



B300 Cluster Power Requirements: Can 10 MW Fit 600 Nodes?

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 · 2026-10-15 · b300 cluster power requirements · b300 power consumption · 600 b300 nodes · 10 mw data center · dgx b300 · gpu financing · critical it capacity · pue · gpu due diligence

Original article: <https://gpusmith.com/articles/en/b300-nodes-10mw-power-capacity-worksheet>



In this article

1. Executive Summary
 2. Introduction and Background
 3. Key Changes: From Equipment Counts to Capacity Evidence
 4. Measurement Boundaries: MW, MWh and PUE
 5. Implementation Considerations and Process Changes
 6. Maintenance, Loss of Feed and Cooling-Limited Capacity
 7. Data Analysis and Evidence
 8. Implications and Future Directions
 9. Frequently Asked Questions (FAQs)
 10. Conclusion
 11. References
-

Executive Summary

B300 cluster power requirements depend on the electrical boundary, the exact server, and the operating state. A nominal **10 megawatt (MW)** facility cannot be reconciled to **600 servers with eight graphics processing units (GPUs) each** by dividing utility capacity by a GPU chip rating. NVIDIA documents eight GPUs per DGX B300 system, so the proposed inventory represents **4,800 GPUs**, calculated as servers multiplied by GPUs per server. Its guide lists **14.5 kilowatts (kW) power consumption** in the physical specifications and a separate **15 kW system maximum** in the power specifications. These are documented ratings, not measurements of the proposed installation. (docs.nvidia.com) (docs.nvidia.com)

Illustrative arithmetic: at an assumed **power usage effectiveness (PUE) of 1.20**, a 10,000 kW facility-meter envelope translates to **8,333 kW of information technology (IT) load**, before any separately retained capacity reserve. The 14.5 kW DGX reference requires **8,700 kW for 600 servers alone** and **10,440 kW at the facility boundary**, before additional network, storage and management IT. That scenario therefore exceeds the facility envelope. The illustrative 15 kW input sensitivity increases server-only facility demand to **10,800 kW**. These calculations use the published DGX inputs and an illustrative overhead ratio; they do not establish that every HGX B300 configuration fails. (docs.nvidia.com) (www.energy.gov)

If the same **10 MW means usable critical IT**, the 14.5 kW server case instead leaves **1,300 kW** before supporting IT and reserves. The distinction matters because PUE compares facility energy with IT energy, and the IT boundary includes more than compute servers. NVIDIA's reference architecture separately identifies compute, storage and management infrastructure; the federal metering guide also separates IT equipment from facility-support energy. A measured annual PUE does not by itself certify peak electrical or cooling capacity. (www.thegreengrid.org) (docs.nvidia.com) (datacenters.lbl.gov) (www.energy.gov)

The lender's decision should therefore follow an evidence chain: **contracted and energized utility MW → usable critical IT kW → support loads and reserve → exact operating server count → surviving capacity in required maintenance and fault states**. Dominion describes staged service and bridging power for commissioning; those states must be distinguished from the final supply envelope. Lambda's public

financing announcement illustrates the relevance of commissioning-linked draws, without disclosing a reusable private acceptance standard. Capacity that remains unresolved should become an explicit funding condition, with an owner, required evidence and release threshold. (www.dominionenergy.com) (lambda.ai)

Introduction and Background

The practical question is whether financed GPU equipment can operate inside the power capacity actually available when funds are released. A proposed server count, a utility letter and an equipment schedule describe different parts of that question. The reconciliation must connect them through a common boundary and a declared operating state. Dominion Energy's description of staged data-center service provides a concrete reason to distinguish initial energization from the final power plan. (www.dominionenergy.com)

This report addresses **bank and private-credit credit officers, equipment lessors, infrastructure funds and their electrical reviewers**. It treats GPU Smith as an adjacent independent engineering advisor, rather than a server manufacturer, utility or cloud provider. The worksheet is a proposed review method. It does not claim access to a borrower configuration, utility agreement, commissioned site or private loan agreement.

The timing is relevant to equipment financing. In its announcement dated **October 1, 2026**, Lambda reported a **\$1.008 billion delayed-draw term loan** and said the structure aligns proceeds with cluster commissioning milestones. This is evidence of one company's publicly described financing structure, not evidence that all lenders require the same tests or that its terms apply to another transaction. (lambda.ai) (lambda.ai)

The analysis is anchored to sources retrieved **October 15, 2026 UTC**. Observations below are statements found in those sources. **Illustrative arithmetic** is reproducible calculation using declared inputs.

