

PART 4 · SECOND EDITION · SEPTEMBER 2026 (FIRST PUBLISHED JULY 2026)

Electrical & Energy Infrastructure

*Every conversion stage between the grid and the die is a tax.
Every rating you pour into the slab is a bet on a rack density
you can't yet see.*

JULY 10, 2024 · NORTHERN VIRGINIA

~1,500 MW

One failed lightning arrestor. Six faults in 82 seconds. 1,500 MW of data-center demand off the grid.

Every facility did exactly what it was designed to do — protect its own electronics. Together, they became the grid event.

NERC, INCIDENT REVIEW: LARGE LOAD LOSS · AS-OF 2024-07

THE REFRAME

The AI hall breaks three electrical instincts at once.

AI loads can step from idle to full power in milliseconds and in phase. Active-PFC reduces distortion, but verify compliance at the PCC; UPS transfer can also create a fast grid-load change.


ACT I

The chain: grid to die

A conventional chain includes several voltage conversions from utility service to the GPU core rail. Busway does not convert voltage; declare the count and loss for the specific one-line.

THE TERRITORY The grid-to-die chain — every stage is a tax and a lead time

Fewer packages, less loss, at the same boundary



Compare losses, fault behaviour and delivery dates, not stage counts alone.

Guide scenario · Source: guide derivation · as of September 2026

Compare AC rack input, AC-fed DC sidecar and direct MV/DC at the same final-output and failure boundaries; package count alone does not establish loss or delivery.

WHAT COLLAPSING THE CHAIN BUYS

The 800 VDC case, in four numbers

4 → 3 functional blocks Guide fixture: four AC versus three integrated MV/DC packages, including final conversion GUIDE ILLUSTRATIVE FUNCTIONAL LEDGER DATED SEPTEMBER 8, 2026 · AS-OF 2026-09	About 80–87% AC case (82% representative) Guide AC fixture, 82% representative — not the analyst 800 VDC model; losses separate GUIDE DERIVATION, SEPTEMBER 8, 2026 · AS-OF 2026-09-08	~98% ~98% at 400 kW is a US DOE prototype (INTELEC 2025); qualify offered SSTs at the boundary KOLAR ET AL., EFFICIENT SSTs FOR FUTURE DATACENTERS, INTELEC 2025 KEYNOTE · AS-OF 2025-10-13	+157% NVIDIA claims +157% over the same wire, 415 VAC → 800 VDC (Oct 2025; basis unpublished) NVIDIA, BUILDING THE 800 VDC ECOSYSTEM · AS-OF 2025-10-13
--	--	---	---

THE ARCHITECTURAL DECISION

AC to the rack, or DC from the perimeter?

Legacy AC-to-the-rack

- 415/480 VAC distribution; rectify once per rack
- Four guide packages: about 80–87% utility-to-VRM (82% representative; distribution loss separate)
- Mature: every vendor, every code official
- Accept only within offered current, thermal and surviving-feed ratings

DC-disaggregated ($\pm 400/800$ VDC)

- Rectify once per hall, at the facility perimeter
- Three packages in the direct-MV/DC guide fixture; an AC-fed sidecar retains upstream plant
- Ramping: DC breakers, busway, SST listings still maturing
- Qualify the named roadmap rack, insulation and fault/service envelope

SO WHAT

Retain AC where its delivered capacity, losses, service access and cost win; select DC where the qualified offer resolves a current, routing or loss constraint. Rack kW alone does not decide.

THE MOVING TARGET

Rack power by named product — the voltage architecture must outlive all three

**132 kW nominal TDP:
115 kW liquid + 17 kW
air**

HPE GB200 NVL72 (2024): 132 kW nominal rack TDP

HPE, NVIDIA GB200 NVL72 BY HPE DATA SHEET, OCTOBER 2026 EDITION. · AS-OF 2026-10

**188 kW (Max Q) / 228
kW (Max P)**

Pegatron Vera Rubin NVL72 (2026): 188 kW Max Q / 228 kW Max P; 330 kW DSX facility basis

PEGATRON RA4803-72N3 DATASHEET · AS-OF 2026-08

~600 kW

Kyber NVL144 (NVIDIA target H2 2027, contested): ~600 kW per-rack plan on 800 VDC

NVIDIA GTC · AS-OF H2 2027 ROADMAP (GTC 2025); 2028 REPORTED JULY 2026

THE SECOND FORK

Which 800 V rail topology: bipolar ± 400 V or unipolar 800 V?

± 400 VDC bipolar — OCP Mt. Diablo 1.0.0

- Backed by Google, Meta, Microsoft via OCP
- EV components are candidates; qualify installed DC stress and fault duty
- Ideal solid midpoint: nominal ± 400 V to earth; study actual transients
- Supply-chain economics as the headline driver

800 VDC nominal — NVIDIA reference

- NVIDIA plus Eaton, ABB, Schneider, Vertiv, Siemens
- Equal 800 V spans carry equal ideal current at equal load
- 1,200 V-class SiC candidates; qualify actual topology and interrupting duty
- Vertiv's October 2025 H2 2026 forecast is not proof of offered-SKU shipment


SO WHAT

Read each named output scheme: solid midpoint, HRG or isolated. Its fault current, rail-to-earth stress, detection and service isolation decide the interface.

THE FORK, DRAWN **Compare the full DC interface: reference, fault behavior and service access**

Equal rail span gives equal load current

Illustrative · same power P across the same span U



Guide scenario · Source: Guide derivation · as of Sep 2026

OCP Diablo and NVIDIA reference designs require an offered grounding scheme. Equal full-span voltage gives equal ideal load current; it does not give equal first-fault behavior.

