

How Fast Does AI Load Move? What 321 Public Sources Do and Don't Measure

Data period: Public-source review · 2026-10-13

Generated: 2026-10-13 22:47 (Europe/Riga)

Compiled exclusively from publicly available sources. See Disclaimer on the last page.

How Fast Does AI Load Move?

What 321 Public Sources Do and Don't Measure

Based on a review of 321 public sources (920 extracted findings), 2026-10-13. The specific findings below cite 55 of those sources by number; the full list with links is in Sources at the end. Figures are quoted only where a source gives them; claims from vendors or press are marked, and points nobody has measured are stated as such.

In one page

- **Across 321 sources, the quantified training-load picture collapses to a single production telemetry set.** Microsoft/OpenAI/NVIDIA's "Power Stabilization for AI Training Datacenters" (arXiv:2508.14318) is the one paper with real production measurements, and what it reports is a **frequency band** — training-power fluctuation energy concentrated around **0.2–3 Hz** — **not a rate of change**. Almost every other serious number points back to it, is a calculation, a vendor model, or an unsourced claim.
- **Seconds, not milliseconds — and that is the bad news.** The measured band (~0.2–3 Hz, i.e. cycles of seconds) sits exactly where bulk-grid components have their own oscillation modes (turbine-generators, long lines). A millisecond spike a capacitor absorbs; a 0.2–3 Hz swing with tens of MW behind it is in the grid's own danger band.
- **Amplitude is calculated, not measured.** The widely-quoted **~66.7 MW** for a 100,000-GPU H100 cluster is an explicit *upper-bound calculation* (Michigan State, arXiv:2606.22096: $100k \times (700\text{ W} - \sim 100\text{ W}) \div 0.9$; real shapes smoother), and even its "1–10 s period" is cited to the same Microsoft/OpenAI/NVIDIA paper. That paper itself states amplitude only generally — "tens or hundreds of megawatts."
- **Rate of change (kW/s, MW/s) is measured nowhere.** The widely-repeated "**>1,000 MW/s at gigawatt scale**" traces to one **sponsored** article (T&D World, sponsor EGM) that attributes it to the Microsoft/OpenAI/NVIDIA paper — **but that paper contains no such figure**. Treat it as unsourced. The only other rate numbers are vendor models (battalion.energy, ~150–260 MW/s) and a rack figure (~132 kW/s, NVIDIA GB200 NVL72, vendor). Two auto-captured "measurements" were discarded outright — a current-slew figure (200 A/ms $\neq$ a power ramp) and a grid-policy limit (0.33%/s).
- **The rack arithmetic still stands — as arithmetic.** Taking the vendor 132 kW/s per rack against the one explicit regulatory limit anywhere (US MISO, 30 MW/min = 500 kW/s for a whole site): **$500 \div 132 \approx 3.8$** — fewer than four such racks moving in sync would exhaust a whole site's ramp allowance. Arithmetic from a vendor figure, not a measurement.
- **Grid operators *have* measured oscillations near large loads — but at higher frequencies, a different mechanism.** Dominion Energy reported ~14.7 Hz (~4% peak-to-peak) at a 115 kV substation; ERCOT/Texas operators observed ~23 Hz near a large electronic load (Oct 2024); Meta saw 49/71 Hz (2017). These are real field measurements of control/power-electronics interactions — **not** the 0.2–3 Hz training-load band — and should not be conflated with it.

- **Colocation contracts price capacity (kW), not rate-of-change (kW/s).** Of 54 contract-related findings, none is an actual contract setting a rate-of-change obligation — the higher-grade items are vendor reference designs. **Absence of evidence:** the risk appears unpriced, but that is inferred.
- **Europe:** Across the European grid-code material captured — Ireland, Poland, France, Denmark — not one entry sets a rate-of-change limit on large loads. Finland, Sweden, Germany, Netherlands were not captured in this review, so Europe is a confirmed negative only for the codes listed. The only explicit large-load ramp limit anywhere in the data is in the US (MISO, 30 MW/min).

