




INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

NVIDIA HGX Platform: Data Center Physical Requirements Guide

Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2025-10-23 · nvidia hgx · data center design · ai infrastructure · gpu server · high-density computing · liquid cooling · power distribution · dgx h100 · blackwell · gb200 nvl72 · dgx b200

Original article: <https://intuitionlabs.ai/articles/nvidia-hgx-data-center-requirements>



In this article

1. [NVIDIA HGX Platform: Data Center Physical Requirements Guide](#)
 2. [Executive Summary](#)
 3. [Introduction and Background](#)
 4. [NVIDIA HGX Platform Overview](#)
 5. [Space and Rack Planning](#)
 6. [Power Infrastructure](#)
 7. [Cooling and Environmental Controls](#)
 8. [Network Infrastructure](#)
 9. [Safety, Security, and Compliance](#)
 10. [Case Studies and Examples](#)
 11. [Implications and Future Directions](#)
 12. [Conclusion](#)
-

NVIDIA HGX Platform: Data Center Physical Requirements Guide

Executive Summary

The deployment of NVIDIA's HGX™ platform – a high-density GPU server architecture – in a data center imposes exceptionally stringent physical requirements. HGX-based systems (such as NVIDIA DGX™ servers) pack 8–16 of the latest [high-performance GPUs](#) into a single chassis (typically 8U–10U for the DGX H100 and DGX B200) **each drawing on the order of hundreds of watts per GPU**. For example, NVIDIA's DGX H100 system (8×H100 GPUs) delivers ~32 PFLOPS of AI performance ([nvidianews.nvidia.com](#)), but requires *roughly* 10–11 kW of power under load ([semianalysis.com](#)). The newer DGX B200 (8×B200 Blackwell GPUs) pushes this to ~14.3 kW per system ([resources.nvidia.com](#)), while NVIDIA's GB200 NVL72 rack-scale system – housing 72 Blackwell GPUs – demands up to **120–140 kW per rack** with mandatory liquid cooling ([sunbirdcim.com](#)). At four H100-class servers per rack, a single rack can demand >40 kW of IT power – vastly exceeding the 10–12 kW/rack limit of most traditional colocation facilities ([semianalysis.com](#)). With Blackwell-generation hardware, those figures jump to 120–140 kW per rack and beyond. This necessitates comprehensive design changes: heavy-duty racks with tall height and reinforced anchoring, multi-megawatt power distribution with high-voltage (e.g. 415 VAC or DC) three-phase feeds, specialized cooling (airflow optimization and direct liquid or immersion cooling) to extract tens to hundreds of kilowatts of heat per rack, and rigorous site infrastructure (power redundancy, environmental controls, floor loading, etc.).

In sum, deploying an HGX platform requires planning **every facet** of physical infrastructure: floor space and layout for high-density racks, mechanical handling (server lifts), seismic bracing and structural support, electrical capacity (PDU and facility UPS design), high-capacity cooling plant or innovative cooling solutions, advanced networking cabling and switching, plus strict security and safety provisions. This report provides an exhaustive, end-to-end analysis of all *physical* requirements – from site selection to rack installation and

power/cooling systems – needed to support HGX GPU servers. We draw on NVIDIA's official DGX SuperPOD design guides (docs.nvidia.com) (docs.nvidia.com) (docs.nvidia.com), industry technical reports, product specifications, and press analyses. Multiple case studies and vendor examples (e.g. liquid-cooled racks by Supermicro (dnnmagazine.com), immersion-cooling facilities (www.cnbc.com)) illustrate real-world solutions. We also examine trends and projections (e.g. DOE reports predicting tripled data center power demand by 2028 (www.reuters.com), or Rack power targets up to 1 MW by 2030 (www.techradar.com)) to contextualize future directions.

Key findings include: modern AI racks can reach 20–30 times the power density of traditional servers (www.techradar.com), forcing data centers to adopt high-voltage DC power distribution and advanced cooling (rear-door heat exchangers, direct-to-chip liquid, immersion, microfluidics) (www.techradar.com) (www.techradar.com). Detailed specifications – such as rack dimensions, sensor placement, PDU circuit ratings, weight/load limits, environmental setpoints (ASHRAE guidelines), and safety standards (Tier-3 redundancy, fire suppression) – are compiled to form a concrete checklist for data center planners. We also discuss broader implications: the massive power draw and heat output of HGX systems drive innovation in facility design (e.g. NVIDIA's work on digital twins (blogs.nvidia.com) (blogs.nvidia.com)) and regulatory concern (e.g. DOE studies (www.reuters.com)).

Overall, this comprehensive report serves as a definitive guide for architects, engineers, and decision-makers preparing a data center for NVIDIA's HGX platform. It details **every** aspect from foundational site criteria to rack-level implementation, supported by extensive data and citations.

Introduction and Background

The past decade has seen a seismic shift in data center workloads. **Machine learning and AI** have driven unprecedented demand for specialized compute: **dense GPU clusters** rather than traditional CPU farms. NVIDIA's GPUs have especially dominated AI training workloads, and the industry has standardized on multi-GPU server designs. In this context, NVIDIA's **HGX platform** (Hyperscale GPU Accelerated Compute platform) has emerged as the reference architecture for high-performance **AI/HPC servers** (nvidianews.nvidia.com) (nvidianews.nvidia.com). First introduced in 2017–2018 for Volta GPUs, and now in its fifth generation with Blackwell GPUs (following Hopper), the HGX platform bundles 8 (or 16) GPUs connected by NVIDIA NVSwitch/NVLink fabrics, to operate as a unified “super-GPU” (nvidianews.nvidia.com) (nvidianews.nvidia.com). Leading OEMs and hyperscale providers worldwide (e.g. Foxconn, Lenovo, Quanta, Inspur, Microsoft, Meta) have adopted HGX-based systems (nvidianews.nvidia.com) (nvidianews.nvidia.com).

As Jensen Huang noted, NVIDIA's HGX-2 (2018) “ups the ante” with sixteen Tesla V100 GPUs delivering over 2 petaflops (nvidianews.nvidia.com). In 2022, NVIDIA's DGX H100 (an 8×H100 HGX system) achieved 32 PFLOPS (FP8) in a single chassis (nvidianews.nvidia.com). In 2024–2025, the Blackwell generation raised the bar further: the B200 GPU delivers up to 1,000 W TDP (air-cooled) or 1,200 W (liquid-cooled), and the B300 “Blackwell Ultra” pushes to 1,100–1,400 W per GPU with 15 PFLOPS FP4 performance (tomshardware.com). NVIDIA's GB200 NVL72 – a full-rack system integrating 36 Grace Blackwell Superchips (72 GPUs) – achieves 1.4 exaFLOPS of FP4 compute in a single rack (nvidia.com). Such systems are envisioned for next-gen **AI “supercomputers”** – e.g. NVIDIA's own Eos system (576 DGX H100 servers, 4,608 GPUs) projected to deliver ~18.4 exaFLOPS (nvidianews.nvidia.com), while DGX SuperPOD configurations based on GB200 scale to

2,000+ nodes. These numbers hint at the magnitude of infrastructure needed. A single DGX H100 system is a multi-kilowatt server weighing 287.6 lb (130.45 kg), before the weight of its rack and supporting equipment ([NVIDIA DGX SuperPOD H100 white-space infrastructure](#)).