Assumptions are scenario values without project evidence. **Hypotheses** identify propositions that require testing before they can support a capacity conclusion. This separation is particularly important where a power-supply rating differs from operating demand, a distinction explicitly recognized by the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE). (handbook.ashrae.org)

The intended output is a capacity schedule that a sponsor, lender and independent engineer can reconcile. Broader facility planning is covered in the related [GPU rack power planning guide](#); the calculations here focus on which financed inventory fits the evidenced operating envelope.

Key Changes: From Equipment Counts to Capacity Evidence

Change the question from advertised MW to available MW

A capacity statement should identify the meter, delivery point, service stage, energization status and operating restrictions. Dominion states that service can arrive in stages and describes a bridging solution for commissioning and initial load while the final power plan is constructed. Consequently, evidence that a site can begin commissioning should not be treated as evidence that its eventual full operating allocation is already available. (www.dominionenergy.com)

Utility terms also need their own review. Duke Energy's **July 23, 2026** announcement describes agreements that can include [temporary curtailment for limited grid events](#). This is a utility-specific example of why an apparently sufficient normal supply figure may need a separate curtailed-operation case. It establishes no

curtailment obligation for the hypothetical facility considered here. (investors.duke-energy.com)

The proposed capacity schedule should leave these project inputs blank until supplied:

- **Delivery boundary:** Utility revenue meter, landlord allocation, hall intake or critical IT output: _____.
- **Available service:** Contracted MW, currently energized MW and effective dates: _____.
- **Restrictions:** Curtailment conditions, ramp restrictions and shared-site allocations: _____.
- **Operating commitment:** Inventory and workload required during each contractual state: _____.
- **Evidence owner:** Utility, landlord, sponsor or design authority assigned to each entry: _____.

These blanks are deliberate. AEP Ohio's public process distinguishes planning from energization after infrastructure is in place. Its detailed large-load procedures have their own applicability thresholds and should not be imported as the governing tariff for an unidentified site. The review needs the actual serving utility's agreement. (www.aepohio.com)

Replace the GPU label with the complete configured system

HGX B300, DGX B300 and GB300 NVL72 are different equipment boundaries. The worksheet requires an **original equipment manufacturer (OEM)** configuration. Dell identifies an eight-GPU HGX B300 configuration and separately lists its server's power supplies. GIGABYTE identifies GB300 NVL72 rack configurations and separate HGX offerings. None of those product descriptions authorizes substitution of a rack rating for a server rating or of power-supply unit (PSU) capacity for measured draw. (www.dell.com) (www.gigabyte.com)

Table 1 records the public examples used in this report. All were accessed October 15, 2026; values remain specific to the cited configuration and document boundary.

Source and configuration	Published power boundary	Maximum versus observed	Worksheet treatment
NVIDIA DGX B300, introduction guide	14.5 kW physical-table consumption; separate 15 kW system maximum. (docs.nvidia.com)	Documented specifications; no project workload observation.	Retain both illustrative inputs pending written OEM resolution of the sizing basis, including peak demand and AC/DC boundaries.
NVIDIA DGX B200, introduction guide	14.3 kW maximum system input. (docs.nvidia.com)	Maximum for the documented DGX system.	Comparator only; never substitute it for a B300 purchase configuration.
GIGABYTE GB300 NVL72 PAC2	140 kW per rack in its configuration table. (www.gigabyte.com)	OEM rack specification, not an eight-GPU node measurement.	Separate rack-scale inventory and management-rack boundary.
Dell PowerEdge XE9780	Twelve 3,200 W PSU ratings listed separately from GPU configuration. (www.dell.com)	PSU capacity, not a published measured whole-server load.	Actual configured server input: _____ kW.
Lenovo ThinkSystem SR680a V4	Configuration consumption calculated through Lenovo Capacity Planner. (lenovopress.lenovo.com)	Configuration-dependent calculation.	Obtain the saved configuration, voltage and sizing output.