The size of the swings

What AI clusters actually do to their power draw, as measured or described in the sources:

- A 160% of nominal fast overshoot at the edge of a compute phase lasting 400 microseconds. [_\(measured; xpcc.com \[1\]\)_](#)
- A 136% of nominal sustained burst across a phase lasting 50 milliseconds. [_\(measured; xpcc.com \[1\]\)_](#)
- Meta's data centers experienced oscillations at 49 Hz and 71 Hz in 2017. [_\(measured; sentientenergy.com \[2\]\)_](#)
- Texas grid operators observed 23 Hz oscillations in real power measurements near a large electronic load in October 2024. [_\(measured; sentientenergy.com \[2\]\)_](#)
- Dominion Energy reported 14.7 Hz oscillations with a peak-to-peak magnitude of approximately 4% in voltage measurements at a 115 kV substation. [_\(measured; sentientenergy.com \[2\]\)_](#)
- PMU data reported an oscillation frequency of approximately 7.5 Hz with a peak-to-peak magnitude of 25 MW. [_\(measured; linkedin.com \[3\]\)_](#)
- High-resolution waveform data showed an actual oscillation frequency of approximately 23 Hz and a peak-to-peak magnitude of 50 MW. [_\(measured; linkedin.com \[3\]\)_](#)
- The aggregate power therefore traces a deterministic square wave at the training iteration period, whose reciprocal sets the forcing frequency. [_\(measured; arxiv.org \[4\]\)_](#)
- For H100 clusters of 100 000 GPUs behind a 0.9-efficient power chain, the swing reaches 66.7 MW at the point of common coupling. (Upper-bound calculation, $100k \text{ GPU} \times (700W \sim 100W) / 0.9$ — not a facility measurement.) [_\(calculation, not measured; arxiv.org \[4\]\)_](#)
- Datacenter Scale-up Scenario Synchronization Mean Peak Peak-to-Average 25,000 H100 Fully sync 13.1 MW 18.9 MW 1.44 [_\(model from measured traces; github.com \[5\]\)_](#)
- Datacenter Scale-up Scenario Synchronization Mean Peak Peak-to-Average 100,000 H100 Fully sync 52.4 MW 75.6 MW 1.44 [_\(model from measured traces; github.com \[5\]\)_](#)
- Datacenter Scale-up Scenario Synchronization Mean Peak Peak-to-Average 1 GW DC Fully sync 877 MW 1,266 MW 1.44 [_\(model from measured traces; github.com \[5\]\)_](#)
- Datacenter Scale-up Scenario Synchronization Mean Peak Peak-to-Average 1 GW DC Fully random 877 MW 879 MW 1.00 [_\(model from measured traces; github.com \[5\]\)_](#)
- Power changes exceeding 10 MW/second during training burst transitions [_\(vendor/press; ipqdf.com \[6\]\)_](#)

How fast the load moves

The rate of change is described qualitatively everywhere but quantified nowhere for an AI cluster — the central open question:

- At 1 GW the load falls from about 970 MW to about 425 MW and comes back every 20 seconds or so, with rises of up to 147 MW in 100 ms. [_\(measured; battalion.energy \[7\]\)_](#)
- Eight independent jobs move about 260 MW in a second. [_\(measured; battalion.energy \[7\]\)_](#)
- Three training halls among seven inference halls move about 150 MW in a second. [_\(measured; battalion.energy \[7\]\)_](#)
- The data center cluster reduced consumption gradually with a 15-minute ramp down, maintained the 25% power reduction over three hours, then ramped back up without exceeding its original baseline consumption. [_\(measured; blogs.nvidia.com \[8\]\)_](#)
- Eaton UPS units provided an instantaneous response to a sudden loss of 800+ MW of production capacity, with an overall load of less than 400 kW but a response of approximately 1 MW. [_\(measured; linkedin.com \[9\]\)_](#)
- A single transmission fault in the Eastern Interconnection caused 1,800 MW of data center demand to disconnect simultaneously in milliseconds. [_\(measured; techtimes.com \[10\]\)_](#)
- Four more Eastern Interconnection events followed in the same year: 428 MW, 227 MW, 540 MW, and 1,300 MW. [_\(measured; techtimes.com \[10\]\)_](#)
- ERCOT recorded nine separate cryptocurrency mining load-loss events exceeding 100 MW each. [_\(measured; techtimes.com \[10\]\)_](#)
- A major voltage-sensitive load reduction event involved a 1,500 MW load reduction consisting entirely of data center-type load. [_\(measured; datacenterfrontier.com \[11\]\)_](#)
- A North American data center ramped down from about 450 MW to about 40 MW within 36 seconds. [_\(measured; linkedin.com \[12\]\)_](#)
- Training jobs can ramp tens of megawatts in seconds. [_\(measured; xpcc.com \[1\]\)_](#)
- The mean facility active power is 409.14 kW and the peak facility active power is 723.33 kW in the 2026 file, with a mean PUE of 1.552 and annual energy consumption of approximately 3.584 GWh. [_\(vendor/press; ieee-dataport.org \[13\]\)_](#)
- The 2024 file reports a mean facility active power of approximately 390 kW, a peak of approximately 600 kW, and a mean PUE of approximately 1.50, with annual energy consumption of approximately 3.41 GWh. [_\(vendor/press; ieee-dataport.org \[13\]\)_](#)
- EasyRider can smooth millisecond-scale transients to a slow change over 30 seconds or more. [_\(vendor/press; ar5iv.labs.arxiv.org \[14\]\)_](#)