Historically, most enterprise data centers were not built for this scale of power density. Conventional racks (at most 10–15 kW) are dwarfed by AI racks (40–140 kW today, with next-generation systems targeting 250–900 kW). Indeed, industry analysis warns that **AI workloads will dominate future data center design** ([www.techradar.com](#)) ([www.reuters.com](#)). The 1 megawatt rack is no longer a distant projection – NVIDIA and partners like ABB, Eaton, and Schneider Electric are actively developing 800 VDC power architectures for 1 MW racks, targeted for deployment with the Rubin Ultra platform in 2027 ([blog.se.com](#)). Critical power capacity to support data centers globally is expected to nearly double to reach 96 GW by 2026 ([mckinsey.com](#)). In response, data centers are innovating with liquid cooling, high-voltage DC power, and digital design tools.

This report focuses on the *physical deployment* aspects of NVIDIA's HGX platform. We systematically cover:

- **Site and floor planning** (space, expansion, weight/load limits, seismic).
- **Rack infrastructure** (size, clearance, blanking, airflow).
- **Electrical power** (capacity, PDU/PDUs, circuit schemes, redundancy, UPS).
- **Cooling and HVAC** (airflow layout, temperature/humidity standards, advanced cooling).
- **Network connectivity** (high-speed fabrics, cable routing).
- **Safety and compliance** (tier standards, fire suppression, security).

Wherever possible, specific numerical design criteria are given (e.g. rack dimensions, kW per rack, floor load in kg/m²) with authoritative citations (NVIDIA design guides ([docs.nvidia.com](#)) ([docs.nvidia.com](#)), standards documents, technical blogs, etc.). We also include **case studies and examples**: e.g. immersion-cooled data centers ([www.cnbc.com](#)), liquid-cooled HGX server racks ([dcnnmagazine.com](#)), and super-rack solutions with 80 kW of D2C cooling ([dcnnmagazine.com](#)). The discussion culminates with future implications (elevated PUE targets, digital twins ([blogs.nvidia.com](#)) ([blogs.nvidia.com](#)), high-voltage DC ([www.tomshardware.com](#))).

NVIDIA HGX Platform Overview

The NVIDIA HGX family is **NVIDIA's reference architecture for GPU-accelerated servers**. It defines the “building block” of high-performance AI/HPC systems. Key features include:

- **Multi-GPU PCB**: HGX boards integrate 4 or 8 GPUs on one board with high-bandwidth NVLink/NVSwitch interconnects ([nvidianews.nvidia.com](#)) ([nvidianews.nvidia.com](#)). For example, HGX-2 (2018) linked 16 V100 GPUs with NVSwitch ([nvidianews.nvidia.com](#)). The latest HGX-H100 links 8 Hopper H100 GPUs with 900 GB/s NVLink per GPU ([nvidianews.nvidia.com](#)).
- **Purpose-Built for Scale**: Detailing HVDC, network, and storage readiness, HGX designs (in partnership with OEMs/ODMs) ensure the servers can scale massively in PODs/SuperPODs ([nvidianews.nvidia.com](#)) ([nvidianews.nvidia.com](#)). NVIDIA's DGX H100 (official HGX-based product) introduced an NVLink Switch to create up to 70 TB/s all-to-all bandwidth across GPUs ([nvidianews.nvidia.com](#)).

- **Variants:** Over time, several HGX variants have appeared. HGX-H100 is aimed at AI training (with 8 H100 GPUs) (nvidianews.nvidia.com). Historical versions (HGX-1,2) used earlier GPUs (Tesla V100, etc) (nvidianews.nvidia.com). In 2024–2025, NVIDIA launched the **HGX B200** (8×B200 Blackwell GPUs, up to 1,000 W each air-cooled) and **HGX B300** (8×B300 Blackwell Ultra GPUs, up to 1,100 W each), alongside the rack-scale **GB200 NVL72** which integrates 36 Grace Blackwell Superchips in a single liquid-cooled rack (supermicro.com). NVIDIA has also announced the **Rubin** architecture (expected late 2026–2027), featuring HBM4 memory and up to 1,800 W TDP per GPU (tweaktown.com).
- **Integration with DPU and NICs:** Modern HGX servers often include NVIDIA DPUs (BlueField) and InfiniBand NICs for offloading networking and storage. For instance, each DGX H100 has two BlueField-3 DPUs and eight ConnectX-7 400 Gb/s InfiniBand adapters (nvidianews.nvidia.com). The B300 generation introduces the CX-8 NIC with 800 Gb/s InfiniBand throughput.

These architectural choices maximize GPUs' interconnect bandwidth and minimize latency, but also raise physical demands. Each GPU is extremely power- and heat-dense. An NVIDIA H100 GPU operates at up to ~700 W (www.tomshardware.com), while the B200 Blackwell GPU pushes to 1,000 W (air-cooled) or 1,200 W (liquid-cooled), and the B300 Blackwell Ultra reaches 1,100–1,400 W (lenovopress.lenovo.com). For H100-based systems, the server's 8 GPUs alone draw ~5.6 kW, not counting CPUs, memory, fans, and power supply losses. In sum, an HGX H100 server can consume ~**10–11 kW** at full load (semianalysis.com), while a DGX B200 draws up to ~**14.3 kW** (resources.nvidia.com). The GB200 NVL72 rack-scale system demands **120–140 kW** for the entire 72-GPU rack (trendforce.com). DGX H100 is 8U and DGX B200 is 10U; each weighs several hundred pounds (docs.nvidia.com).

Table 1 below compares example NVIDIA DGX systems by generation (all HGX-based) for illustration. It shows how design has progressed in compute density and power draw:

System	GPUs (Type)	GPU Mem	AI Perf (Tensor PFLOPS)	Max Power (kW)	Chassis U
DGX A100 (2020) (developer.nvidia.com)	8 × NVIDIA A100 (40/80 GB)	640 GB	~5 (FP16) (developer.nvidia.com)	~6.5 (www.netapp.com)	6U
DGX H100 (2022) (nvidianews.nvidia.com)	8 × NVIDIA H100 (80 GB)	640 GB	32 (FP8) (nvidianews.nvidia.com)	~10.2 (semianalysis.com)	8U
DGX B200 (2024) (resources.nvidia.com)	8 × NVIDIA B200 (180 GB)	1440 GB	~72 (FP8 training); ~144 (FP4 inference)	~14.3	10U
GB200 NVL72 (2024) (nvidia.com)	72 × Blackwell GPUs (186 GB HBM3E per GPU; 372 GB per Grace Blackwell Superchip)	13.4 TB HBM3E	1,440 (FP4, rack)	120–140 (rack)	Full rack
GB300 NVL72 (2025) (nvidianews.nvidia.com)	72 × B300 Ultra (288 GB)	20.7 TB	~1,100 (FP4 dense, rack)	~150+ (rack)	Full rack

Table 1: Example HGX-based server specifications across generations. Systems are broadly balanced for maximum AI throughput. **Note:** Table values combine sources (NVIDIA docs, OEM specs). For DGX A100, NVIDIA quotes five petaFLOPS (FP16) (developer.nvidia.com) with 6.5 kW draw (www.netapp.com). DGX H100 (8×H100) delivers 32 PFLOPS (FP8) (nvidianews.nvidia.com) at ≈10.2 kW (semianalysis.com). The GB200 NVL72 and GB300 NVL72 are full-rack systems requiring mandatory liquid cooling. (Exact power varies by configuration and workload.)

