



GB200 NVL72 Day-2 Operations: Failure Domains

Plan and build private AI compute with GPU Smith. We connect workload requirements with GPU hardware, networking, deployment and operating decisions so your infrastructure fits the work it must perform.

Published by GPU Smith · 2026-10-29 · gb200 nvl72 power and cooling requirements · gb200 nvl72 maintenance · nvl72 failure domain blast radius · nvlink domain health monitoring nmx · nvl72 compute tray replacement · supermicro nvidia nvl72 shipping · nvl72 leak detection liquid cooling operations

Original article: <https://gpusmith.com/articles/en/gb200-nvl72-day-2-operations-failure-domains>



In this article

1. [Why NVL72 changes day-2 operations](#)
 2. [Physical architecture](#)
 3. [Failure domain map](#)
 4. [Cooling and leak detection](#)
 5. [NVLink domain health](#)
 6. [Out-of-band management](#)
 7. [Partitioning and degraded operation](#)
 8. [Service procedures and spares](#)
 9. [Firmware and OS updates](#)
 10. [Monitoring and alerting stack](#)
 11. [Runbook outline and SLA checklist](#)
 12. [Sources](#)
-

Why NVL72 changes day-2 operations

An 8-GPU HGX server is a self-contained failure domain: a [DGX B200](#) carries its own 2 NVSwitch chips, so its NVLink fabric fails with the server and nowhere else. GB200 NVL72 is "a 72-GPU NVIDIA NVLink domain that acts as a single, massive GPU" ([NVIDIA GB200 NVL72](#)). To get there, every compute tray shares switch trays, a copper backplane, a fabric control plane, power shelves and liquid manifolds.

For each component, operators need to know how many GPUs lose NVLink when it fails and which maintenance takes the whole rack offline. The general spares and SLA method is in our [GPU cluster reliability budget](#); this article covers what is specific to the rack.

Scope note: NVIDIA's DGX user guide covers both DGX GB200 and DGX GB300. GB200 trays use 4x ConnectX-7 400G NICs and GB300 trays use 4x ConnectX-8 800G ([DGX GB200 user guide: Hardware](#)). OEM NVL72 racks follow the same reference architecture, but their service procedures and support terms come from the OEM.

Physical architecture

Per NVIDIA's DGX GB200 user guide ([Hardware](#)), an NVL72 rack contains:

- **18x 1RU compute trays**, "each with 2 Grace CPUs and 4 Blackwell GPUs", each running its own OS image, with 4x 3.84TB E1.S NVMe cache drives on DGX.
- **9x 1RU NVLink switch trays**. Each has 2 NVSwitch chips with 72 NVLink ports per chip, NVOS 2.0 or later, a BMC "for attestation and recovery", and the Fabric Manager (FM) and Subnet Manager (SM) control planes running on the switch tray CPU.
- **2x top-of-rack management switches**.

- **Power shelves** that convert AC to "nominal 50V-51V DC" on a bus bar. The guide specifies "six air-cooled 5.5kW PSUs in eight power shelves that provide N+N redundancy and the required input power of 33kW per power shelf", and puts rack consumption at "approximately 120kW".
- An **"NVLink passive copper cable cartridge backplane"** joining the trays.

NVIDIA's OCP contribution blog (October 15, 2024) adds that ([NVIDIA OCP blog](#)) the rack houses "18 compute trays, nine switch trays, and four NVLink cartridges". The cartridges hold "over 5,000" copper cables. The blog calls 120 kW the "cooling capacity required for the rack"; the user guide gives the same number as power consumption. NVIDIA's product page lists 13.4 TB HBM3E and 130 TB/s of NVLink bandwidth per rack ([GB200 NVL72](#)).

Port arithmetic that matters for failures

The NVLink port counts close exactly, which explains why a switch tray failure reaches every GPU in the rack:

Quantity	Arithmetic	Result
NVSwitch ports in the rack	9 trays x 2 chips x 72 ports	1,296
GPU NVLink links in the rack	72 GPUs x 18 links	1,296
Links per GPU per switch tray	18 links / 9 trays	2
Per-GPU NVLink bandwidth	18 links x 50 GB/s per direction	900 GB/s per direction (1.8 TB/s bidirectional)

The link figures come from NVIDIA's MNNVL verification procedure, which shows 18 active links per GPU to nine switch trays and requires that "All NVLinks must be up and reporting 50 GB/s bandwidth" ([MNNVL user guide: Verifying](#)). The 1.8 TB/s figure is from the [DGX GB200 networking page](#). Every GPU has links into every switch tray.