The useful comparison is therefore **whole configured input against whole configured input**. ASHRAE states that nameplate rating does not represent actual usage draw. Summing redundant PSU outputs can substantially misstate the load to be served; conversely, a favorable workload observation does not demonstrate the maximum envelope required by the financing assumptions. (handbook.ashrae.org)

Expand server power into whole-cluster power

The **600-node inventory** in the illustrative case contains eight GPUs per server, hence **4,800 GPUs**. It is neither a 600-GPU purchase nor a count of NVL72 racks. This arithmetic follows the DGX server's published GPU count; it should be repeated against the actual OEM bill of materials before use in a transaction. (docs.nvidia.com)

Supporting IT requires a separate inventory. NVIDIA's reference design separates storage and compute fabrics and includes distinct storage systems. Supermicro's published B300 cluster example likewise describes **eight compute racks, three networking racks and two coolant distribution units (CDUs)**. That example demonstrates supporting scope; it is not a bill of materials for this proposed cluster. (docs.nvidia.com) (www.supermicro.com)

Measurement Boundaries: MW, MWh and PUE

Capacity and annual energy answer different questions

The U.S. Energy Information Administration (EIA) defines **one megawatt as 1,000 kilowatts** and distinguishes momentary power from electricity measured over time. Accordingly, the supplied 10 MW envelope is **10,000 kW**, a power quantity. Megawatt-hours (MWh) describe energy accumulated over a period; an annual energy model cannot establish whether the site carries its peak concurrent load. (www.eia.gov) (www.eia.gov)

The Department of Energy (DOE) defines PUE using **annual facility energy divided by annual IT energy**. The Green Grid describes the same facility-to-IT energy boundary. A sponsor can use that ratio for an energy model, but converting a capacity limit into IT power requires an overhead assumption applicable to the operating point being assessed. That conversion is an analytical assumption until supported by a load- and weather-matched engineering model or measurement. (www.energy.gov) (www.thegreengrid.org)

The public description of **ISO/IEC 30134-2:2026**, from the International Organization for Standardization and the International Electrotechnical Commission, includes PUE measurement categories and rules for unmeasured energy. The standard's existence should not be used to imply that a brochure PUE, a utility-meter observation and an uninterruptible power supply (UPS) output measurement share identical boundaries. Record the measurement category, included equipment, interval, load and environmental conditions. (www.iso.org) (www.iso.org)

Apply overhead once, with a traceable denominator

For the **illustrative PUE 1.20 case**, the calculation is $10,000 \div 1.20 = 8,333.33 \text{ kW IT}$. The implied facility-support allowance is **1,666.67 kW**, calculated as facility power less IT power. It is already inside that conversion. Deducting a second generic cooling allowance for the same equipment would count overhead

twice. This example applies the published PUE boundary, not a measured site efficiency. (www.energy.gov) (datacenters.lbl.gov)

An equally valid worksheet may itemize distribution losses, pumps, fans, heat rejection and other facility loads instead of applying a single overhead ratio. Its boundary must reconcile to the same facility total. The DOE metering guide warns that basic UPS-output measurement can count downstream distribution losses as IT load. Therefore “critical IT at UPS output” and “server-input IT” may require an explicit downstream-loss bridge. (datacenters.lbl.gov)

Use the following measurement controls:

- **Meter map:** Identify each electrical meter and its position on the single-line diagram.
- **Common interval:** Compare facility, IT and support readings over aligned timestamps.
- **Common inventory:** Match powered equipment to the configuration and server count under review.
- **Thermal record:** Capture relevant temperatures, flows and heat-removal conditions alongside electrical readings.
- **Boundary reconciliation:** Explain losses and shared loads between upstream and downstream meters.

These are proposed controls. ENERGY STAR describes intelligent power distribution units (PDUs) reporting PDU and outlet currents, while the DOE metering guide recommends measuring thermal load where possible. Those capabilities support reconciliation; the presence of a meter alone does not establish calibration, completeness or passing capacity. (www.energystar.gov) (datacenters.lbl.gov)

The related [GPU cluster energy calculator guide](#) addresses the separate annual-energy workflow. The [GPU power calculator](#) can help organize an initial scenario; retain source inputs and calculation boundaries in the lender's own evidence schedule.

Implementation Considerations and Process Changes

Build the meter-to-node waterfall

The proposed **B300 data center power capacity worksheet** starts at the declared delivery point and moves downstream. It should record a value, units, source document, revision, date, operating state and responsible party at each step. The Whole Building Design Guide (WBDG) commissioning guidance places expectations, benchmarks and success criteria in the owner's project requirements (OPR), providing a useful documentation framework for those entries. (legacy.wbdg.org)

The normal-state waterfall is:

- **Contracted and energized facility capacity:** _____ kW at the identified meter.
- **Facility-support conversion:** Applicable modeled overhead or itemized non-IT loads: _____.
- **Usable critical IT capacity:** _____ kW after the documented electrical boundary reconciliation.
- **Supporting IT:** Compute fabric, storage fabric, storage, management and other IT: _____ kW.
- **Retained capacity reserve:** _____ kW, with rationale and a boundary that prevents duplication.
- **Server budget:** Critical IT capacity less supporting IT and retained reserve: _____ kW.