These figures underline the challenge: even a single HGX server requires a **hundreds-of-kilograms chassis** and **many kilowatts** of power. A row of 4–8 H100-class servers becomes a micro-cluster consuming 40–80 kW, while a single GB200 NVL72 rack reaches 120–140 kW. With the GB300 NVL72 and upcoming Rubin platforms, forward-thinking operators are provisioning 250–600 kW per rack. Standard data center racks (typically 42–48U, ~800 mm wide) need special preparation (heavy supports, high-current cabling, powerful cooling fans or liquid loops). The following sections detail these requirements from A–Z.

Space and Rack Planning

Floor Space and Layout

Site Selection & Expansion: When planning for HGX deployment, allocate **significant contiguous floor area**. Not only will each high-density rack use far more space (for cable trays, cold/hot aisles, drop-thru from ceilings), but future growth must be considered. NVIDIA's DGX SuperPOD guide emphasizes reserving adjacent space for scalability (docs.nvidia.com). Rack clusters (often 3–5 racks per Pod) must be **within short cable reach** (often ≤30 m) of each other to maintain GPU-internal networks (semianalysis.com). Thus racks for a given POD should be grouped in the same aisle or neighboring aisles.

The **shape of the white space** may be critical. High-power racks generate heat plumes and require uninterrupted airflow paths. Avoid placing GPUs near external walls or under stairwells. Provide **wide cold aisles** (recommended ≥4–5 ft / ~1.2–1.5 m) for intakes, and hot aisles sufficiently deep for return. Overhead cable trays/cable ladders should be planned to minimize intersections and keep cables within the prescribed length limits (e.g. IB/NVLink max ~50 m) (semianalysis.com). Power lines from floor PDUs to rack PDU must also have clear paths. In summary, a well-planned space will cluster the HGX racks tightly (for network latency and power distribution) while providing generous aisle widths and overhead cable trays.

Rack Dimensions and Standards

Rack Size/Footprint: Use **19-inch EIA-310 compliant** server cabinets (docs.nvidia.com). Each HGX server (e.g. DGX) is ~17.75" (0.45 m) wide (with rack ears) and typically ~6U–8U tall (docs.nvidia.com). Cabinets must be at least 24" wide × 48" deep (600 × 1,200 mm) per NVIDIA guidance (docs.nvidia.com). In practice, NVIDIA recommends wider 32"×48" (800 × 1,200 mm) cabinets to accommodate rear power cables and airflow plenums (docs.nvidia.com). Cabinets must be at least 48U tall to hold multiple GPU systems (docs.nvidia.com).

Accessory Requirements: Equip each rack with proper airflow management. As per NVIDIA, install **blanking panels** in any empty U position (docs.nvidia.com) to prevent recirculation of hot exhaust. Side panels should remain in place to maintain the cold/hot aisle barrier (docs.nvidia.com). Remove front/rear doors if policy

allows, since open racks improve airflow (NVIDIA discourages doors) (docs.nvidia.com). Ensure the rack has sufficient cable entry openings (top, side, or rear) fitted with brush grommets (docs.nvidia.com) for clean airflow and cable strain relief. A **grounding and bonding kit** is also essential (may need to be ordered separately) to protect against electrical faults (docs.nvidia.com). In summary, racks must be rugged, EIA-standard, with full side walls and blanking panels, robust cable management, and solid grounding.

Mounting and Rails: The IT equipment should mount on adjustable rack rails meeting EIA-310 spacing (docs.nvidia.com) (docs.nvidia.com). Use **sliding rail kits** or shelves rated for heavy loads. The NVIDIA guide notes DGX servers weigh ~287 lb / 130 kg each, and a fully loaded rack (4 servers + extras) can be ~1500 lb (680 kg) (docs.nvidia.com). Thus, rack rails and shelves must be rated for the full weight of each server (287.6 lb / 130.45 kg for a DGX H100). Always use rack-mount rail kits supplied by the server OEM (if available) to support the full depth; never just hang a heavy GPU chassis by front brackets alone.

Anchoring and Safety: Bolt all racks to the floor. NVIDIA requires cabinets to be fastened to the floor per manufacturer recommendations (for example, bolting through flanges or mount points designed for that purpose) for personnel and equipment safety (docs.nvidia.com). If the facility has a raised floor, seismic bracing under the floor may be required in certain zones (docs.nvidia.com). Engage a structural engineer for high seismic areas. Check floor loading: NVIDIA's static weight and point load table lists a rack with four DGX H100 systems at about 1,500 lb (680 kg) total, with a point load of about 500 lb (226 kg) (docs.nvidia.com). Ensure the server room floor and any transport pathways (e.g. from loading dock) are rated accordingly (docs.nvidia.com).

Server Handling: Because individual HGX servers often weigh >130 kg (docs.nvidia.com), use a **proper server lift** (platform type) for installation (docs.nvidia.com). Do **not** rely on forklifts or pallet jacks: traditional forklifts fork-blades can damage components. Data-center-specific platform lifts should be rated well above the 287.6 lb (130.45 kg) weight of a DGX H100 (docs.nvidia.com). Verify clearance for lifts under any overhead cranes or ductwork.

In summary, design the physical layout and racks as for an "AI pod": floor anchors, high-weight capacity, full integration of airflow controls (blanking panels, side walls), and space for service and cable runs. These measures ensure the rack infrastructure can safely hold and operate the HGX servers without overheating or collapse.

Power Infrastructure

Power Demand and Distribution

IT Power Capacity: HGX GPU servers demand exceptionally high power. Each NVIDIA H100 GPU consumes up to **700 W** (www.tomshardware.com), while the B200 Blackwell GPU draws up to **1,000 W** (air-cooled) or **1,200 W** (liquid-cooled), and the B300 Blackwell Ultra reaches **1,100–1,400 W** (lenovopress.lenovo.com). For H100 systems, an 8×H100 server needs ~5.6 kW just for GPUs. Including CPU, memory, NVSwitch, fans and inefficiencies, total power draw per server is on the order of 10 kW. For example, analyses report **DGX H100** draws ~10.2 kW under load (semianalysis.com), while the **DGX B200** draws up to ~14.3 kW

(resources.nvidia.com). NVIDIA's own DGX SuperPOD design for H100 plans circuits at a peak demand of 10.2 kW per system (20.4 kW per circuit for two systems) (docs.nvidia.com). The SuperPOD reference architecture for B200 has been updated accordingly (docs.nvidia.com).

By contrast, a DGX A100 (8×A100) draws about **6.5 kW** max (www.netapp.com). An older Tesla-era DGX-2 (16×V100) was ~10.6 kW peak (semianalysis.com). The generational trend is stark: GPU power per chip has roughly doubled from A100 (400 W) to H100 (700 W) to B200 (1,000 W), with no sign of slowing. **Racks traditionally hold 1–4 servers**, so a standard DGX rack (4×HGX H100) at 10–11 kW each could draw ~40–45 kW. A B200-based rack approaches 60 kW, and the GB200 NVL72 full-rack system demands 120–140 kW. This far exceeds older racks (10–15 kW) and pushes data centers well beyond their traditional comfort zone (semianalysis.com).