What the equipment can take

Limits of the gear between the chips and the grid (UPS, generators, transformers, power quality):

- Heavy duty UPS systems demonstrate $\pm 15\%$ voltage deviation and 16ms recovery time meeting IEC 62040-3 standards. [_\(measured; bkpowers.com \[15\]\)_](#)

- UPS systems transfer to backup power on a deviation of just 5 to 10 percent sustained beyond 40 to 66 milliseconds. [_\(measured; tdworld.com \[16\]\)_](#)
- UPS circuits that protect servers now destabilize the grid. [_\(measured; techtimes.com \[10\]\)_](#)
- ISO 8528-5 categorizes generator sets into four performance classes (G1 to G4) based on their voltage and frequency requirements. [_\(vendor/press; scribd.com \[17\]\)_](#)
- UPS systems should limit total harmonic distortion (THD) to 3% at 100% load and 5% at 50% load, and mitigate leading power factor to less than 5% at loads less than 10%. UPS step load acceptance is +/- 2%. [_\(vendor/press; mitsubishicritical.com \[18\]\)_](#)
- ISO 8528-5 defines performance classes for generator sets, with G1 being the least stringent, suitable for small, simple generating sets intended to supply unsophisticated loads. [_\(vendor/press; avtronpower.com \[19\]\)_](#)
- ISO 8528-5 defines performance classes for generator sets, with G2 being broadly equivalent to commercially available power. [_\(vendor/press; avtronpower.com \[19\]\)_](#)
- ISO 8528-5 defines performance classes for generator sets, with G3 intended for sets powering strategically critical loads, such as hospitals and data centers. [_\(vendor/press; avtronpower.com \[19\]\)_](#)
- ISO 8528-5 defines performance classes for generator sets, with G4 reserved for performance criteria agreed between the supplier and the buyer. [_\(vendor/press; avtronpower.com \[19\]\)_](#)
- Most generator sets are rated at 0.8 power factor, meaning 80 percent real power and 20 percent reactive. [_\(vendor/press; anvilfield.com \[20\]\)_](#)
- A 2000 kW, 0.8 PF generator requires 2000 kW resistive plus reactive to reach about 2500 kVA. [_\(vendor/press; anvilfield.com \[20\]\)_](#)
- Data center generator sizing has a well-established industry benchmark: total installed generator capacity is approximately 2× the IT load. [_\(vendor/press; primegenset.com \[21\]\)_](#)

The grid connection

What grid codes and connection agreements require of large loads:

- EN 50160 is the European reference standard that describes the main characteristics of the voltage supply delivered by public distribution networks under normal operating conditions. [_\(vendor/press; satec-global.com.au \[22\]\)_](#)
- EN 50160 provides a framework for characterising voltage variations, typically by magnitude and duration. [_\(vendor/press; satec-global.com.au \[22\]\)_](#)
- EN 50160 defines flicker using standardised indices so that assessments are repeatable and comparable. [_\(vendor/press; satec-global.com.au \[22\]\)_](#)
- EN 50160 treats voltage unbalance as a supply characteristic to be monitored and kept within limits for most of the time under normal conditions. [_\(vendor/press; satec-global.com.au \[22\]\)_](#)
- EN 50160 addresses voltage harmonics as part of supply quality. [_\(vendor/press; satec-global.com.au \[22\]\)_](#)
- EN 50160 defines the key voltage characteristics of electricity supplied by public distribution systems under typical operating conditions. [_\(vendor/press; satec-global.com \[23\]\)_](#)

