



# GPU Server PSU Redundancy: 4+2, 5+1, N+N, A/B Feeds

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 · dgx h100 power supply redundancy · dgx b200 5+1 psu · dgx b300 power supplies n+n · gpu server psu redundancy a/b feed · data center power supply architecture · dgx b300 power supply

Original article: <https://gpusmith.com/articles/en/gpu-server-psu-redundancy-and-rack-feed-design>



In this article

- 1. Why PSU topology decides the feed design
- 2. Notation: N+1, 4+2, 5+1, 3+3 and N+N
- 3. DGX H100: six PSUs, 4+2
- 4. DGX B200: six PSUs, 5+1, with a documented GPU derate
- 5. DGX B300: twelve PSUs, N+N
- 6. HGX OEM servers: 3+3 and 5+1 are configuration choices
- 7. Mapping PSUs to A/B feeds
- 8. Phase balance and breaker sizing
- 9. Failure behavior and power capping
- 10. Maintenance and hot-swap
- 11. Design checklist
- 12. Sources

Why PSU topology decides the feed design

A dual-corded server is usually assumed to survive the loss of either feed. That holds for a 1U server whose two supplies each carry the full load, not automatically for eight-GPU systems: a DGX H100 or B200 has six power supply units, and no vendor-documented scheme promises full operation on half of them. The design question is how many PSUs remain energized when one upstream path fails, and what the server does with that number.

This guide compares NVIDIA's documented schemes for three DGX generations plus two OEM HGX servers, then derives single-path failure outcomes. Vendor statements are quoted from documentation fetched on 2026-10-28; feed mappings and arithmetic are this guide's analysis. General rack sizing is covered in the GPU rack power requirements planning guide.

Notation: N+1, 4+2, 5+1, 3+3 and N+N

N is the count of supplies needed for the rated load; the second number is spares. Only 3+3 (generally N+N) lets an even A/B split survive a whole feed loss at full capability. Two other thresholds are often confused with N: the **boot minimum** (three PSUs in all three DGX guides, with only the BMC available below that) and the **shutdown threshold** for a running system. Between N and shutdown lies a degraded band where the system runs but limits power.

A consistency check, derived here rather than stated by NVIDIA, is that N is the smallest PSU count whose combined rating covers the system maximum:

| System        | PSUs installed | Per-PSU rating | System max. input | N × rating        | (N-1) × rating   |
|---------------|----------------|----------------|-------------------|-------------------|------------------|
| DGX H100/H200 | 6 (4+2)        | 3,300 W        | 10.2 kW           | 4 × 3.3 = 13.2 kW | 3 × 3.3 = 9.9 kW |

| System        | PSUs installed  | Per-PSU rating                       | System max. input | N × rating        | (N-1) × rating    |
|---------------|-----------------|--------------------------------------|-------------------|-------------------|-------------------|
| DGX B200      | 6 (5+1)         | 3,300 W                              | 14.3 kW           | 5 × 3.3 = 16.5 kW | 4 × 3.3 = 13.2 kW |
| DGX B300 (AC) | 12 (N+N, N = 6) | 3.2 kW (table) / 3.3 kW (power spec) | 15 kW system max. | 6 × 3.2 = 19.2 kW | 5 × 3.2 = 16.0 kW |

Table 1. PSU counts and ratings from the NVIDIA user guides; the last two columns are derived arithmetic.

For H100 and B200 the (N-1) column falls below the system maximum, consistent with the documented power limiting below N supplies. For B300, N = 6 leaves margin even though five supplies would nominally cover 15 kW.

## DGX H100: six PSUs, 4+2

The DGX H100/H200 user guide states that "the system includes six power supply units (PSU) configured for 4+2 redundancy," each rated "3300 W @ 200-240 V, 16 A, 50-60 Hz" ([NVIDIA DGX H100/H200 User Guide](#)). Its power table pairs the "200-240 volts AC" input with "10.2 kW max.", the same layout the B200 guide uses for its 14.3 kW system maximum, and NVIDIA's SuperPOD design guide sizes circuits on 10.2 kW per system at peak (Table 7: 20.4 kW per circuit for four systems) ([NVIDIA DGX SuperPOD Data Center Design, DGX H100: Electrical Specifications](#)).