Failure domain map

NVIDIA's Multi-Node NVL Partition User's Guide (DU-12143-001 v1.0, September 18, 2026, which supersedes the July 2025 GB200 v0.3 guide and adds GB300 and Vera Rubin) documents how the NVLink control plane reacts to each class of hardware fault ([Multi-Node NVL Partition User's Guide](#)). The answer depends on the NVLink software release, which the guide labels GB200 1.x and GB200/300 2.x. The key statements:

- **GPU or access link failure:** the workload in the partition fails. The control plane marks the GPU NO_NVLINK, and "the administrator can use the partition to run workloads on the remaining NVLink capable GPUs".
- **Compute tray failure:** all four GPUs in the tray are marked NO_NVLINK, the tray is marked UNHEALTHY, and "the health of the partition remains unchanged".
- **Switch tray or NVSwitch failure:** in a single-chassis platform it "initially causes the GPUs in all partitions to lose NVLink connectivity". On 1.x software, partition health goes to UNHEALTHY and administrators "cannot use the partition to run workloads that use NVLink" (standalone GPUs still work).

On 2.x, the rack tolerates "one failed switch tray that does not host NMX-C, at reduced bandwidth. After the initial application failure and Control Plane reconfiguration, applications can restart." Multiple failed trays or a failed NMX-C tray "require taking the domain out of service".

- **Switch tray and cable cartridge maintenance:** the v1.0 guide says both sections "will be updated in a future release". The July 2025 v0.3 guide said to "ensure that all partitions in the NVLink domain are deleted" before either is serviced ([v0.3](#)); keep that rule until NVIDIA or your OEM says otherwise.

NVL72 is a single-chassis platform. The arithmetic says one switch tray carries 2 of each GPU's 18 links (11%). On 1.x software the rack's NVLink capability is lost, not reduced; on 2.x the rack restarts at roughly that reduced bandwidth, unless the failed tray hosts the control plane (NMX-C). Check which release your racks run before using the table below.

Blast-radius table (our analysis)

The table combines NVIDIA's documented behaviour with our arithmetic. NVIDIA does not publish it, and it contains no failure rates.

Component (count per rack)	Documented behaviour	GPUs losing NVLink	NVLink-capable GPUs left
One GPU or access link (72)	Running partition job fails; GPU NO_NVLINK	1	71 (68 if the scheduler drains the whole tray)
Compute tray (18)	4 GPUs NO_NVLINK; partition health unchanged	4	68
Compute tray liquid leak (18)	Tray BMC starts a shutdown timer	4	68
NVLink switch tray or NVSwitch (9)	All partitions lose NVLink; 1.x stays down, 2.x restarts at reduced bandwidth if the tray does not host NMX-C	72 (1.x); 72 until restart (2.x)	0 (1.x); 72 at ~8/9 bandwidth (2.x)
Switch tray liquid leak (9)	BCM powers off the tray via Redfish	as the row above	as the row above
Cable cartridge (4)	Failure handling undocumented; servicing requires deleting all partitions	72 for the service window	0 during service
PSU (48) or power shelf (8)	N+N; PSU swap allowed while 4 or more shelves are healthy	0 if redundancy holds	72
Rack-level leak	Power shelf DC output shut off; BMS cuts breakers and valves	72	0
Row-level leak	BMS cuts all racks in the row and the row CDU	72 x racks in row	0

Leak rows come from NVIDIA Mission Control's BMS integration guide (covered below). PSU rows come from the [DGX GB200 PSU replacement procedure](#), which allows replacement "as long as at least four power shelves are fully active and healthy".

Worked example: capacity loss in GPU-hours

Consider one repair window of T hours. The repair time is a hypothetical for the arithmetic, not a measured or vendor figure:

- Compute tray out for $T = 8$ h: 4 GPUs x 8 h = **32 GPU-hours** of NVLink capacity lost, 5.6% of the rack (4/72).
- Switch tray out for $T = 8$ h on 1.x software: 72 GPUs x 8 h = **576 GPU-hours** of NVLink capacity lost, 100% of the rack's NVLink capacity. Single-GPU work can still run. On 2.x (tray not hosting NMX-C), the loss is one job restart plus about 11% of NVLink bandwidth for 8 h.
- HGX comparison: a DGX B200 node out for 8 h costs $8 \times 8 =$ **64 GPU-hours**, and no other node is affected.