Power Distribution (Voltage/Phase): Meeting these loads typically requires 3-phase power. NVIDIA's guidelines (Table 8) show typical rack PDU schemes: e.g. *415 VAC three-phase at 60 A (80% rated) breakers* can supply ~32.7 kW and support 4 DGX H100 systems (docs.nvidia.com). A preferred design is 415 VAC, 32 A 3Φ (N+1) – yielding ~21.8 kW per circuit (docs.nvidia.com). Multiple PDUs per rack (often 3 PDUs fed from separate breakers) balance the load. For instance, one scheme uses 3Φ Y, 400 VAC, 32 A circuits at ~21 kW per circuit; 20.4 kW is the combined maximum load of two 10.2 kW DGX H100 systems, not a per-server planning assumption (docs.nvidia.com). Four-way (4-circuit) configurations with N+1 redundancy ensure no single circuit exceeds 50% load (docs.nvidia.com).

In simpler terms, **do not attempt to run an 8-GPU HGX on single-phase 120/208 VAC distribution**. Instead, plan high-voltage 3Φ into the room, with rack-level step-down to the 230 VAC used by each server's PSUs. Many DGX/HGX power supplies are rated for 200–240 VAC input (docs.nvidia.com), so the facility can run on either 208Y or 400Y (240) as available. However, higher line voltage (400+ VAC) reduces current and I²R losses, and is strongly recommended in rack-level designs (docs.nvidia.com). Newer trends even push to **±400–800 V DC bus architectures** for ultra-dense racks (www.techradar.com) (www.tomshardware.com). For example, NVIDIA and ABB are prototyping 800 VDC power rails to support 1 MW racks (www.tomshardware.com) (www.tomshardware.com).

PDU and Circuit Layout

Rack Power Configuration: Follow NVIDIA's DGX SuperPOD recommendations meticulously (docs.nvidia.com) (docs.nvidia.com). Each rack typically uses *three* rack PDUs (rPDUs) fed by independent circuits (A, B, C phases). Each rPDU distributes to multiple outlets (e.g. C13/C19 or Honeywell connectors) for the servers. It's crucial to **balance the load**: for N+1 redundancy, no single phase or circuit should exceed ~50% of its capacity (docs.nvidia.com). For example, if a DGX H100 peak draws 10.2 kW, then two circuits per rack would already break 50%. Thus an N+1 scheme uses 3 circuits, each about 20–22 kW capacity (docs.nvidia.com), with one left as standby.

Redundancy: DGX/HGX servers themselves have multiple PSUs (six in DGX H100, four needed) (docs.nvidia.com). To avoid any single point of failure, plan for at least *three* power sources per rack: each PDU fed from a distinct upstream circuit. Each circuit should be UPS-backed and generator-backed. The goal is that losing one PDU (breaker or feed) still leaves ≥4 PSUs energized on each server (docs.nvidia.com). NVIDIA explicitly states: "*minimum of three power sources (rPDUs on separate feeds) to each rack*" with each

source sized to ~50% of peak load (docs.nvidia.com) (docs.nvidia.com). For example, if 45 kW total load, three 25 kW circuits would be typical (2×25 kW carry the load, one spare). **Key point:** implement the three-feed N+1 scheme for DGX racks; NVIDIA notes that traditional 2N redundancy with two power feeds may be used for management racks.

UPS and Generators: All feeds must tie to a UPS before any PDU. Use enterprise-level UPS (+generator) systems sized for the full rack loads. Given extremely high energy draw, facilities should ensure UPS modules have enough battery runtime (at least several minutes) at 10+ kW per rack. Also consider **phase balancing**: distribute racks across subnetworks so total load phases out evenly at the utility. Unbalanced heavy loads can trip breakers. After installation, measure and rebalance PDUs.

Wiring and Power Quality

Wiring: Use heavy-gauge cabling appropriate to current. For a 32 A circuit at 415 VAC, 4 mm² copper or larger is common. Label and color-code circuits per PDU. Route cables cleanly to avoid airflow blocks. The NVIDIA guidelines emphasize keeping cables out of airflow paths (docs.nvidia.com) (docs.nvidia.com).

Power Factor and PF Correction: GPUs often have parallel capacitive loads, leading to lower power factor (PF≈0.95 assumed). Select PDUs and breakers that assume PF in calculations (docs.nvidia.com). If needed, install PF correction banks on transformers or UPS input.

Transient Protection: Given the high economic value of GPU clusters, install surge protectors. Lightning and utility transients can be devastating. Ideally, SPD units at facility entry and at each PDU keep voltage spikes off critical systems. Grounding and bonding of racks (with those grounding kits (docs.nvidia.com)) also helps.

Monitoring: Use smart rack PDUs that measure current/voltage per outlet. Place temperature sensors in racks (e.g. front at 4U and 42U (docs.nvidia.com)) to detect hot spots. Monitoring is essential: at least the PDU should report kW per phase. Many digital PDU models can network these stats. Implement SNMP or DCIM integration.

Cooling and Environmental Controls

Thermal Loads of HGX

HGX GPUs are *intensely hot*. An NVIDIA H100 draws 350–700 W depending on form factor (PCIe versus SXM) and dumps enormous heat into its chassis. Eight SXM H100s can dump up to ~5.6 kW of waste heat – plus CPU and other losses. Thus a DGX H100 server may output up to 10–12 kW of heat into the rack's hot aisle. A full rack of 4 servers could yield ~50 kW of heat. Cooling requirements are product- and configuration-specific. NVIDIA documents the DGX B200 as an air-cooled system with 1,550 CFM airflow; liquid cooling should be planned only for the rack-scale products and configurations that specify it. NVIDIA's GB200 NVL72 rack generates 120–140 kW of heat and requires direct liquid cooling, which OEMs such as Supermicro deliver as integrated rack-scale liquid-cooling solutions ([supermicro.com](https://www.supermicro.com)). By comparison, **traditional enterprise racks** rarely exceed 10–15 kW of heat load. Industry projections warn AI racks *may approach 1 MW* within the next few years, producing heat far beyond legacy cooling schemes (www.techradar.com).

To illustrate the scale: according to Techradar, an AI rack is 20–30× hotter than a typical rack (www.techradar.com). Liquid-cooling adoption is accelerating: for example, Microsoft trials show **microfluidic cooling** (chip-level coolant channels) can triple heat removal efficiency (www.techradar.com). Even rear-door heat exchangers (like Dell's PowerCool eRDHx) are being touted, capturing 100% of equipment heat and cutting cooling energy by ~60% (www.techradar.com). In short, any design must assume *far greater heat load* and plan accordingly.

Airflow Management

Hot/Cold Aisles: The most fundamental strategy is a hot-aisle/cold-aisle layout. Orient HGX servers with front intake, rear exhaust (docs.nvidia.com). Cold aisle (fronts) should receive cool supply air, and hot aisle (backs) should be isolated. Use blanking panels in empty spaces (docs.nvidia.com) and keep the cold aisle sealed from mixing with hot return air. Avoid “chimneys” (vertical exhaust ducts) as these can impede planned airflow paths. (docs.nvidia.com).

Figure 16 in NVIDIA's guide shows grouping: servers at bottom and top with a ventilation gap in the middle to improve airflow (docs.nvidia.com). These details optimize airflow through server inlets. Follow the datasheet: many rackmount chassis allow front-to-back only; do not install any HGX servers with front-to-front (blowing directly into each other). If any network gear has reversible airflow, set them to match: most require **Power-to-Connector (P2C)** front-intake mode by default (docs.nvidia.com).