- **Configured server input:** _____ kW per node at the required operating envelope.
- **Electrical node ceiling:** Round down server budget divided by configured node input.

The support-load line cannot be inferred from server count alone. NVIDIA assigns the storage topology, capacity and components to the storage partner, while Lenovo describes configuration-specific consumption calculations. Collect both the supporting bill of materials and the configured node sizing output rather than applying an unexplained percentage to the accelerator inventory. (docs.nvidia.com) (lenovopress.lenovo.com)

The reserve is an explicit scenario choice, not an asserted industry requirement. It can represent unallocated operating headroom only if its purpose is documented. A blanket reserve should not duplicate a derating already applied to usable path capacity, an overhead allowance already incorporated through PUE, or a spare module already excluded when calculating the surviving operating state.

Reconcile electrical ratings locally

The electrical reviewer should confirm applicable continuous-load rules, equipment listings, voltage, ambient conditions, protective settings and jurisdiction. UL Solutions explains that, unless marked for full continuous operation, a breaker is intended for no more than **80% of rated current** when normal operation continues for at least **three hours**. It also makes **100% continuous ratings conditional on the specified enclosure or cubicle**. These are conditional product rules, not authorization to multiply every facility rating by the same percentage. (www.ul.com) (www.ul.com)

Voltage matters at the equipment boundary too. Lenovo's guide limits a particular **3,200 W PSU to 2,700 W output at 200 V alternating current (AC)**. This specific example shows why an apparently redundant PSU arrangement needs its voltage conditions checked. It does not establish the derating of another server, rack PDU or upstream circuit. (lenovopress.lenovo.com)

For rack density, a **Hypothetical Example** placing **four 14.5 kW systems** in a rack produces **58 kW server load**; using the separate illustrative 15 kW input produces **60 kW**. These are multiplication exercises based on DGX specifications, not approved rack layouts. Outlet, branch, busway, feed, heat-removal and physical constraints still require evidence. The related [data-center electrical design guide](https://docs.nvidia.com) covers the broader design workstream. (docs.nvidia.com) (docs.nvidia.com) (www.energystar.gov)

Keep supporting equipment in the correct category

An external CDU, a storage array and a network switch should not disappear into a general “infrastructure” line with no meter boundary. The DOE metering guide distinguishes IT equipment energy from facility-support energy. Classify each component according to the adopted measurement method, then reconcile it once in the waterfall. Shared plant allocations also need a stated basis so the proposed cluster is not credited with capacity assigned elsewhere. (datacenters.lbl.gov)

For liquid cooling, ASHRAE identifies **configuration-specific flow and pressure-differential requirements** and requires consideration of CDU approach temperature. A cooling nameplate without its water temperatures, flow, pressure and applicable equipment configuration does not establish supported IT load. The related [GPU data-center cooling design guide](https://handbook.ashrae.org) provides adjacent coverage. (handbook.ashrae.org) (handbook.ashrae.org)

Maintenance, Loss of Feed and Cooling-Limited Capacity

Calculate the surviving path rather than a redundancy multiplier

Schneider Electric defines **N** as the power capacity required to feed protected equipment and describes a system-plus-system design as **2N**. Eaton defines **N+1** UPS redundancy as at least one module beyond those needed to support the load. These labels describe arrangements and redundancy objectives; they do not make the compute workload consume twice its input power. (productinfo.se.com) (productinfo.se.com) (www.eaton.com)

A **Hypothetical Example** illustrates the distinction. Suppose two independent paths each have **10,000 kW usable critical IT capacity** and normally share a **9,200 kW load** equally. Each carries **4,600 kW** normally; after losing one path, the survivor must carry **9,200 kW**. The relevant comparison is surviving usable capacity against transferred load, including every downstream bottleneck. The topology label alone establishes none of these assumed ratings. (productinfo.se.com)

Installed spare capacity also needs maintenance-state treatment. If a module or feed is already out for maintenance, the surviving envelope can differ from both the normal state and a single-fault state starting with everything available. Eaton's redundancy definition supplies the counting concept; the actual module capacities, permitted combinations and transfer behavior must come from the site's design and tests. (www.eaton.com)

Review the required states separately:

- **Normal operation:** All equipment intended to be available is operating, with the declared utility allocation.
- **Planned maintenance:** The identified module, path or cooling component is isolated as the maintenance procedure requires.
- **Loss of feed:** The defined electrical source or path is unavailable and loads transfer as designed.
- **Cooling failure:** The defined thermal component or path is unavailable, with electrical power potentially still present.