NVL72's common failure unit is smaller than HGX's (4 GPUs against 8), but each of its 13 shared components (9 switch trays, 4 cartridges) can take NVLink from all 72 GPUs, and on 2.x software a second switch tray failure still does. Size spares and SLA remedies on that asymmetry, not the average.

Job sizing headroom

Parallelism degrees are usually powers of two, so a natural large job uses 64 GPUs (16 trays), leaving 2 trays of headroom for compute tray losses. A third tray loss, or a switch tray loss on 1.x software, forces it onto another rack or a smaller shape. Write the 64-or-72 choice down.

Cooling and leak detection

Compute trays are cooled by liquid that "runs up and down the rack through manifolds, then through the cold plates that are attached to the CPUs and the GPUs". Networking and storage devices "are air cooled" ([DGX GB200 user guide: Hardware](#)), so the room still carries a residual air load. Trays mate to the manifold through blind-mate liquid connections ([NVIDIA OCP blog](#)). First fill and pressure testing are day-0 work, covered in our [cold-plate rack commissioning checklist](#).

For day-2 operations, the key document is NVIDIA Mission Control's BMS integration guide (2.3.1, last updated October 12, 2026) ([Mission Control BMS integration](#)). It defines three detection tiers:

1. **Tray level:** cold plate leak sensor and inner manifold leak sensor, handled by the compute or switch tray BMC.
2. **Rack level:** leak sensing ropes and spot sensors along piping and inside the rack, handled by the customer's building management system (BMS).
3. **Data center level:** sensors in CDUs and along facility piping.

The guide's fault table assigns responses:

Event	BMC	Mission Control (BCM)	BMS
Compute tray leak	Activates shutdown timer	Receives leak event	—

Event	BMC	Mission Control (BCM)	BMS
Switch tray leak	Notifies BCM	Powers off the tray over out-of-band Redfish; opens a service ticket; notifies BMS	—
Rack leak (detected by BCM)	—	Shuts off power shelf DC output immediately; notifies BMS	Opens all rack breakers; closes supply and return valves
Rack leak (detected by BMS)	—	Notified	Opens rack breakers; closes valves
Row leak	—	Notified	Opens breakers for all racks in the row; shuts down the row CDU
Sensor fault	—	Calls for onsite inspection	Calls for onsite inspection

Consequences:

- **It is licensed.** The feature "is not available in standalone BCM deployments" and needs a customer-supplied BMS with an MQTT broker. Without Mission Control, you build the tray-BMC-to-BMS path yourself.
- **A switch tray leak is a rack event.** The leaking tray is powered off, and under the partition guide's rule every GPU in the rack loses NVLink.
- **Sensor faults get no automatic action**, only onsite inspection.

NVLink domain health

NVLink health has more layers than on HGX, each with its own tool:

Layer	Check	Healthy result	Source
GPU fabric registration	<code>nvidia-smi -q grep 'Fabric' -A 4</code>	State Completed, status Success	MNNVL verifying
Per-GPU links	<code>nvidia-smi nvlink --status</code>	18 links up at 50 GB/s each	same
Peer-to-peer topology	<code>nvidia-smi topo -p2p n</code>	OK for all GPUs; otherwise <code>nvidia-smi -r</code>	same
Persistence and IMEX daemons	<code>systemctl status nvidia-persistenced, systemctl status nvidia-imex; nvidia-imex --version</code> identical on every tray	running, same version	same
Domain membership	<code>nvdomaininfo</code>	Success per GPU (In progress while joining)	Mission Control NVLink info
Fabric control plane	<code>nvfabricinfo --raw</code>	nm-x-controller and nm-x-telemetry status ok on the leader	same
Partitions	<code>nvsdnpartitioninfo</code>	partition health healthy (not degraded)	same

Layer	Check	Healthy result	Source
Rack hardware	<code>sudo nvsm show health</code>	all checks and overall status Healthy	DGX GB200 health check

The control plane runs on one switch tray, the leader. `nvfabricinfo --start <switch>` moves it: "The NMX controller starts on the new leader and stops on the previous one." Run it before servicing the leader's tray. If the control plane is down, compute trays can still be moved or replaced, and the control plane reconciles against its saved topology snapshot when it returns ([Multi-Node NVL Partition User's Guide v1.0](#)). On 2.x software a failed tray hosting NMX-C is the one switch tray failure the rack cannot ride through, so know which tray is the leader.