Temperature Ranges: Set up the facility's CRAC units to maintain inlet air at 18–27 °C (65–80 °F) (www.sunbirdcim.com). ASHRAE recommends ~20–24 °C for best balance (www.sunbirdcim.com). Avoid letting hot aisles exceed about 27–32 °C (80–90 °F), since above ~60 °C (140 °F) electrical equipment and cables overheat (docs.nvidia.com). If high ambient is unavoidable, consider reconfiguring some ventilation: for extreme cases NVIDIA describes a “connector-to-power” (C2P) mode where some gear ingests hot air from the front and exhausts to the rear (docs.nvidia.com), but this is atypical.

Humidity: Maintain 40–60% RH at inlet (www.sunbirdcim.com). Too dry (<40%) risks static, too humid (>60%) risks condensation on hot components. Ensure humidification/dehumidification as needed. This is critical given GPU hot exhaust can locally elevate humidity. ASHRAE TC9.9 guidelines standby should be followed (www.sunbirdcim.com).

Cooling Infrastructure

Given the enormous heat flux, **air cooling alone may be insufficient** unless carefully engineered. Even with optimized airflow, CRAC units must remove tens of kW per cabinet. Many data centers are now adopting **liquid cooling** for GPU servers.

- **CPU Rear-Door Heat Exchangers (RDHx):** As Techradar reports, companies like Dell are offering heat-exchanger door units that capture 100% of rack heat using warm water (32–36 °C) (www.techradar.com). Dell claims up to 60% cooling energy savings. Installing these on HGX racks allows rack exhaust air to be directly cooled back into supply, reducing chiller load.

- **Direct-to-Chip (D2C) Liquid Cooling:** Most server OEMs now build liquid-cooling loops, and for Blackwell-generation GPUs this has become the dominant approach. Supermicro's HGX H100 nodes use internal cold plates and an 80 kW coolant distribution unit ([dcnmagazine.com](https://www.dcnmagazine.com)), while their GB200 NVL72 systems ship with 1.8 MW in-row coolant distribution units ([supermicro.com](https://www.supermicro.com)). D2C solutions keep GPUs up to 35 °C cooler than air cooling, enabling higher inlet temperatures and more efficient operation. For B200 liquid-cooled systems, DLC-2 technology captures 98% of heat, achieving ~40% data center power savings ([corvex.ai](https://www.corvex.ai)). PUE values near 1.1 are achievable with liquid cooling, cutting facility power use by up to 27%. ZutaCore has also announced waterless (two-phase) liquid cooling solutions specifically for the HGX B300 platform (blog.zutacore.com). For a comparison of which server vendors engineer their own direct-liquid-cooling stacks in-house, see GPU Smith's [NVIDIA server vendor index](#).
- **Immersion Cooling:** Some environments forego air entirely. In immersion, servers (or blades) are submerged in dielectric oil baths. For instance, Sustainable Metal Cloud's "HyperCubes" in Singapore use immersion to reduce energy use by up to 50% versus an air-cooled DC (www.cnbc.com). This level of efficiency enables dramatically higher rack density. However, it requires specialized infrastructure (sealed tanks, oil pumps, leak detection).
- **Innovations:** Cutting-edge approaches like microfluidic coldplates (tiny channels etched on GPUs) are being trialed. Microsoft reports that in-chip microfluidic cooling removed heat up to three times better than conventional cold plates in its tests (www.techradar.com). While not yet mainstream, such methods point to future directions.

Rack Layout: In optimized air-cooled racks, place servers with no vertical block. NVIDIA suggests placing two servers at bottom, a 3–6U empty gap, then two servers above (docs.nvidia.com). This "ventilation gap" improves air mixing. Also, put horizontal PDUs near the top, leaving space below. For neutral airflow, keep power cables off the bottom if possible, as floor-level cables can create stalls.

Cooling Capacity: Given ~50–140 kW heat per rack with current hardware (and projections of 250–900 kW for next-generation systems), each server room's cooling plant must be sized accordingly. Many colocation centers cap out at 12–15 kW/rack. Air cooling alone can handle 30–40 kW per rack with optimized design ([semianalysis.com](https://www.semianalysis.com)), while direct-to-chip liquid cooling extends this to 60–120 kW. For rack-scale systems like the GB200 NVL72, operators should provision cooling infrastructure capable of dissipating 120+ kW per rack. Forward-thinking deployments are provisioning 600 kW cooling capacity even for current 120 kW systems, in preparation for successor platforms ([introl.com](https://www.introl.com)). Make sure chillers, pumps, CDUs, and airflow fans are rated for these loads. Redundant CRAC/CDU units are advised so that if one fails the load can be shifted.

Environmental Monitoring

Install **temperature and humidity sensors** in each rack. NVIDIA suggests one temp sensor at 4U height and one at 42U (rear side) per rack (docs.nvidia.com). These should tie into building management or DCIM to alert on hotspots. High-pressure ULVs or alarms should trigger if any rack gets too hot. In volatile AI workloads, real-time thermal monitoring is critical.

Given the stakes, using a modern **DCIM/digital twin** platform is recommended to predict and visualize thermal behavior. For instance, some operators use CFD modeling (per NVIDIA's case study (blogs.nvidia.com)) to simulate the HGX installation before committing.

In summary, **treat each rack as a cooling challenge on par with small process industries**. Traditional design rules loosen: be prepared to handle 5–10× the traditional rack heat. Leverage optimized layouts, blanking, airflow containment (e.g. cold-aisle curtains if needed), and advanced cooling methods to stay within temperature limits.

Network Infrastructure

Internal GPU Networking

HGX servers rely on ultra-fast internal links (NVLink/NVSwitch) to operate GPUs as one coherent mesh. This means the **topology** of GPU pods is very restrictive: servers in a pod must be physically close. SemiAnalysis notes that GPU clusters should ideally be within ~30 m of leaf switches (semianalysis.com), because the optics used (e.g. multimode QSFP56 for NVLink fabrics) have limited reach (~100 m max, but hubs are often closer). For external NVLink switch fabrics, SemiAnalysis cites an even tighter maximum cable length (~20 m) (semianalysis.com).

In practice, plan racks so that each multi-GPU pod is contiguous. This means grouping the 4–8 racks of a DGX SuperPOD within one row or adjacent rows, minimizing long cable runs. Overhead or underfloor cable trays must accommodate large numbers of 400 Gb/s fiber or copper cables. Avoid tangles: NVIDIA's guidelines emphasize that IPC (inter-rack cable length) should be a key design criterion (docs.nvidia.com) (semianalysis.com).

Uplinks and Switches

Beyond internal GPU links, provision high-bandwidth external network. Modern AI servers use 200–800 Gb/s Ethernet or InfiniBand to storage/compute networks. For example, each DGX H100 includes **eight 400 Gb/s InfiniBand links** (NVIDIA Quantum-2) (nvidianews.nvidia.com). The B300 generation introduces the **ConnectX-8 NIC** with **800 Gb/s InfiniBand** throughput, paired with **Quantum-X800** switches (nvidianews.nvidia.com). A cluster of racks, therefore, needs infrastructure (leaf/spine switches) capable of tens to hundreds of Tb/s aggregate. Ensure the data center has space and cooling for such network gear (which often require 2–4U per switch, with its own power and temperature considerations).

