf

13. NVIDIA DGX B300 User Guide: Introduction to NVIDIA DGX B300 System. NVIDIA, last updated 4 November 2026, viewed 7 November 2026. https://docs.nvidia.com/dgx/dgxb300-user-guide/introduction-to-dgxb300.html

14. NVIDIA DGX H100/H200 User Guide: Introduction to NVIDIA DGX H100/H200 Systems. NVIDIA, last updated 17 October 2026. https://docs.nvidia.com/dgx/dgxh100-user-guide/introduction-to-dgxh100.html

15. NVIDIA DGX GB300. NVIDIA, viewed 24 October 2026. https://www.nvidia.com/en-us/data-center/dgx-gb300/

16. NVIDIA DGX Rubin NVL8. NVIDIA, viewed 24 October 2026. https://www.nvidia.com/en-us/data-center/dgx-rubin-nvl8/

17. NVIDIA DGX Platform. NVIDIA, viewed 24 October 2026. https://www.nvidia.com/en-us/data-center/dgx-platform/

18. NVIDIA MGX. NVIDIA, viewed 24 October 2026. https://www.nvidia.com/en-us/data-center/products/mgx/

19. NVIDIA MGX Gives System Makers Modular Architecture to Meet Diverse Accelerated Computing Needs of World's Data Centers. NVIDIA Newsroom, 28 May 2023. https://nvidianews.nvidia.com/news/nvidia-mgx-server-specification

20. Building the Modular Foundation for AI Factories with NVIDIA MGX. NVIDIA Technical Blog, 16 May 2025. https://developer.nvidia.com/blog/building-the-modular-foundation-for-ai-factories-with-nvidia-mgx/

21. NVIDIA-Certified Systems. NVIDIA, viewed 24 October 2026. https://www.nvidia.com/en-us/data-center/products/certified-systems/

22. List of NVIDIA-Certified Systems. NVIDIA Certification Programs Documentation, last updated 30 October 2026, viewed 7 November 2026. https://docs.nvidia.com/certification-programs/latest/nvidia-certified-systems.html

23. NVIDIA-Certified Systems Configuration Guide. NVIDIA Certification Programs Documentation, last updated 30 September 2026. https://docs.nvidia.com/certification-programs/latest/nvidia-certified-configuration-guide.html

24. NVIDIA AI Enterprise Licensing. NVIDIA Enterprise Licensing Guide, last updated 2 September 2026. https://docs.nvidia.com/ai-enterprise/planning-resource/licensing-guide/latest/licensing.html

25. DGX Resources. NVIDIA Docs, viewed 7 November 2026. https://docs.nvidia.com/dgx-resources/index.html

26. NVIDIA Mission Control. NVIDIA, viewed 7 November 2026. https://www.nvidia.com/en-us/data-center/mission-control/

27. Support Details. NVIDIA Enterprise Support and Services User Guide, viewed 24 October 2026. https://docs.nvidia.com/enterprise-support-and-services-user-guide/support-details

28. SYS-A21GE-NBRT GPU SuperServer. Supermicro, viewed 24 October 2026. https://www.supermicro.com/en/products/system/gpu/10u/sys-a21ge-nbrt

29. Dell PowerEdge XE9780 Technical Guide, Rev. A03. Dell Technologies, August 2026. https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/poweredge-xe9780-technical-guide.pdf

30. Lenovo ThinkSystem SR680a V4 Server Product Guide (LP2264). Lenovo Press, updated 5 October 2026. https://lenovopress.lenovo.com/lp2264.pdf

About this report

Published by Rillor on 7 November 2026. Superintelligence (SI) tools assist research and drafting, and every figure cites its source. Rillor builds and secures SI systems, supplies the data they learn from and delivers the infrastructure they run on. We apply SI to financial and compute markets.

Cite as: Rillor. DGX, HGX and MGX explained for GPU system buyers. 7 November 2026. https://rillor.com/insights/dgx-hgx-mgx-explained

Notices

Rillor is not a registered investment adviser or commodity trading advisor and does not provide investment or trading advice. Research, forecasts, data and software described on this site are for research and engineering use. Nothing here is an offer or recommendation to buy or sell any security, commodity interest or digital asset. Past or simulated results do not indicate future results.

NVIDIA, AMD, Ceph, Kubernetes and related product names are trademarks of their owners and are used only to identify products. No affiliation or endorsement is implied.

Work with Rillor

Send the workload, system, quantity, delivery window and destination. Contact us at rillor.com/contact?type=gpu-systems or contact@rillor.com. More on GPU systems: rillor.com/compute/gpu-systems