IMEX, which maps GPU memory across trays over NVLink ([DGX GB200 user guide: Software](#)), has its own failure semantics. By default it "will not do any clean up until the connection is established again". If `IMEX_NODE_DISCONNECTED_GRACE_TIME` is set, cleanup runs after that period and marks "all jobs that had imported from that node" invalid ([IMEX guide: Connections](#)). The page states no default. Set one explicitly so that a dead tray fails jobs instead of hanging them.

Out-of-band management

Out-of-band management on a compute tray is "provided by a combination of Board Management Controller (BMC) and Host Management Controller (HMC) microcontrollers". Each switch tray has its own BMC ([DGX GB200 user guide: Hardware](#)). The [networking port tables](#) put these ports on the `ipminet` management network:

- Compute tray: BMC 1GbE plus BlueField-3 BMC RJ45s. The hardware table lists "2x 1GbE x RJ45 from the BlueField-3 BMC interface" (one per DPU); the port table shows one. Size for the hardware table: 3 ports.
- Switch tray: BMC 1GbE plus two 1GbE switch management ports, so 3 ports. Each switch also has an RS232 console at 115200 baud.

Per rack: $18 \times 3 + 9 \times 3 = \mathbf{81 \text{ management ports}}$ (63 if only one BlueField-3 BMC port per tray is cabled) into the 2 management ToR switches, before power-shelf management. Leak-driven Redfish power-offs travel over this network, so treat it as part of the safety path.

Partitioning and degraded operation

NVLink partitions are the scheduler's view of the rack. A Default Partition is created at boot (`MNNVL_ENABLE_DEFAULT_PARTITION`, default 1). Its default type is UID-based (`MNNVL_DEFAULT_PARTITION_TYPE`, default 2), and its default resiliency mode is adaptive bandwidth (`MNNVL_DEFAULT_RESILIENCY_MODE`, default 2) ([Multi-Node NVL Partition User's Guide v1.0](#)). On 2.x a GPU no longer needs a reset after an access link returns; on 1.x it does. The guide's health states give a scheduler three classes:

- HEALTHY: full bandwidth, full capacity.

- `DEGRADED` or `DEGRADED_BANDWIDTH`: operational with parked GPUs or reduced bandwidth. These states arise from trunk-link loss in multi-chassis designs.
- `UNHEALTHY`: not operational. Loss of a switch on 1.x software is one of the listed causes; on 2.x a switch tray reboot moves an affected partition to `DEGRADED_BANDWIDTH` instead.

On Kubernetes, NVIDIA's DRA driver exposes the domain as `ComputeDomains` and, in its default driver-managed mode, runs one IMEX DaemonSet per `ComputeDomain`. In host-managed mode, "A Ready `ComputeDomain` therefore does not report the health of the host service" ([NVIDIA DRA driver: ComputeDomains](#)), so admission must check IMEX health separately. Gang admission is covered in [gang scheduling on private clusters](#). The NVL72-specific rule: partition and switch-tray health gate placement for the whole rack, not per node.

Service procedures and spares

The public DGX GB200 service manual lists only three customer-replaceable units: the power supply module, the power supply management module and the E1.S cache drive. For everything else it says to "Contact NVIDIA Enterprise Support for replacement instructions", and "When replacing a component, use only the replacement supplied to you by NVIDIA" ([DGX GB200 service manual](#)). The same manual recommends `sudo nvsm stress-test --force` (about 20 minutes) before a system goes into production.

The partition guide documents the logical side of a compute tray swap. Move the tray into an "Out For Repair" (OFR) partition with `RemoveGpusFromPartition(..., reroute=false)` and `AddGpusToPartition`, optionally power it off and remove it, reinsert it, then return it to its partition, while the rest of the domain keeps running. For physical procedures the guide points to a "GB200 NVL Service Flow User Guide" available from the OEM. We found no public document that says whether a compute tray can be hot-swapped on a live liquid loop, or that gives repair-time targets, so we make no hot-swap or MTBF claim.

Field reliability data is thin. In August 2025 SemiAnalysis wrote that "the NVLink copper backplane still is not that reliable" and that backplane debug tools were "behind and sub-optimal" ([SemiAnalysis, August 20, 2025](#)). Treat that as a dated reason to write strict acceptance terms, not as a current failure rate.

Spares follow the blast-radius table: PSUs and E1.S drives on site, ideally one compute tray too, and switch trays and cartridges under contracted time targets, since they set the rack's worst-case outage.