Do not infer a Tier classification from an N+1 label. Uptime Institute describes its classification as outcome based, and Tier III includes maintenance without shutdowns and redundant delivery paths for power and cooling. A server PSU configuration alone cannot establish those facility outcomes. (journal.uptimeinstitute.com) (journal.uptimeinstitute.com)

Treat thermal survival as its own constraint

Available electrical kW do not establish available thermal kW. The cooling-limited node ceiling should use the supported heat-removal capacity at the relevant supply temperatures and flow conditions. ASHRAE's configuration-specific requirements and CDU approach-temperature guidance make those boundary conditions material. DOE also identifies outdoor temperature and humidity as influences on cooling opportunities. (handbook.ashrae.org) (handbook.ashrae.org) (www.energy.gov)

The proposed underwriting rule is to use the lowest evidenced node ceiling across the states that the operating and funding assumptions require. This is a review method, not a universal loan covenant. A **hypothesis** that **power capping** preserves the financed service must be tested against both electrical demand and the required workload outcome; a lower power reading alone does not establish equivalent delivered service.

Commissioning must connect components to system behavior. Vertiv describes integrated testing under designed load and emergency conditions. WBDG identifies test reports as records of testing and field observation. Request the corresponding raw traces, meter map, inventory and pass criteria, rather than accepting an unqualified completion certificate as a numerical post-fault capacity value. (www.vertiv.com) (legacy.wbdg.org)

Data Analysis and Evidence

Work the 600-node case explicitly

All values in this subsection and Table 2 are **illustrative arithmetic**, except the cited DGX input specifications. They are not observations of a real facility. At **600 × 14.5 kW**, servers require **8,700 kW**. Applying assumed **PUE 1.20** gives **10,440 kW** at the facility boundary, a **440 kW shortfall** against 10,000 kW before additional IT. At **15 kW per node**, the corresponding values are **9,000 kW IT** and **10,800 kW facility**, a **800 kW shortfall**. (docs.nvidia.com) (www.energy.gov)

The assumptions for the complete illustrative cluster are **500 kW supporting IT** and **400 kW retained IT reserve**. Neither is sourced as a typical or recommended allowance. The meter case then leaves **7,433.33 kW** for servers and supports **512 whole nodes** at 14.5 kW. The same assumed support and reserve applied to **10,000 kW usable critical IT** leave **9,100 kW** and support **627 nodes**. Round down, because a fractional node cannot satisfy the inventory schedule. (docs.nvidia.com)

Table 2 shows the sensitivities. Maintenance, feed-loss and cooling figures are assumed usable IT ceilings after their respective boundary and derating calculations. They are not derived from the normal-state PUE. The table applies the DGX reference power solely to expose the decision's sensitivity; it does not demonstrate a peak operating envelope. (docs.nvidia.com) (docs.nvidia.com)

Illustrative boundary or state	Usable IT kW	Support / reserve kW	Server kW per node	Whole-node ceiling
10 MW facility, PUE 1.20, servers only	8,333.33	0 / 0	14.5	574
10 MW facility, PUE 1.20, complete scenario	8,333.33	500 / 400	14.5	512
10 MW facility, PUE 1.10 sensitivity	9,090.91	500 / 400	14.5	564
10 MW facility, PUE 1.30 sensitivity	7,692.31	500 / 400	14.5	468
10 MW usable critical IT, normal	10,000.00	500 / 400	14.5	627
Maintenance, assumed surviving IT ceiling	8,000.00	500 / 400	14.5	489
Loss of feed, assumed surviving IT ceiling	7,000.00	500 / 400	14.5	420

Illustrative boundary or state	Usable IT kW	Support / reserve kW	Server kW per node	Whole-node ceiling
Cooling failure, assumed supported IT ceiling	6,500.00	500 / 400	14.5	386
10 MW facility, PUE 1.20, illustrative 15 kW input sensitivity	8,333.33	500 / 400	15.0	495