The two NVIDIA documents conflict on failure behavior. The user guide says that if three PSUs lose power from "a data center issue or power distribution unit failure, the system continues to function, but at a reduced performance level." The SuperPOD guide says "Four of the six power supplies must be energized for the system to operate," and grades a traditional two-feed (2N) layout "Not acceptable for DGX H100 systems" because a feed loss drops a system "below four, resulting in a shutdown." Both pages were rechecked on 2026-10-28 and still disagree; design to the SuperPOD guide's four-PSU floor, the basis of NVIDIA's facility guidance.

## DGX B200: six PSUs, 5+1, with a documented GPU derate

The DGX B200 keeps six 3,300 W supplies at 200-240 V and 16 A each but moves to 5+1, with "14.3 kW max." in the input row ([NVIDIA DGX B200 User Guide](#)); the product page lists system power usage as "~14.3 kW max" ([NVIDIA DGX B200](#)). The B200 guide is the only one of the three that tabulates what happens as AC cables are pulled:

| AC cables removed | Operating PSUs | Effect (as documented)      | Max. power per GPU |
|-------------------|----------------|-----------------------------|--------------------|
| 1                 | 5              | No effect on the DGX system | 1,000 W            |
| 2                 | 4              | BMC alerts on AC power loss | 800 W              |
| 3                 | 3              | BMC alerts on AC power loss | 800 W              |
| 4                 | 2              | DGX system shuts down       | Not applicable     |

Table 2. DGX B200 behavior on AC cable removal, from the [NVIDIA DGX B200 User Guide](#); the guide notes the per-GPU maximum is reported by `nvidia-smi`.

Losing two or more supplies costs 200 W per GPU, 1.6 kW of GPU limit across eight GPUs (derived). The derate does not deepen between four and three PSUs, so a fourth surviving supply buys margin above shutdown, not performance. NVIDIA does not publish system wattage in the 800 W state, so the derate cannot be converted into a reduced feed load.

DGX B300: twelve PSUs, N+N

The AC version of the DGX B300 carries "12 power supply units (PSU) configured for N+N redundancy," with "the minimum number of required PSUs" stated as six and the same three-PSU boot minimum ([NVIDIA DGX B300 User Guide](#)). The guide lists power consumption as 14.5 kW and "15 kW system max." for both the 200-240 VAC and 54 VDC inputs. The supply rating appears as both "12 x 3.2 kW" and "3.3 kW @ 200-240 V, 16 A" (unchanged in the 2026-10-21 revision); this guide uses the lower 3.2 kW. The page gives no per-GPU limits by PSU count and does not mention dual-feed wiring; a 6/6 split across two feeds is an inference from the notation.

A busbar version is fed at 54 VDC from in-rack power shelves; NVIDIA's B300 SuperPOD reference architecture describes "six 5.5kW Power Shelves configured as N redundancy" delivering "up to 33kW," and up to four systems per rack ([NVIDIA DGX SuperPOD B300 Reference Architecture: Key Components](#)). Redundancy then moves to the shelves and their feeds, outside this AC-focused guide; the [B300 10 MW capacity worksheet](#) covers facility-scale B300 sizing.

HGX OEM servers: 3+3 and 5+1 are configuration choices

Supermicro's SYS-A21GE-NBRT (HGX B200 8-GPU) lists "6x 5250W Redundant (3 + 3) Titanium Level (96%) Hot-plug power supplies" at 200-240 VAC ([Supermicro SYS-A21GE-NBRT](#)); three surviving supplies give 3 × 5,250 W = 15.75 kW of nameplate (derived). Dell's XE9680 technical guide (Rev. A09, July 2025) offers six 2800 W or 3200 W supplies in either mode and spells out the difference: "Not Redundant (5+1) FR (System can get full performance)" versus "PSU Redundant (3+3) FTR (System throttles when more than one PSU is down)" ([Dell PowerEdge XE9680 Technical Guide](#)).