- DCCOPP V3 requires new dispatchable generation or storage equal to 100% of MIC (de-rated) before a connection offer can issue for data centres. [_\(vendor/press; philiplee.ie \[24\]\)_](#)
- The low-voltage ride-through envelope for Poland (PSE Wymogi 2025) specifies a linear ramp from 0 pu at 150 ms to 0.85 pu at 2.5 s. [_\(vendor/press; bess.courses \[25\]\)_](#)
- Voltage quality is monitored according to EN 50160 [_\(vendor/press; a-eberle.de \[26\]\)_](#)
- Approximately 20% of planned data center projects face significant delays due to grid congestion and interconnection backlogs. [_\(vendor/press; techplustrends.com \[27\]\)_](#)
- MISO caps new large-load ramp rates at 30 MW per minute. [_\(vendor/press; theenergystoragewire.com \[28\]\)_](#)
- NERC is developing registration criteria and reliability standards for computational loads with a proposed threshold of 20 MW or more of aggregated computational load at an interconnection point of 60 kV and above. [_\(vendor/press; zeroemissiongrid.com \[29\]\)_](#)
- Grid support from inverter-based resources (IBRs) is enabled by bidirectional power converters and governed by interconnection standards such as IEEE 1547. [_\(vendor/press; ashrae.org \[30\]\)_](#)
- EN 50160 standard is referenced for voltage flicker Pst/Plt. [_\(vendor/press; powerquality.blog \[31\]\)_](#)

Where the problem gets absorbed

Which layer is being used to smooth the swings:

- The battery absorbs the fast transients in real time, turning starting GPU power swings into steady predictable demand. [_\(measured; linkedin.com \[32\]\)_](#)
- Software-based approach enables AI data centers to operate as flexible grid resources. [_\(measured; aixenergy.io \[33\]\)_](#)
- NVIDIA has moved energy storage into the power shelf: electrolytic capacitors taking about half the PSU volume, 65 joules per GPU, which on an instrumented GB200 rack cut the peak power the grid saw by up to 30% while training Megatron. [_\(vendor/press; bess.courses \[34\]\)_](#)
- Meta runs an always-on software smoother that launches dummy kernels during power dips — nearly 800 W per GB200 accelerator of deliberately wasted work — for under 3% impact on application performance. [_\(vendor/press; bess.courses \[34\]\)_](#)
- The GB300 power shelf includes electrolytic capacitors that charge during low-demand intervals and discharge during peaks. [_\(vendor/press; aimagazine.com \[35\]\)_](#)
- The GB300 power shelf reduces grid-facing peak power by 30% while maintaining the same GPU output pattern. [_\(vendor/press; aimagazine.com \[35\]\)_](#)
- The GB300 power supply unit allocates about half its volume to capacitors, delivering 65 joules per GPU of energy storage co-designed with LITEON Technology and managed by a new charge management controller. [_\(vendor/press; developer.nvidia.com \[36\]\)_](#)
- Capacitor-based energy storage in the GB300 NVL72 system provides 65 joules of energy storage per GPU. [_\(vendor/press; futurumgroup.com \[37\]\)_](#)
- NVIDIA GB300 NVL72 features provide steady power for AI workloads, indicating power smoothing at the chip level. [_\(vendor/press; forums.developer.nvidia.com \[38\]\)_](#)

- The GB300 NVL72 power supply unit integrates energy storage elements, specifically electrolytic capacitors, to smooth rapid, short-term power fluctuations during steady-state operation. [_\(vendor/press; news.nvinio.com \[39\]\)_](#)
- Eaton is selling the XLHV, a supercapacitor bank the size of a single unit of a server rack, capable of dynamically delivering up to 420 kilowatts of power. [_\(vendor/press; spectrum.ieee.org \[40\]\)_](#)
- Siemens Energy is selling the E-statcom, a supercapacitor bank that can cycle up to 75 megawatts per unit. [_\(vendor/press; spectrum.ieee.org \[40\]\)_](#)

What colocation contracts say

Whether contracts address rate-of-change at all:

- The data center must minimally provide N+1 power, where N equals two power sources. [_\(vendor/press; docs.nvidia.com \[41\]\)_](#)
- The contract term mentions 200 kW as a contracted power level. [_\(vendor/press; computelaw.blog \[42\]\)_](#)
- PPAs can span 10 to 20 years to provide long-term price stability and support the financing for the construction of the power generating facility. [_\(vendor/press; pillsburylaw.com \[43\]\)_](#)
- Financial security of \$50,000 per MW of contracted peak demand. [_\(vendor/press; computelaw.blog \[44\]\)_](#)
- A tenant commits to 2 MW of liquid-cooled capacity across 20 racks at 100 kW each, for five years. [_\(vendor/press; altstreet.investments \[45\]\)_](#)
- Ramp schedule specifies 25 percent of capacity billable in month one, 100 percent by month seven. [_\(vendor/press; altstreet.investments \[45\]\)_](#)
- Colocation pricing in 2026 is structured predominantly around kilowatts of committed power rather than rack units or square footage. [_\(vendor/press; encoradvisors.com \[46\]\)_](#)
- High-density AI colocation runs at 20-100 kW or more per rack, requiring purpose-built infrastructure. [_\(vendor/press; encoradvisors.com \[46\]\)_](#)
- Colocation contracts typically have terms of 5–10 years. [_\(vendor/press; thechiefnegotiators.com \[47\]\)_](#)
- Anchor tenant leases commonly run 10 to 20 years at \$100 to \$150 per kilowatt per month, with the tenant paying base rent plus all operating costs. [_\(vendor/press; computelaw.blog \[48\]\)_](#)

When one tenant affects the neighbours

Whether one tenant's swinging load degrades others on a shared electrical block:

- more than three fourths of highly-distorted power readings across the country are within 50 miles of significant data center activity. [_\(measured; linkedin.com \[49\]\)_](#)
- One Virginia fault dropped 1,500 megawatts of AI load in under a second. [_\(measured; tdworld.com \[16\]\)_](#)
- Harmonic currents can cause interference with sensitive load equipment across tenant rooms. [_\(vendor/press; centiel.com \[50\]\)_](#)

- The complex interaction between multiple power conversion stages creates cascaded dynamic limitations that constrain system response capabilities. [_\(vendor/press; arxiv.org \[51\]\)_](#)
- One rack's droop can become upstream variability that affects neighbors and shared plant. [_\(vendor/press; ure.us \[52\]\)_](#)
- A voltage fluctuation in northern Virginia impacted multiple data centers simultaneously. [_\(vendor/press; belfercenter.org \[53\]\)_](#)
- One tenant's rectifier load can shape the power quality another tenant sees. [_\(vendor/press; cnbygele.com \[54\]\)_](#)
- Harmonic disturbances can extend beyond the data center itself, affecting neighboring facilities connected to the same electrical infrastructure. [_\(vendor/press; datacenterfrontier.com \[55\]\)_](#)

What nobody has measured yet

- A primary, independent facility measurement of ramp (kW/s or MW/s) with a stated method. None exists in these findings: the only rate numbers are vendor models ($\approx 150\text{--}260$ MW/s) and a per-rack vendor figure (~ 132 kW/s), while the widely-quoted ">1,000 MW/s at gigawatt scale" is an unsourced attribution (see "Rate of change... measured nowhere" above) — none is an independent measurement.
- Whether one tenant's oscillation measurably degrades a neighbour's power quality on a shared block.
- Any European grid-code or connection requirement that caps how fast a large load may change.

What this study cannot tell you

These are the structural gaps — not missing sources, but things the public record does not establish. They bound every conclusion above.

- **Load shape is never laid against an equipment-class limit.** ISO 8528-5 generator classes (G1–G4) and UPS step-load specs appear only as definitions. No source puts a measured AI load shape next to a class's accepted load step, frequency dip and recovery time. So "rate of change is the binding constraint" is asserted, not quantified against the thing it supposedly binds.
- **UPS operating mode is a blank.** Nothing on eco-mode vs double-conversion. A double-conversion UPS with a DC bus absorbs part of the swing; eco-mode passes it straight to the grid — and many large sites run eco-mode for efficiency. Whether a swing even reaches the distribution network depends on a switch position the corpus never mentions.
- **It is all training, not inference.** Every number here — 0.2–3 Hz, 66.7 MW, the vendor rates — is a training load. Inference is demand-driven, not iteration-synchronised, and may not oscillate the same way. As the market tilts toward inference, the whole risk profile may change; this study cannot say how.
- **The value side is unquantified.** A site that can move ~ 150 MW/s is also the fastest frequency-response resource on the system — with markets and a price, especially in the Nordics/Baltics. The corpus barely touches ancillary-service / flexibility value (fast frequency response, FCR/aFRR $\approx$ absent). The same characteristic that is a risk may be an asset nobody is buying.