Firmware and OS updates

NVIDIA firmware is "typically released as a recipe to update all applicable components simultaneously" ([DGX GB200 user guide: System Control](#)). With `nvfwupd` 2.0.5 or later, firmware can be pushed to multiple trays in parallel. Activation is disruptive: compute trays are activated with `PWR_OFF` followed by `RF_AUX_PWR_CYCLE`, and switch trays with `NVUE_PWR_CYCLE` ([nvfwupd parallel update](#)). Under the switch tray rule above, activating switch tray firmware is a rack-wide NVLink outage. Plan it as a rack maintenance window, not a rolling update.

OS updates are lighter: a DGX compute tray checks for image updates at boot ([System Control](#)), which allows tray-by-tray rolls through OFR partitions.

Worked example: cold reboot duration

NVIDIA's documented cold reboot sequence includes these fixed waits ([DGX GB200 rack reboot sequence](#)):

- 3 to 5 minutes for PSU capacitors to discharge
- 2 minutes for tray BMCs to come up
- 2 minutes for switch trays to boot NVOS
- 5 minutes after powering on compute nodes

Total: (3 to 5) + 2 + 2 + 5 = **12 to 14 minutes**, before switch tray provisioning (L11 diagnostics or NMX-C setup) and post-reboot checks. The warm reboot skips the power shelves: it powers compute trays off through the BMC, restarts NMX-C and powers them back on.

Monitoring and alerting stack

Layer the stack:

1. **GPU telemetry and errors:** DCGM, which NVIDIA describes as managing and monitoring "datacenter GPUs and NVSwitches" ([DGX GB200 software](#)). Route XID events through the [ECC and XID triage decision tree](#).
2. **Fabric:** NMX-Telemetry from the switch trays, plus scheduled runs of the checks in the domain health table above.
3. **Rack and facility:** BMC Redfish events (leak, power shelf, thermal) feeding Mission Control or your own BMS bridge.
4. **Evidence for support cases:** NVDebug, which collects BMC out-of-band data and host logs for support teams ([NVDebug guide](#)).

Recommended alert severities:

- **Page immediately:** any partition UNHEALTHY, any leak event, a power shelf count approaching the four-healthy-shelf limit.
- **Ticket:** a single GPU NO_NVLINK, a link below 50 GB/s, or nvidia-imex version drift between trays.

Runbook outline and SLA checklist

Runbook skeleton for a rack:

1. Detect: classify the alert as GPU, tray, switch tray, cartridge, power or leak.
2. Contain: compute tray goes to an OFR partition. For a switch tray or cartridge, drain all NVLink jobs on the rack and delete partitions before service.
3. Evidence: NVDebug bundle, `nvsm show health`, fabric and partition output.
4. Escalate: open an NVIDIA or OEM case. Anything outside the three CRUs needs vendor-supplied parts.
5. Restore: replace, reinsert, re-add to the partition, run the MNNVL verification checks, return to the scheduler.
6. Account: log NVLink GPU-hours lost per incident, using the worked example above.

Questions to put in the support contract:

- Which parts can our staff replace, and may we see the OEM's Service Flow guide?
- What are time-to-part and time-to-repair targets for switch trays and cartridges, separately from compute trays?
- Can a compute tray be swapped on a live liquid loop? Get the answer in writing.
- How do the eight power shelves map to A/B facility feeds, and what happens on loss of one feed?
- Which NVLink software release (1.x or 2.x) and firmware recipe is supported, and does switch tray activation need a full-rack window?
- Is Mission Control licensed, and if not, who builds the leak-event path to the BMS?
- What acceptance tests (burn-in duration, NVLink error thresholds) must pass before handover?

GPU Smith can turn these questions into an acceptance and support specification.

Sources

1. [NVIDIA DGX GB200 user guide: Hardware](#)
2. [NVIDIA DGX GB200 user guide: Networking](#)
3. [NVIDIA DGX GB200 user guide: Software](#)
4. [NVIDIA DGX GB200 user guide: System Control and Configuration](#)
5. [NVIDIA DGX GB200 user guide: Rack Reboot Sequence](#)
6. [NVIDIA DGX GB200 user guide: System Health Check](#)
7. [NVIDIA DGX GB200 service manual: Introduction](#)
8. [NVIDIA DGX GB200 service manual: Power Supply Replacement](#)
9. [NVIDIA Multi-Node NVL Partition User's Guide, DU-12143-001 v1.0 \(September 18, 2026\); earlier GB200 NVL Partition User Guide v0.3 \(July 2025\)](#)
10. [NVIDIA Multi-Node NVLink user guide: Verifying](#)
11. [NVIDIA IMEX guide: Connections, Quorum, and Clean up](#)
12. [NVIDIA nvfwupd guide: Parallel update](#)
13. [NVIDIA NVDebug guide](#)
14. [NVIDIA Mission Control: NVLink info commands](#)
15. [NVIDIA Mission Control 2.3.1: Integration with Building Management Systems](#)
16. [NVIDIA Technical Blog: NVIDIA contributes GB200 NVL72 designs to OCP \(October 15, 2024\)](#)
17. [NVIDIA GB200 NVL72 product page](#)
18. [NVIDIA DGX B200 product page](#)
19. [NVIDIA DRA driver for GPUs: ComputeDomains](#)
20. [SemiAnalysis: H100 vs GB200 NVL72 Training Benchmarks \(August 20, 2025\)](#)