Cabling Infrastructure

High-density multi-mode fiber runs should be planned with slack. Use OM4/OM5 fiber in trunks. It's critical to respect the distance limits: Aim to keep NIC-to-switch cables <15–20 m to use cost-effective optical transceivers (semianalysis.com). This typically means placing the high-speed switch in the same row or by the end-of-row. If racks are deep or cables long, strand path must be carefully engineered to avoid exceeding the NVLink/InfiniBand limits. Cable trays should be rated for both power and data cables (can interfere thermally, so separate them if possible).

Power Cabling: The many PDUs and server cables also need robust passes. NVIDIA's guidelines state cable ingress ports must be brush-grommited (docs.nvidia.com). Plan overhead busways if available. Label all power cables clearly by phase & PDU.

Power-over-Ethernet / Misc

HGX platforms typically do not use PoE or low-voltage; everything is AC-driven. However, don't neglect management networks. Provide a few 1GbE (or 10GbE) ports to each rack for IPMI and telemetry. For example, most servers have one dedicated management port (as noted by vendors like ASRock (www.storagereview.com)). Ensure separate LAN cables for management traffic. Provide KVM-over-IP or console access for out-of-band management.

Network Redundancy

Given the importance of GPU compute nodes, use redundant interconnect paths. Typical HPC practice is leaf/spine dual fabrics. Each server's NICs (often dual-port) should connect to separate switches. Avoid having a single switch or single cable as a single point. Monitoring and Link Aggregation (LACP/bonding) are recommended. On the physical side, use topologically separate cable trays for redundant paths if possible (to avoid both failing under a floor breach).

In summary, network design for HGX clusters must prioritize extremely high bandwidth and low latency, with strict cable length discipline. The physical arrangement of racks and switches should minimize optical path lengths while providing redundancy.

Safety, Security, and Compliance

Regulatory Standards

Tier/Redundancy Level: Aim for data center Tier 3 or higher (TIA-942 Rated 3, EN50600 class 3) standards (docs.nvidia.com). This means concurrent maintainability: any single component (CRAC, UPS, generator, power feed) can be serviced without shutting down the cluster. At least N+1 on all major subsystems (power, cooling). Achieving this for 40–60 kW racks demands industrial-grade equipment.

Environmental Compliance: Fire and safety codes vary by region. At minimum, install fire detection (smoke/heat detectors) and suppression per local codes (docs.nvidia.com). Many AI installations use clean agent systems (e.g. FM-200 or Novec 1230) to avoid water. Ensure fire zones align with server rows. Label all high-voltage gear and provide clear emergency shutoffs.

Physical Security

HGX clusters will host "crown-jewel" IP (AI models, HPC data). Implement robust security: locked cages or biometric access to GPU racks (docs.nvidia.com). Per NVIDIA guidance, data center entry should be SOC-compliant, with badge/pass controls and CCTV (docs.nvidia.com). Restrict rack access: if other IT racks are in the same room, isolate HGX racks with partitions/fencing that require extra clearance to reach (docs.nvidia.com). Unsupervised personnel should have no access to GPU systems. Disable front-panel USB/Thunderbolt ports to prevent unauthorized code injection.

Worker Safety (Noise and Ergonomics)

NVIDIA states that DGX H100 systems can generate noise levels greater than 98 dB at 1 m distance (docs.nvidia.com). This is well above typical office or DC safeguards. Install sound-reducing enclosures or provide hearing protection for technicians when inspecting racks. If possible, locate HGX clusters in dedicated sound-controlled areas. Follow OSHA or local regulations for machine noise. Lightning/EMS safety is also important: proper cable shielding and grounding mitigate static discharges (www.sunbirdcim.com).

Environmental Controls

Maintain room environmental monitoring (CRAC controls) to respond to excursions. Consider raising the chilled-water setpoint when possible (especially if using rear-door cooling from 32–36 °C water (www.techradar.com)). This reduces energy use. Implement airflow alarms: for instance, if rack inlet temp > threshold (e.g. 27 °C) or PDU load spikes. Use DCIM tools to correlate sensor data; NVIDIA's Omniverse blog highlights using digital twins for «AI supercomputing» facilities (blogs.nvidia.com). (We discuss this in Future Directions.)

Case Studies and Examples

The challenges above are not hypothetical. Broad industry examples illustrate these requirements:

- **Hyperscale Adopters:** NVIDIA reports that in 2017–2018, global cloud and hyperscale vendors (AWS, Microsoft, Google, Facebook) incorporated HGX design into their data centers (nvidianews.nvidia.com). For example, Microsoft's *Project Olympus* and Facebook's *Big Basin* were HGX-inspired designs (nvidianews.nvidia.com).
- **NVIDIA Eos Supercomputer:** Internally, NVIDIA's own blueprint is instructive. NVIDIA's *Eos* system uses **576 DGX H100 nodes (4,608 H100 GPUs)** (nvidianews.nvidia.com). Conservatively estimating 10 kW per node, Eos alone needs ~5.76 MW of IT power, plus cooling. Such scale required NVIDIA to utilize new data center designs. (NVIDIA headlines note Eos will be a blueprint for OEMs.)
- **Supermicro Liquid-Cooled Racks:** Supermicro launched fully liquid-cooled 8-GPU HGX racks. They reported ~40% total data center power savings and ~86% reduction in cooling energy use (versus air cooling) (dcnmagazine.com). In one case, Supermicro delivered a rack with an 80 kW CDU capable of direct liquid cooling (dcnmagazine.com). Such racks come pre-integrated (piping, leak detection, remote pump units) to simplify site setup.
- **Immersion Cooling (Sustainable Metal Cloud):** A startup in Southeast Asia built “HyperCube” flexible AI containers by immersing NVIDIA GPUs in cooling oil (www.cnbc.com). This reduced energy consumption up to 50% vs. air-cooling, cutting capex by ~28%. They can deploy in unused spaces of standard data halls (www.cnbc.com) (www.cnbc.com). While full oil-immersion remains unusual, this case shows how high-density GPU pods can be supported when traditional cooling fails.
- **Colocation AI Hubs:** Major colocation providers now offer purpose-built AI infrastructure. Digital Realty, Equinix, and others have launched “AI-ready” data hall configurations with pre-installed power (30–50 kW/rack for H100-class, 120+ kW/rack for GB200 NVL72), liquid cooling infrastructure, and high-

bandwidth connectivity, catering to customers needing turnkey GPU clusters. This reflects a broader trend: colos routinely offer 30–50 kW per rack for AI workloads, far above prior 10–20 kW norms, with some hyperscale facilities now provisioning 100+ kW per rack.

- **Game-Changing Startups:** TechRadar coverage highlights new tech: for example, Dell's eRDHx rear-door cooler (www.techradar.com) and microfluidic startups (www.techradar.com). Also, press coverage of NVIDIA's 800 VDC project with ABB (www.tomshardware.com) shows how supply chain is adapting power infrastructure.

Each of these examples reinforces that deploying HGX hardware is not a plug-and-play upgrade – it often requires custom facility solutions. We next distill the general design lessons from these cases.