Mapping PSUs to A/B feeds

This table is analysis, assuming PSUs are split evenly across independent paths and one whole path is lost.

| System   | Layout             | PSUs left after one path loss | Outcome (derived from vendor text)                                                           |
|----------|--------------------|-------------------------------|----------------------------------------------------------------------------------------------|
| DGX H100 | A/B, 3+3           | 3                             | User guide: runs at reduced performance. SuperPOD guide: shuts down. Treat as not acceptable |
| DGX H100 | Three paths, 2+2+2 | 4                             | Meets the four-PSU operating floor; NVIDIA's documented design                               |

| System        | Layout                                    | PSUs left after one path loss | Outcome (derived from vendor text)                                                    |
|---------------|-------------------------------------------|-------------------------------|---------------------------------------------------------------------------------------|
| DGX B200      | A/B, 3+3                                  | 3                             | Runs, BMC alerts, GPUs capped at 800 W; one further PSU failure shuts the system down |
| DGX B200      | Three paths, 2+2+2                        | 4                             | Runs, GPUs capped at 800 W; tolerates one further PSU failure                         |
| DGX B200      | Any layout with two or more PSUs per path | $\leq 4$                      | 800 W per GPU; holding 1,000 W needs at most one PSU per path                         |
| DGX B300 (AC) | A/B, 6+6                                  | 6                             | Meets the stated six-PSU minimum; no documented derate                                |
| Dell XE9680   | A/B, 3+3 FTR                              | 3                             | Runs, throttled per Dell                                                              |

Table 3. Single-path failure outcomes (analysis).

The less obvious result is the B200: because four PSUs already trigger the 800 W cap, no practical multi-path layout preserves full GPU power through a path loss. NVIDIA's H100 guide adds an upstream caveat: in its plain N+1 pattern two of three rPDUs "are supplied by the same upstream UPS power source," so that UPS failing still powers the system off; only "Enhanced N+1" with three discrete UPS paths is graded "Optimal" (NVIDIA SuperPOD H100 electrical). Path independence must be traced to the UPS.

### Worked example 1: DGX H100, three paths (reproducing NVIDIA's table)

Assumptions: four systems per rack, 10.2 kW each at peak.

1. Rack peak:  $4 \times 10.2 = 40.8$  kW; normal load per path  $40.8 / 3 = 13.6$  kW.
2. One path lost:  $40.8 / 2 = 20.4$  kW per surviving path (the 50% rule).
3. Circuit at 415 V, 32 A, three-phase, 100% breaker rating, 0.95 power factor:  $\sqrt{3} \times 415 \times 32 \times 0.95 \approx 21.85$  kW (listed 21.8 kW).
4. Margin:  $21.8 - 20.4 = 1.4$  kW, NVIDIA's "stranded capacity".

### Worked example 2: two DGX B200 per rack on A/B (illustrative)

Two systems per rack follows NVIDIA's B200 SuperPOD reference architecture (NVIDIA DGX SuperPOD B200 Reference Architecture: Key Components); figures are illustrative.

1. Rack peak:  $2 \times 14.3 = 28.6$  kW; normal load 14.3 kW per feed,  $14.3 / 6 \approx 2.38$  kW per PSU.
2. Feed A lost: carrying 14.3 kW on three PSUs would need 4.77 kW each, above the 3,300 W rating, hence the derate.
3. Nameplate ceiling on feed B:  $6 \times 3.3 = 19.8$  kW (approximate; the guide does not say whether 3,300 W is input or output).
4. Conservative sizing uses the full 28.6 kW per feed, since the derated draw is unpublished. One 415 V/32 A circuit (21.8 kW) is insufficient; a 415 V/63 A circuit on a 100%-rated assembly gives  $\sqrt{3} \times 415 \times 63 \times 0.95 \approx 43.0$  kW, or each side can use two 32 A circuits.
5. Sizing to the 19.8 kW ceiling would rely on undocumented firmware limiting; confirm with NVIDIA first.