The table's normal critical-IT case fits the proposed inventory, but its assumed maintenance and fault cases do not. No actual site's answer follows from those invented capacity ceilings. Their purpose is to show why the approved node count must state the required operating condition. The governing method remains surviving capacity and collective system behavior. (productinfo.se.com) (www.vertiv.com)

For a different OEM system to fit **600 nodes** inside the meter case with the assumed support and reserve, its allowable configured node input would be at most $(10,000 \div 1.20 - 500 - 400) \div 600 = 12.389 \text{ kW}$. Without either additional allowance, the server-only threshold is **13.889 kW**. These are calculated ceilings, not assertions about any OEM's achievable consumption. Configuration-specific sizing is therefore essential. (lenovopress.lenovo.com)

Reconcile documents to claims before release

Table 3 converts the calculation into a proposed evidence schedule. Project values remain blank until supplied. The structure follows WBDG's distinction between requirements and test records; no public source establishes that a particular lender must use this exact schedule. (legacy.wbdg.org) (legacy.wbdg.org)

Document or record	Claim it must substantiate	Project value	Suggested unresolved-item treatment
Utility agreement, allocation and energization evidence	Meter boundary, currently available service, ramp and curtailment conditions	_____ MW; effective _____	Utility or sponsor closes capacity conditions before the relevant draw.
As-built single-line and rating schedule	Usable path limits and surviving maintenance/feed-loss capacity	Normal _____ kW; surviving _____ kW	Electrical design authority reconciles ratings, settings and applicable rules.
OEM configuration and input-power basis	Exact server count, GPU count, voltage and power envelope	Stock-keeping unit (SKU) _____; node _____ kW	Supplier resolves rating ambiguity before inventory capacity approval.
Fabric, storage and management bill of materials	Supporting IT within the adopted measurement boundary	_____ kW	Sponsor supplies component-level schedule and meter mapping.
Cooling curves and equipment schedule	Supported IT heat load at required weather, temperature and failure conditions	Normal _____ kW; surviving _____ kW	Cooling reviewer closes the thermal capacity bridge.
Synchronized facility and rack metering	Measured demand reconciles across upstream and downstream boundaries	Interval _____; residual _____	Commissioning authority explains unallocated losses and shared loads.

Document or record	Claim it must substantiate	Project value	Suggested unresolved-item treatment
Integrated test report and raw traces	Required workload survives defined maintenance and fault states	State _____; nodes _____	Relevant funding condition remains open until evidence meets the declared criterion.

Each document answers a different question. A utility allocation does not prove downstream delivery, an OEM sizing tool does not prove cooling, and a test report needs stated conditions to support a reusable capacity conclusion. WBDG distinguishes commissioning review from technical peer review, so the engagement scope should identify which work was actually performed. (legacy.wbdg.org)

Implications and Future Directions

For **GPU infrastructure loan due diligence**, the practical improvement is an explicit reconciliation between equipment eligible for funding and capacity eligible for operation. Lambda's commissioning-linked financing announcement supports the relevance of milestone evidence, while leaving its detailed private acceptance requirements undisclosed. The proposed worksheet should therefore be presented as an independent review framework, not as a reconstruction of that transaction. (lambda.ai)

Suggested conditions should identify measurable deliverables rather than broad promises:

- **Capacity condition:** Executed evidence of the required energized allocation at the named boundary.
- **Configuration condition:** Frozen OEM inventory and an accepted maximum-power basis.
- **Infrastructure condition:** Complete supporting IT and facility-load reconciliation.
- **Resilience condition:** Required node count demonstrated in the agreed maintenance and fault states.
- **Measurement condition:** Raw synchronized electrical and thermal records with explained residuals.
- **Closeout condition:** Responsible owners, approved exceptions and a controlled as-built evidence package.

Seasonal limitations deserve explicit disposition. WBDG recommends multi-season testing, and Vertiv notes that construction testing cannot simulate every environmental condition. Where the initial test cannot reproduce the required weather, record what remains modeled, what operating restriction applies, and what later evidence must close the item. Do not label an untested seasonal envelope as an observed result. (legacy.wbdg.org) (www.vertiv.com)

The distinction also matters for future expansion: power availability can change in stages, and equipment configurations can change independently of utility service. A later capacity increase should trigger a revised reconciliation rather than automatic approval of additional nodes. Dominion's staged-service description and Lenovo's configuration-dependent sizing explain why those two evidence chains must remain linked. (www.dominionenergy.com) (lenovopress.lenovo.com)

This report does not reproduce interconnection procedures or rack PSU failure-test protocols. The companion topics **GPU Data Center Power Access** and **GPU Rack Power Failure Matrix** can be cross-referenced when published. For the wider acceptance and responsibility framework, see the related [NVIDIA DSX due-diligence guide](#).