- **The author's own region is absent.** Nordic/Baltic grids (Fingrid, Litgrid, Elering, Svenska kraftnät) are essentially not in the data, yet the Nordic synchronous zone has inter-area modes near the training band.
- **There is no skeptic in the corpus.** 800+ findings, one recorded contradiction — and it is about tariffs, not oscillation. Either the field genuinely agrees, or the search surfaced only those who see a problem, most of whom sell a solution. For a body of evidence where nearly every number comes from someone selling something, the absence of a counter-voice is a weakness, not a confirmation — weight the conclusions accordingly.

About this research

Sources were gathered from public material (papers, standards, vendor and operator publications, grid-code documents). Measurements are kept separate from vendor or press claims throughout, and where no figure exists the brief says so rather than estimating. This is a findings brief, not an opinion piece.

Sources

Numbers match the [n] markers in the findings above. 55 sources cited here, drawn from 321 reviewed in total.

1. xpcc.com — <https://xpcc.com/ai-power-transients/>
2. sentientenergy.com — <https://sentientenergy.com/insights/managing-the-grid-ai/>
3. linkedin.com — <https://www.linkedin.com/pulse/monitoring-data-center-sub-synchronous-oscillations-address-kamal-bctyf?tl=en>
4. arxiv.org — <https://arxiv.org/html/2606.22096v1>
5. github.com — <https://github.com/IsabellaLiu99/genai-gpu-power-analysis>
6. ipqdf.com — <https://ipqdf.com/case-studies/harmonics/ai-data-centres-and-power-quality-a-new-category-of-grid-disturbance/>
7. battalion.energy — <https://battalion.energy/solutions/ai-data-centers/>
8. blogs.nvidia.com — <https://blogs.nvidia.com/blog/ai-factories-flexible-power-use/>
9. linkedin.com — <https://www.linkedin.com/pulse/data-centers-source-fast-frequency-response-virtual-inertia-paananen>
10. techtimes.com — <https://www.techtimes.com/articles/319695/20260704/ai-data-centers-triggered-1800-mw-grid-drop-nerc-issues-highest-alert.htm>
11. datacenterfrontier.com — <https://www.datacenterfrontier.com/energy/article/55376679/why-nerc-now-sees-ai-data-centers-as-grid-actors>
12. linkedin.com — <https://www.linkedin.com/pulse/evolving-nature-large-electrical-loads-impact-power-system-ahmad-jxwbf>
13. ieee-dataport.org — <https://ieee-dataport.org/documents/multi-scale-hourly-electrical-load-dataset-ai-driven-data-centers-active-and-reactive>
14. ar5iv.labs.arxiv.org — <https://ar5iv.labs.arxiv.org/html/2604.15522>