### Worked example 3: two DGX B300 (AC) per rack on A/B (illustrative)

1. Rack peak:  $2 \times 15 = 30$  kW; normal 15 kW per feed, 1.25 kW per PSU.
2. Feed lost: six PSUs per system carry 15 kW, 2.5 kW each, below the 3.2 kW rating.
3. Surviving feed: 30 kW, above 21.8 kW, so a larger circuit or two per side.
4. Per phase, 12 PSUs at four per phase:  $4 \times 2.5 = 10$  kW, or  $10,000 / (240 \text{ V} \times 0.95) \approx 43.9$  A line-to-neutral.

## Phase balance and breaker sizing

rPDUs "typically derive 200-240 VAC single phase power by dividing a three-phase input power circuit," and NVIDIA asks that "each PSU of each system should be connected to a different 'leg' (or phase)," noting that with the three-path N+1 pattern "it takes two racks of systems to balance the phases" ([NVIDIA DGX SuperPOD Data Center Design](#), [DGX H100: Electrical Specifications](#)). A 3+3 A/B layout maps one PSU per phase per feed. Check imbalance in the failover state, where 100% load can trip what 50% tolerated. DGX cords are C19/C20, 250 VAC, 20 A: six C19 outlets per H100 or B200 and twelve per AC B300 across both feeds.

In North America, NEC 210.20(A) requires the overcurrent device rating to be at least the noncontinuous load plus 125% of the continuous load, i.e. a standard breaker loaded to 80%, except for assemblies listed for 100% operation ([IAEI Magazine: 100% vs. 80%](#), written on the 2014 NEC; the 2023 and 2026 editions keep 210.20(A) with the same 125% rule and 100% exception per [UpCodes](#); use the edition your jurisdiction adopted). Schneider Electric notes that "UL has mandated a de-rating of all components to 80% continuous load" in rack PDUs ([Schneider Electric FAQ FA156194](#)). NVIDIA's Table 7 applies 100% to the 32 A IEC rows and 80% to the 60 A rows.

One row does not reproduce: for 208 V, 60 A, three-phase delta at 80%, NVIDIA lists 32.8 kW, but  $\sqrt{3} \times 208 \times 60 \times 0.8 \times 0.95 \approx 16.4$  kW, too little for the 20.4 kW per-circuit peak of four DGX H100 systems. The 32.8 kW figure is close to the 415 V/60 A row (32.7 kW); the row was unchanged on 2026-10-28, so recompute 208 V designs. Inrush is not specified in the three DGX guides; request it from the PSU or rPDU vendor before choosing trip curves.

## Failure behavior and power capping

On B200, losing two or more PSUs raises a BMC alert and the lower per-GPU maximum appears in `nvidia-smi` ([NVIDIA DGX B200 User Guide](#)); the H100 guide says only "reduced performance level," and the B300 guide gives no derate data. Alarm on PSU input loss from the BMC (Redfish or IPMI); one capped node can slow a synchronous multi-node job. The cost of an 800 W cap is workload-dependent; see [GPU power cap economics](#).

Commissioning should include NVIDIA's test of "turning off the individual circuit breakers feeding each rack power strip to verify that power redundancy has been achieved in each rack." On an A/B B200, opening one feed under load should yield a BMC alert and an 800 W limit, not a shutdown.

## Maintenance and hot-swap

All three DGX guides require replacing a failed PSU "immediately" and forbid operating "with PSUs depopulated." Live replacement is limited in degraded states: on H100, "if only three PSUs have power, shut down the system before replacing an operational PSU"; the B200 threshold is two. B300 supplies are "front-loaded with rear power connections," so swaps need front and rear access. In practice (analysis), never overlap feed maintenance and PSU swaps on a rack: an A/B B200 with one feed open is one supply from shutdown.

## Design checklist

| #  | Check                                                                                       |
|----|---------------------------------------------------------------------------------------------|
| 1  | PSU count, N, boot minimum and shutdown threshold per server model                          |
| 2  | PSUs energized after the worst single upstream failure, traced to the UPS                   |
| 3  | Documented behavior at that count (full power, derate, shutdown)                            |
| 4  | Per-feed load at failover: full system maximum unless the vendor publishes a derated figure |
| 5  | Circuit capacity with the correct derating (80% NEC or 100% listed) and power factor        |
| 6  | Per-phase current in the failover state                                                     |
| 7  | C19 outlet count per feed and cord type                                                     |
| 8  | BMC alarms for PSU input loss wired into monitoring                                         |
| 9  | Breaker-by-breaker failover test recorded at commissioning                                  |
| 10 | Spares and a rule that feed work and PSU swaps never overlap                                |

*Table 4. Per-rack PSU and feed worksheet.*

Upstream topology and UPS sizing are treated in [electrical design for GPU clusters](#). GPU Smith can review a proposed PSU-to-feed map on request.