Source links

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

1. <https://www.nvidia.com/en-us/data-center/dgx-b200>
2. <https://www.nvidia.com/en-us/data-center/gb200-nvl72>

3. <https://docs.nvidia.com/dgx/dgxgb200-user-guide/hardware.html>
4. <https://developer.nvidia.com/blog/nvidia-contributes-nvidia-gb200-nvl72-designs-to-open-compute-project>
5. <https://docs.nvidia.com/multi-node-nvlink-systems/mnnvl-user-guide/verifying.html>
6. <https://docs.nvidia.com/dgx/dgxgb200-user-guide/networking.html>
7. <https://docscontent.nvidia.com/f4/26/62771ba2405d8f283c18194b46d7/partition-guide-v1-0.pdf>
8. <https://docs.nvidia.com/multi-node-nvlink-systems/partition-guide-v1-2.pdf>
9. <https://docs.nvidia.com/dgx/dgxgb200-service-manual/psu-replacement.html>
10. <https://docs.nvidia.com/mission-control/docs/nmc-software-installation-guide/2.3.1/integration-of-bms-with-bcm.html>
11. <https://docs.nvidia.com/base-command-manager/product-manuals/nvidia-mission-control/nvlinkinfo>
12. <https://docs.nvidia.com/dgx/dgxgb200-user-guide/health-check.html>
13. <https://docs.nvidia.com/dgx/dgxgb200-user-guide/software.html>
14. <https://docs.nvidia.com/multi-node-nvlink-systems/imex-guide/connections.html>
15. <https://dra-driver-nvidia-gpu.sigs.k8s.io/docs/concepts/compute-domains>
16. <https://docs.nvidia.com/dgx/dgxgb200-service-manual/introduction.html>
17. <https://newsletter.semianalysis.com/p/h100-vs-gb200-nvl72-training-benchmarks>
18. <https://docs.nvidia.com/dgx/dgxgb200-user-guide/control-config.html>
19. <https://docs.nvidia.com/multi-node-nvlink-systems/nvfupd-guide/appendix/full-rack-parallel-update.html>
20. <https://docs.nvidia.com/dgx/dgxgb200-user-guide/rack-reboot-sequence.html>
21. <https://docs.nvidia.com/multi-node-nvlink-systems/nvdebug-guide/index.html>

THE PUBLISHER

About GPU Smith

GPU Smith provides independent engineering and research for private AI infrastructure. We help technical buyers and operators reason about GPU workload sizing, hardware selection, procurement, deployment and inference operations, from the compute system to the networking, power and cooling requirements around it.

Start with the workload

Useful infrastructure decisions begin with the models, data, throughput, latency and operating constraints a team actually has. GPU Smith helps connect those requirements with system design and deployment choices. The goal is a privately controlled AI environment whose capacity and operational demands are understood before a purchasing decision.

Hardware and supplier research

Our public reference library covers GPUs, complete systems and networking components. The server-vendor directory and manufacturing research help teams investigate suppliers, compare options and follow sources behind technical claims. We distinguish manufacturer specifications from measured benchmarks and advertised capabilities from tested configurations. Reference pages are research resources, not a live inventory listing or a binding equipment quote.

Deployment and operations

GPU Smith's engineering scope includes infrastructure integration, networking, power, cooling and the practical operation of inference workloads. Our published articles explain the tradeoffs behind deployment, procurement and ongoing operations so teams can ask better questions and document decisions.

Work with GPU Smith

Explore the [hardware library](#), [GPU references](#), [networking references](#) and [vendor research](#). [Contact GPU Smith](#) with your workloads and deployment constraints to discuss a project and confirm current scope, pricing and availability.

Inclusion of a manufacturer or product in our research does not imply a partnership, certification or endorsement.

Website: <https://gpusmith.com>

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.