Limitations

No borrower documents or site measurements were supplied. Published specifications are observations about the cited documentation; all project capacities, support allowances, reserves and node-count sensitivities remain illustrative. PUE assumptions are not warranties. Applicable electrical rules and thermal limits require project-specific professional assessment against the actual installation. UL's conditional breaker guidance and ASHRAE's configuration-specific cooling requirements explain why a generic percentage cannot close those reviews. (www.ul.com) (handbook.ashrae.org)

The public ISO description was used for its stated metric scope, not as evidence of access to every provision of the complete standard. No benchmark, private loan term, customer experience, conversion result or market-demand estimate is inferred from these sources. The conclusions concern capacity arithmetic and evidence boundaries.

Frequently Asked Questions (FAQs)

Can 10 MW support 600 B300 nodes?

For the illustrative DGX reference at 14.5 kW and PUE 1.20, 10 MW at the facility meter is insufficient: server-only demand is **10.44 MW**. If 10 MW instead means usable critical IT, those servers leave **1.3 MW** before other IT and reserve. Approval should require the exact OEM configuration, resolution of peak demand and AC/DC boundaries, and evidence for the required post-fault states. (docs.nvidia.com) (www.energy.gov)

Is B300 node power consumption the GPU power multiplied by GPU count?

That calculation omits the rest of the server and its conversion and cooling loads. Use an exact whole-system power basis. Dell publishes GPU and PSU specifications separately; ASHRAE explains why PSU nameplates do not equal operating draw. Neither summed GPU ratings nor summed redundant PSU ratings establish configured server consumption. (www.dell.com) (handbook.ashrae.org)

Does 2N require twice the utility load?

2N describes duplicated capacity, not a fixed multiplier on useful IT consumption. Calculate normal loading, losses and surviving-path capacity from the installed arrangement. Schneider's terminology defines N by the protected load and distinguishes a system-plus-system design. Both paths' actual losses still belong inside the facility boundary. (productinfo.se.com)

Conclusion

The answer to B300 cluster power requirements begins with the location and availability of the capacity figure. A lender should first establish whether the proposed envelope describes utility-meter power or usable critical IT, then reconcile the exact inventory through supporting loads, overhead and retained reserve. The distinction follows directly from the published PUE boundary and metering guidance. (www.thegreengrid.org) (datacenters.lbl.gov)

The **600-node DGX illustration** does not fit a **10 MW facility envelope at assumed PUE 1.20**, even before additional IT. It can fit the same numerical **critical-IT envelope** under the stated normal assumptions, while still requiring separate maintenance and fault-state evidence. The documented 14.5 kW consumption and separate 15 kW user-guide maximum must remain visible as illustrative inputs until the equipment supplier resolves the applicable sizing basis, including peak demand and AC/DC boundaries. (docs.nvidia.com) (docs.nvidia.com)

The proposed credit decision is therefore conditional and reproducible: fund an inventory whose configured demand fits the contracted, commissioned and required surviving envelope, and identify every unresolved input before treating it as capacity. Integrated testing and traceable commissioning records provide the evidence bridge between a design calculation and an operating conclusion. (www.vertiv.com) (legacy.wbdg.org)

For an independent pre-close review of that capacity reconciliation, use GPU Smith's GPU financing due-diligence contact.

References

All cited pages were retrieved October 15, 2026 UTC. Citations point to independent source organizations; publisher links above are navigation or a call to action. Primary documentation establishes specifications and methods, while the worksheet calculations remain explicitly illustrative.