Implications and Future Directions

NVIDIA HGX deployments represent a **broad industry inflection point**. The physical requirements transcend any single vendor solution: power grids, HVAC engineering, and data center planning must evolve. We highlight key implications and trends:

- **Extreme Power Density:** The pace of power density growth has accelerated beyond earlier projections. Current Blackwell-generation racks (GB200 NVL72) already demand 120–140 kW, and NVIDIA's roadmap shows this trajectory continuing: B300-based systems push higher still, while the upcoming Rubin generation (Rubin in 2026, Rubin Ultra in 2027) is expected to require 250–900 kW per rack (datacenterknowledge.com). This pushes the industry toward high-voltage distribution. NVIDIA and partners (ABB, Eaton, Schneider Electric, Vertiv) are actively developing **800 VDC** power architectures for 1 MW racks, with ecosystem readiness targeted for 2027 to coincide with the Rubin Ultra platform (datacenterfrontier.com). Vertiv's 800 VDC power portfolio is planned for release in H2 2026 (vertiv.com). The trend is clear: **data centers will adopt distribution architectures akin to power substations** to feed GPU farms.
- **Advanced Cooling Adoption:** Air cooling alone is reaching limits. The industry is rapidly iterating: rear-door exchangers, full immersion, microfluidics, and liquid loops all aim to cut PUE (power usage effectiveness). NVIDIA itself is experimenting with simulation platforms (digital twins) to optimize cooling layout (blogs.nvidia.com) (blogs.nvidia.com). Predictive tools will become essential to avoid costly trial-and-error.
- **Renewables and Energy Management:** With GPU facilities consuming MW of power, grid impact is huge. Some AI operators (xAI, OpenAI) are already securing dedicated power plants (www.cnbc.com) (www.tomshardware.com). Data centers might co-locate with solar/wind, or use on-site generation (gas turbines as Musk's team has done (www.tomshardware.com)). There will also be more emphasis on **dynamic load scheduling** (trading flexibility for cheaper rates). Monitoring these massive loads in real-time and tying into energy markets will be a new discipline.
- **Sustainability:** Pressure is mounting to improve efficiency. TrendForce noted that by late 2024 only ~10% of new AI chips used liquid cooling (www.trendforce.com), but cooling requirements vary across Blackwell systems; NVIDIA documents DGX B200 as air-cooled. With GB200 NVL72 shipments

ramping through 2025 (TrendForce projected ~60,000 racks), liquid cooling is becoming the default for high-end AI infrastructure. Techniques like heat reuse (capturing GPU heat for facility heating) and integrating data center cooling with building HVAC are gaining attention. The goal will be not just to keep systems running, but to do so within tight CO₂ budgets.

- **Standardization and New Guidelines:** Industry standards are evolving rapidly. NVIDIA has published updated SuperPOD reference architectures for B200-based deployments (docs.nvidia.com) alongside the existing H100 guides (docs.nvidia.com). These documents, along with the Open Compute Project (OCP) ORV3 rack standard used by the GB200 NVL72, are serving as templates for new industry codes. UPS and power distribution manufacturers are designing modules capable of 150+ kW per rack, and CDU vendors are scaling to multi-megawatt cooling capacity.
- **Digital Twins and Automation:** As hinted by NVIDIA's Omniverse blog (blogs.nvidia.com) (blogs.nvidia.com), future data centers will use virtual models to simulate airflows, power flows, and networking before physical deployment. This is especially valuable for HGX clusters, where hot spots and power transients can be predicted. Sensors will be integrated and AI-driven control loops may adjust cooling in real time.
- **Resilience and Business Continuity:** With racks consuming tens of kilowatts each, even a brief outage incurs enormous risk. Facilities must plan for very high availability. For instance, backup generators must handle MW loads – requiring multiple gensets or load shedding plans. More features like redundant HVAC (N+1 chillers, CRAHs) and fire suppression fail-safes (e.g. dry-pipe pre-action sprinklers) will be mandated.
- **Economic and Regulatory Impact:** The scale of AI data centers is drawing increasing regulatory attention. Multiple U.S. states have introduced legislation around data center energy use and grid impact, and European regulators are enforcing energy efficiency reporting requirements. Building permits for new GPU centers increasingly require grid impact assessments and efficiency guarantees. Electric tariffs may include demand charges based on 15-min peaks. The global data center infrastructure market is projected to reach \$1 trillion in spending by 2030, driven primarily by AI (iot-analytics.com). All these factors feed back into the physical design: e.g., energy storage (batteries or flywheels) and on-site generation are increasingly placed on-site to shave peaks and ensure power availability.

In conclusion, deploying NVIDIA's HGX platform is not a simple server upgrade – it heralds a new era of data center design. Physical planning must incorporate cutting-edge power distribution, cooling, and automation to handle the extreme loads.

Realizing this vision will demand cross-disciplinary collaboration (IT, mechanical, electrical, and civil engineers working together) and adherence to emerging best practices. It will also create challenges (cost, complexity) but the industry response is rapid and innovative.

Conclusion

NVIDIA's HGX platform represents the **frontier of data center hardware** – massive multi-GPU systems that unlock unprecedented compute performance. But harnessing that performance requires **meticulous engineering of the physical site**. This report has detailed every aspect “A–Z” of such engineering for HGX deployments: from site selection, rack and power layout, to cooling systems and safety. The key takeaways for practitioners are:

- **Space:** Allocate enough dedicated floor area, with proper aisle design and expansion room.
- **Rack Infrastructure:** Use robust, well-organized racks (48U+, proper grounding, blanking). Ensure floors can bear hundreds of kilograms per rack (docs.nvidia.com).
- **Power:** Install multi-megawatt 3-phase distribution. Plan at least 40–50 kW per rack for H100-class systems, 60+ kW for B200-class, and 120–140 kW for GB200 NVL72 rack-scale systems. Use multiple UPS-backed circuits per rack (N+1) (docs.nvidia.com) (docs.nvidia.com).
- **Cooling:** Engineer for extreme heat loads via hot-aisle/cold-aisle containment and modern cooling (direct liquid cooling, immersion). Cooling requirements are product- and configuration-specific: DGX B200 is air-cooled, while rack-scale products or configurations that specify liquid cooling require appropriate CDU infrastructure and plumbing. Upgrade CRAC/CDU capacity accordingly.
- **Networking:** Cluster racks closely to keep high-speed fabric cable runs short (semianalysis.com). Provide ample InfiniBand/Ethernet switching with redundant paths.
- **Safety & Security:** Adhere to Tier-3+ standards for reliability (docs.nvidia.com). Equip with fire suppression, surveillance, and physical access controls (docs.nvidia.com). Account for noise and provide protection (docs.nvidia.com).
- **Monitoring & Automation:** Deploy sensors and DCIM to watch temperature, load, and use digital-twin modeling (blogs.nvidia.com) (blogs.nvidia.com) where possible.

Each claim in this report is backed by current references – from NVIDIA's own design guides to industry news and technical analyses. As we have shown with cited data and examples, the requirements are concrete and quantifiable. In deploying HGX, data centers are effectively running *small power plants*. Success requires the same rigor as a heavy industrial project.