15. bkpowers.com — <https://bkpowers.com/service/output-voltage-stability-under-load-step-changes-heavy-duty-ups-test-data/>
16. tdworld.com — <https://www.tdworld.com/sponsored/article/55400084/how-ai-data-centers-are-changing-power-demand-on-the-grid>
17. scribd.com — <https://www.scribd.com/document/894443027/04-09-Transient>
18. mitsubishicritical.com — <https://mitsubishicritical.com/resources/minimize-ups-downtime-with-critical-power-system-design/component-integration/>
19. avtronpower.com — <https://avtronpower.com/resources/articles/testing-diesel-gen-sets-to-iso8528/>
20. anvilfield.com — <https://anvilfield.com/field-guides/datacenter/load-bank-test-acceptance-criteria/>
21. primegenset.com — <https://primegenset.com/diesel-generator-sizing-guide-for-data-centers/>
22. satec-global.com.au — <https://satec-global.com.au/en-50160-power-quality-frequency-voltage-variation-s-flicker-unbalance-harmonics/>
23. satec-global.com — <https://www.satec-global.com/en-50160-reports-power-quality/>
24. philiplee.ie — <https://www.philiplee.ie/eirgrid-dccopp-version-3-what-data-centre-operators-need-to-know/>
25. bess.courses — <https://bess.courses/grid-codes/poland-pse-wymogi-2025/>
26. a-eberle.de — <https://www.a-eberle.de/en/special-publications/power-quality-monitoring-jupiter-data-centre/>
27. techplustrends.com — <https://techplustrends.com/power-requirements-ai-data-centers/>
28. theenergystoragewire.com — <https://theenergystoragewire.com/miso-caps-new-largeload-ramp-rates-at-30-mw-per-minute-forci/>
29. zeroemissiongrid.com — <https://www.zeroemissiongrid.com/iso-rto/miso-llwg-05-14/>
30. ashrae.org — <https://www.ashrae.org/technical-resources/topics-and-initiatives/ai-data-center/grid-inter-active-design-demand-flexibility>
31. powerquality.blog — <https://powerquality.blog/2021/02/01/voltage-disturbances-flicker-measurement/>
32. linkedin.com — https://www.linkedin.com/posts/alan-cooper-33b1a62a_data-centers-are-increasingly-vibrated-as-grid-activity-7432816325867237376-PTf
33. aixenergy.io — <https://www.aixenergy.io/dcflex/>
34. bess.courses — <https://bess.courses/articles/sizing-bess-for-ai-data-center-load-ramps/>
35. aimagazine.com — <https://aimagazine.com/news/how-nvidia-gb300-nvl72-provides-data-centres-steady-ai-power>
36. developer.nvidia.com — <https://developer.nvidia.com/blog/how-new-gb300-nvl72-features-provide-steady-power-for-ai/>
37. futurumgroup.com — <https://futurumgroup.com/insights/is-nvidias-gb300-nvl72-power-shelf-the-answer-to-ais-grid-strain/>
38. forums.developer.nvidia.com — <https://forums.developer.nvidia.com/t/how-new-gb300-nvl72-features-provide-steady-power-for-ai/340494>
39. news.nvinio.com — <https://news.nvinio.com/how-new-gb300-nvl72-features-provide-steady-power-for-ai-20029.html>
40. spectrum.ieee.org — <https://spectrum.ieee.org/supercapacitor-2671883490>

41. docs.nvidia.com — <https://docs.nvidia.com/dgx-superpod/design-guides/dgx-superpod-data-center-design-h100/latest/electrical.html>
42. computelaw.blog — <https://computelaw.blog/construction/key-terms-ai-data-center-construction-contract/>
43. pillsburylaw.com — <https://www.pillsburylaw.com/en/news-and-insights/power-purchase-interconnection-agreements-data-centers.html>
44. computelaw.blog — <https://computelaw.blog/power/ercot-large-load-interconnection-texas/>
45. altstreet.investments — <https://altstreet.investments/reference/mechanics/colocation-lease>
46. encoradvisors.com — <https://encoradvisors.com/data-center-colocation-pricing/>
47. thechiefnegotiators.com — <https://www.thechiefnegotiators.com/data-center-capacity-sourcing>
48. computelaw.blog — <https://computelaw.blog/deals/anchor-tenant-colocation-leases-ai-data-centers/>
49. linkedin.com — <https://www.linkedin.com/top-content/technology/data-center-operations/grid-impact-of-high-power-data-centers/>
50. centiel.com — <https://www.centiel.com/ups-efficiency-colocation-data-centers/>
51. arxiv.org — <https://arxiv.org/html/2502.01647v2>
52. ure.us — <https://ure.us/articles/predictive-power-conditioning-gpu-clusters/>
53. belfercenter.org — <https://www.belfercenter.org/research-analysis/ai-data-centers-us-electric-grid>
54. cnbygele.com — <https://cnbygele.com/blog/colocation-data-center-tenant-power-quality/>
55. datacenterfrontier.com — <https://www.datacenterfrontier.com/sponsored/article/55383213/harmonic-distortion-in-data-centers-a-power-quality-planning-guide>

Appendix: collected evidence

Disclaimer

This report was compiled exclusively from publicly available sources, as accessed on the generation date shown on the cover. No confidential, proprietary or insider information was used or sought. Source URLs are listed in the Sources section; excerpts and third-party data are reproduced for research and citation purposes only, and all trademarks and product names remain the property of their respective owners.

The information is provided "as is", without warranty of any kind as to accuracy, completeness or fitness for any purpose. Underlying sources may themselves contain errors; where evidence is weak or contradictory this is indicated in the text (evidence grades, confidence labels, preserved contradictions). All estimates, models and forward-looking statements are analytical interpretations by the author, not statements of fact, and may prove incorrect.

Nothing in this report constitutes investment, legal, engineering or other professional advice. Any decisions based on this material are made solely at the reader's own risk. If you believe any statement concerning you or your organisation is inaccurate, corrections will be made promptly upon substantiated request via gailitis.ceo.