- **System specifications:** NVIDIA DGX B300 and DGX B200 user guides, GIGABYTE AI Factory configurations, Dell PowerEdge XE9780 specifications, Lenovo SR680a V4 guide and Supermicro B300 cluster datasheet, cited in Table 1 and the surrounding discussion. (docs.nvidia.com) (docs.nvidia.com) (www.gigabyte.com) (www.dell.com) (lenovopress.lenovo.com) (www.supermicro.com)
- **Cluster architecture:** NVIDIA DGX SuperPOD B300 components reference, cited in the supporting-IT discussion. (docs.nvidia.com)
- **Measurement:** EIA electricity units, DOE cooling guidance, the DOE/LBNL metering guide, ENERGY STAR PDU guidance, ISO's PUE standard description and The Green Grid PUE definition, cited in Measurement Boundaries. (www.eia.gov) (www.energy.gov) (datacenters.lbl.gov) (www.energystar.gov) (www.iso.org) (www.thegreengrid.org)
- **Capacity and commissioning:** Schneider Electric, Eaton, Uptime Institute, ASHRAE, UL Solutions, Vertiv and WBDG guidance, cited in the electrical, thermal and evidence sections. (productinfo.se.com) (www.eaton.com) (journal.uptimeinstitute.com) (handbook.ashrae.org) (www.ul.com) (www.vertiv.com) (legacy.wbdg.org)
- **Utility and financing examples:** Dominion Energy, AEP Ohio, Duke Energy and Lambda public statements, cited at the relevant claims. (www.dominionenergy.com) (www.aepohio.com) (investors.duke-energy.com) (lambda.ai)

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/dgxb300-user-guide/introduction-to-dgxb300.html>

2. <https://www.energy.gov/cmei/femp/cooling-water-efficiency-opportunities-federal-data-centers>
3. <https://www.thegreengrid.org/node/372>
4. <https://docs.nvidia.com/dgx-superpod/reference-architecture/scalable-infrastructure-b300/latest/dgx-superpod-components.html>
5. https://datacenters.lbl.gov/sites/default/files/DataCenterMeteringandResourceGuide_02072017_0.pdf
6. <https://www.dominionenergy.com/virginia/large-business-services/data-center-requests>
7. <https://lambda.ai/blog/lambda-closes-1-billion-senior-secured-fixed-rate-financing>
8. https://handbook.ashrae.org/Handbooks/A23/IP/A23_Ch20/a23_ch20_ip.aspx
9. <https://gpusmith.com/articles/en/gpu-rack-power-requirements-planning-guide>
10. <https://gpusmith.com/articles/en/ai-data-center-demand-response-interconnection-plan>
11. <https://investors.duke-energy.com/news/news-details/2026/Duke-Energy-Data-center-growth-will-deliver-billions-of-dollars-in-customer-savings/default.aspx>
12. <https://www.aepohio.com/company/about/rates/data-center-tariff/>
13. <https://www.dell.com/en-us/shop/ipovw/poweredge-xe9780?hve=shop+now>
14. <https://www.gigabyte.com/topics/nvidia-ai-factory>
15. <https://docs.nvidia.com/dgx/dgxb200-user-guide/introduction-to-dgxb200.html>
16. <https://lenovopress.lenovo.com/lp2264-thinksystem-sr680a-v4-server>
17. https://www.supermicro.com/datasheet/datasheet_SuperCluster_B300_2OU_Systems.pdf
18. <https://www.eia.gov/energyexplained/electricity/measuring-electricity.php>
19. <https://www.iso.org/standard/30134-2>
20. https://www.energystar.gov/products/data_center_equipment/16-more-ways-cut-energy-waste-data-center/reduce-energy-losses-power-distribution-units-pdus
21. <https://gpusmith.com/articles/en/gpu-cluster-energy-calculator>
22. <https://gpusmith.com/tools/gpu-power-calculator>
23. <https://legacy.wbdg.org/building-commissioning/commissioning-documents>
24. <https://www.ul.com/thecodeauthority/knowledge/circuit-breaker-guide>
25. <https://gpusmith.com/articles/en/data-center-electrical-design-gpu-clusters>
26. <https://gpusmith.com/articles/en/gpu-data-center-cooling-design>
27. https://productinfo.se.com/0100db2301_elecdistfunddesignguide/0100db2301-electrical-distribution-fundamentals-design-guide/English/0100DB2301%20Electrical%20Distribution%20Fundamentals%20Design%20Guide%20%28bookmap%29_DD00648155.xml/%24/EmergencyAndStandbyPowerSystems-A64217F8
28. https://www.eaton.com/us/en-us/catalog/backup-power-ups-surge-it-power-distribution/Eaton_BladeUPS_UPS/-modular-ups-technology--eaton-.html
29. <https://journal.uptimeinstitute.com/explaining-uptime-institutes-tier-classification-system/>
30. <https://gpusmith.com/articles/en/gpu-power-cap-economics-inference-slos>
31. <https://www.vertiv.com/en-cn/services-catalog/services/project-services/commissioning/>
32. <https://gpusmith.com/articles/en/nvidia-dsx-due-diligence-guide>
33. <https://gpusmith.com/gpu-financing-due-diligence>

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.