## Sources

1. [Introduction to NVIDIA DGX H100/H200 Systems](#) - NVIDIA, 4+2 redundancy, 3300 W per PSU, "10.2 kW max.", boot minimum, replacement rules, cords (page last updated 2026-10-17; accessed 2026-10-28)
2. [Introduction to NVIDIA DGX B200 Systems](#) - NVIDIA, 5+1, 14.3 kW max., cable-removal table (1,000 W/800 W per GPU) (page last updated 2026-10-17; accessed 2026-10-28)
3. [Introduction to NVIDIA DGX B300 Systems](#) - NVIDIA, 12 PSUs N+N, six-PSU minimum, 3.2/3.3 kW, 15 kW max. (page last updated 2026-10-21; accessed 2026-10-28)
4. [NVIDIA DGX B200](#) - NVIDIA, system power usage "~14.3 kW max" (accessed 2026-10-28)

5. [NVIDIA DGX SuperPOD Data Center Design, DGX H100: Electrical Specifications](#) - NVIDIA, four-PSU requirement, feed-pattern grading, Table 7, phase balancing, breaker test (page last updated 2025-11-19; accessed 2026-10-28)
6. [NVIDIA DGX SuperPOD B200 Reference Architecture: Key Components](#) - NVIDIA, two DGX B200 systems per rack (accessed 2026-10-28)
7. [NVIDIA DGX SuperPOD B300 Reference Architecture: Key Components](#) - NVIDIA, 5.5 kW power shelves in N redundancy, 33 kW, four shelves and up to four systems per rack (accessed 2026-10-28)
8. [Supermicro SYS-A21GE-NBRT](#) - Supermicro, HGX B200 server with 6x 5250W 3+3 redundant supplies at 200-240 VAC (accessed 2026-10-28)
9. [Dell PowerEdge XE9680 Technical Guide](#) - Dell, Rev. A09 July 2025, 5+1 FR vs 3+3 FTR (accessed 2026-10-28)
10. [50A 208v 3 Phase Rack PDU Load Ratings and Calculations Explained \(FA156194\)](#) - Schneider Electric, UL 80% rack PDU derating (accessed 2026-10-28)
11. [100% vs. 80%: Choosing the Right OCPD Solution](#) - IAEI Magazine, NEC 210.20(A) 125% rule and 100%-rated exception (2014 NEC basis) (accessed 2026-10-28)

## Source links

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

1. <https://docs.nvidia.com/dgx/dgxb200-user-guide/introduction-to-dgxb200.html>
2. <https://docs.nvidia.com/dgx-superpod/design-guides/dgx-superpod-data-center-design-h100/latest/electrical.html>
3. <https://docs.nvidia.com/dgx/dgxb200-user-guide/introduction-to-dgxb200.html>
4. <https://www.nvidia.com/en-us/data-center/dgx-b200>
5. <https://docs.nvidia.com/dgx/dgxb300-user-guide/introduction-to-dgxb300.html>
6. <https://docs.nvidia.com/dgx-superpod/reference-architecture/scalable-infrastructure-b300/latest/dgx-superpod-components.html>
7. <https://www.supermicro.com/en/products/system/gpu/10u/sys-a21ge-nbrt>
8. <https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/poweredge-xe9680-technical-guide.pdf>
9. <https://docs.nvidia.com/dgx-superpod/reference-architecture/scalable-infrastructure-b200/latest/dgx-superpod-components.html>
10. <https://iaeimagazine.org/magazine/2016/may2016/100-vs-80-choosing-the-right-ocpd-solution>
11. <https://up.codes/s/continuous-and-noncontinuous-loads>
12. <https://www.se.com/ca/en/faqs/FA156194>

## 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.