Looking ahead, such advanced GPU infrastructures will only grow in scale and sophistication. The transition from Hopper (H100/H200) to Blackwell (B200/B300) and soon to Rubin (expected in the second half of 2026, with Rubin Ultra in 2027) represents a continuous escalation in power and cooling demands – from ~10 kW per server to 14+ kW, and from 40 kW per rack to 120–140 kW today with projections of 250–900 kW by 2027. Power and cooling technologies are evolving in parallel: the ABB/NVIDIA/Vertiv 800 VDC ecosystem is targeting 2027 deployment (datacenterfrontier.com), liquid cooling is required for the specific rack-scale products and configurations that specify it, while NVIDIA documents DGX B200 as air-cooled; the industry is moving from megawatt to gigawatt-scale facility planning. Regulatory and sustainability pressures will tighten. Data center professionals must stay informed of the latest developments (e.g. NVIDIA's GTC announcements, updated SuperPOD reference architectures, DOE reports) and adapt their designs iteratively.

In sum, NVIDIA HGX platform deployment fundamentally changes the game: it is not about merely adding servers, but rethinking the entire data center ecosystem. By following the design principles, planning guidelines, and examples documented here, architects can build facilities that safely and efficiently support the HPC workloads of tomorrow.

References: (Inline citations above link to authoritative sources including NVIDIA design documents (docs.nvidia.com) (docs.nvidia.com), industry technical blogs and news (nvidianews.nvidia.com) (www.cnbc.com), and standards/guidelines summaries (www.sunbirdcim.com) (docs.nvidia.com).) Each key claim is supported by one or more cited references as indicated.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. <https://intuitionlabs.ai/articles/openai-amd-ai-hardware-partnership>
2. <https://nvidianews.nvidia.com/news/nvidia-announces-dgx-h100-systems-worlds-most-advanced-enterprise-ai-infrastructure>
3. <https://semianalysis.com/2024/03/13/ai-datacenter-energy-dilemma-race/>
4. <https://resources.nvidia.com/en-us-dgx-systems/dgx-b200-datasheet>
5. <https://www.sunbirdcim.com/blog/your-data-center-ready-nvidia-gb200-nvl72>
6. <https://docs.nvidia.com/dgx-superpod/design-guides/dgx-superpod-data-center-design-h100/latest/infrastructure.html>
7. <https://docs.nvidia.com/dgx-superpod/design-guides/dgx-superpod-data-center-design-h100/latest/electrical.html>
8. <https://dcnmagazine.com/infrastructure/cooling/supermicro-nvidia-servers/>
9. <https://www.cnbc.com/2024/08/27/nvidia-partner-sustainable-metal-cloud-ai-data-center-energy-consumption.html>
10. <https://www.reuters.com/business/energy/us-data-center-power-use-could-nearly-triple-by-2028-doe-backed-report-says-2024-12-20/>
11. <https://www.techradar.com/pro/security/this-graph-alone-shows-how-global-ai-power-consumption-is-getting-out-of-hand-very-quickly-and-its-not-just-about-hyperscalers-or-openai>
12. <https://www.techradar.com/pro/overheating-is-a-big-problem-for-ai-hardware-as-demand-rises-and-dell-thinks-it-might-have-the-answer>
13. <https://blogs.nvidia.com/blog/omniverse-digital-twin-data-center/>
14. <https://intuitionlabs.ai/articles/key-innovations-behind-chatgpt>
15. <https://intuitionlabs.ai/articles/hardware-requirements-gpt-oss-20b>
16. <https://nvidianews.nvidia.com/news/nvidia-introduces-hgx-2-fusing-hpc-and-ai-computing-into-unified-architecture-6696445>
17. <https://nvidianews.nvidia.com/news/nvidia-partners-with-world-s-top-server-manufacturers-to-advance-ai-cloud-computing>
18. <https://www.tomshardware.com/pc-components/gpus/nvidia-announces-blackwell-ultra-b300-1-5x-faster-than-b200-with-288gb-hbm3e-and-15-pflops-dense-fp4>
19. <https://www.nvidia.com/en-us/data-center/gb200-nvl72/>
20. <https://intuitionlabs.ai/articles/openai-stargate-datacenter-details>
21. <https://blog.se.com/datacenter/2025/10/16/the-1-mw-ai-it-rack-is-coming-and-it-needs-800-vdc-power/>
22. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand>

23. <https://www.tomshardware.com/tech-industry/big-tech/nvidia-800-vdc-power-rollout-for-1-megawatt-server-racks-to-be-supported-by-abb-company-says-collaboration-will-create-new-power-solutions-for-future-gigawatt-scale-data-centers>
24. <https://www.supermicro.com/en/accelerators/nvidia>
25. <https://www.tweaktown.com/news/103938/nvidia-gtc-2025-gb300-ai-gpu-with-1-4kw-power-new-details-on-rubin-cpo-tech-and-more/index.html>
26. <https://www.tomshardware.com/tech-industry/nvidias-h100-gpus-will-consume-more-power-than-some-countries-each-gpu-consumes-700w-of-power-35-million-are-expected-to-be-sold-in-the-coming-year>
27. <https://lenovopress.lenovo.com/lp2226-thinksystem-nvidia-b200-180gb-1000w-gpu>
28. <https://www.trendforce.com/presscenter/news/20240730-12232.html>
29. <https://developer.nvidia.com/blog/defining-ai-innovation-with-dgx-a100/>
30. <https://www.netapp.com/blog/nvidia-dgx-a100/>
31. <https://nvidianews.nvidia.com/news/nvidia-blackwell-ultra-ai-factory-platform-paves-way-for-age-of-ai-reasoning>
32. <https://docs.nvidia.com/dgx-superpod/reference-architecture-scalable-infrastructure-b200/latest/dgx-superpod-architecture.html>
33. <https://www.sunbirdcim.com/glossary/data-center-temperature-and-humidity-standards>
34. <https://www.corvex.ai/blog/liquid-cooling-nvidia-b200-ai-infrastructure>
35. <https://blog.zutacore.com/press-releases/zutacore-asrock-rack-waterless-liquid-cooling-for-nvidia-hgx-b300>
36. <https://gpusmith.com/vendors>
37. <https://introl.com/blog/gb200-nvl72-deployment-72-gpu-liquid-cooled>
38. <https://www.storagereview.com/news/asrock-rack-unveils-powerful-ai-training-systems-with-nvidia-hgx-h100-h200-8-gpu-support>
39. <https://www.techradar.com/pro/ai-workloads-are-reshaping-infrastructure-heres-what-data-centers-need-to-know>
40. <https://www.datacenterknowledge.com/energy-power-supply/from-mw-to-gw-how-ai-is-forcing-a-complete-rethink-of-data-center-power>
41. <https://www.datacenterfrontier.com/energy/article/55323139/preparing-for-800-vdc-data-centers-abb-eaton-support-nvidias-ai-infrastructure-evolution>
42. <https://www.vertiv.com/en-us/about/news-and-insights/corporate-news/from-vision-to-readiness-vertiv-collaborates-with-nvidia-to-advance-800-vdc-platform-designs-to-power-the-next-generation-of-ai-factories/>
43. <https://www.tomshardware.com/tech-industry/artificial-intelligence/elon-musk-says-xai-is-targeting-50-million-h100-equivalent-ai-gpus-in-five-years-230k-gpus-including-30k-gb200s-already-reportedly-operational-for-training-grok>
44. <https://iot-analytics.com/data-center-infrastructure-market/>

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore [AI enablement](#), [pharma and biotech software development](#), [data engineering and BI](#), and [Veeva services](#). [Contact IntuitionLabs](#) to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: <https://intuitionlabs.ai>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.