THE PHYSICS **Current, not voltage, is the enemy**


At 600 kW, 48 V needs 12,500 A; 800 V needs 750 A

Same delivered power. Loss at the same resistance scales with current squared.

Ideal current before losses ($I = P/V$)



Copper at equal current density (schematic)



What 48 V costs you

- Copper mass
- Contact + connector heating
- Many parallel feeds

What 800 V costs you

- Insulation + shock protection
- DC arc interruption
- Touch-safe access

Current, connectors and cooling decide the usable rating.


Guide scenario · Source: guide derivation · as of September 2026

Illustrative 600 kW DC comparison: calculate current and losses, then qualify conductor and connector temperature, cooling, protection and surviving-feed duty.

THE YARD Closing the tie changes the fault duty

Closing the tie can exceed the breaker duty

400 V · 630 kVA / 6% each · common 500 MVA source



Keep the tie open and interlocked, or buy gear rated for the closed-tie duty. Making, withstand, selectivity and arc flash are separate checks.

Guide scenario · Source: guide derivation · as of September 2026

One transformer gives about 15 kA; the tied pair about 29 kA against a 25 kA screen. Keep the tie open or interlocked, or qualify the higher duty — a shared source impedance is not halved.

THE PROCUREMENT CLOCK

The schedule is set by the yard, not the building

Order the long-lead electrical items at design freeze; most of them outlast a GPU generation.

Host-utility large-load service path	Studies, agreement, and upgrades
HV / substation power transformer	>160 wk (US, WoodMac 2026)
Switchgear (voltage class not specified)	~44 wk (US, WoodMac Q2 2025)
Aeroderivative gas turbines (bridge power)	~40 mo (Enverus, Sept 2026)
Reciprocating engines (bridge power)	15–24 months


GPU generation cycle those numbers must survive ~12–24 months

The MV bus rating is a one-way density bet: a yard sized to today's racks cannot be re-rated by buying a bigger breaker next quarter.

THE SECOND CLOCK Behind every turbine order sits a second critical path: the gas train

Every link in the gas train has its own owner and calendar

The first two links run on other companies' calendars, not yours.



Contract fork: firm transport pays a reservation charge; interruptible needs the dashed backups.

Schematic · Source: Guide synthesis · as of August 2026

Tie-in, metering, conditioning, compression, and a firm-or-interruptible contract each carry their own lead time. Lock the fuel chain with the prime movers, or the engines arrive to sit idle.

ACT II


The load that moves like one machine

Thousands of active-PFC rectifiers can step together. Synchronized transients are a major design case, while harmonics remain a measured PCC compliance obligation rather than disappearing.

THE INVERSION Measure both harmonic compliance and synchronized transients

Heating and load steps need separate tests

Schematic traces; the measured load and source set the actual duty.



Test the PSU spectrum at the load point; the PCC limit is a separate test.

K-rating tolerates winding heat; it does not suppress harmonics.

Schematic · Source: ABB + Eaton · as of September 2026

Active-front-end PSUs can draw clean current, but verify aggregate harmonics, voltage distortion and transformer heating at the boundary. A K-rating tolerates heating; it does not suppress harmonics.

THE SIZING BASIS

~400 J/GPU

On-rack energy storage is the engineered transient budget — Vera Rubin's reported target is ~6× GB300 per GPU


Size the chain from the named product's published peak and transient envelope, never a TDP multiplier. A TDP-summed chain passes a steady load-bank test and trips on the real all-reduce peak.

NVIDIA TECHNICAL BLOG, NVIDIA VERA RUBIN POD: SEVEN CHIPS, FIVE RACK-SCALE SYSTEMS, ONE AI SUPERCOMPUTER · AS-OF 2026-03


THE MITIGATION SPINE **A relay of buffers: chip → rack → facility**

A relay of storage layers absorbs the GPU load step

Arrows pass the residual step upstream; the sketch above each shows it smaller, slower.



Residual = load – source



Solid: load · dashed: source · shaded area = energy per event

Test every handoff for

- power
- delay
- energy per event
- recharge before the next step

Never spend the outage reserve on pulses.

Drop a layer and the step lands on the next, slower one.

Guide scenario · Source: guide synthesis · as of September 2026

Select each interface by its residual pulse, reachable power and energy, response and recharge; preserve outage reserve if supervisory smoothing fails.

ABSORPTION, BY THE NUMBERS

How much energy it takes to hide a GPU from the grid

65 J/GPU

Energy storage in each GB300 NVL72 power shelf — half the PSU volume is now capacitance

NVIDIA DEVELOPER BLOG · AS-OF 2025-07

~400 J/GPU

Reported on-rack storage target for Vera Rubin — a ~6× jump in one generation

NVIDIA TECHNICAL BLOG, NVIDIA VERA RUBIN POD: SEVEN CHIPS, FIVE RACK-SCALE SYSTEMS, ONE AI SUPERCOMPUTER · AS-OF 2026-03

30%

Peak grid-demand reduction demonstrated on a real Megatron-LLM training run

NVIDIA DEVELOPER · AS-OF 2025-07

THE DECISION THAT SIZES YOUR INTERCONNECTION

Smooth hard at the rack, or buy a bigger connection?

Smooth at the rack

- Over-provision capacitance, BBU, and supercaps (à la 400 J/GPU)
- Demonstrate bounded grid demand with reserved outage energy and failed-smoothing response
- Offer the utility ramp-rate compliance
- Costs silicon area, PSU volume, and rack capex

Pass the transient upstream


- Cheaper racks; the facility BESS and headroom absorb the swing
- Qualify the larger residual against actual source, storage and connection envelopes
- Utility may impose ramp-rate limits anyway
- Interconnection is the scarcest asset you own

SO WHAT

Decide early — the answer sizes the one thing you cannot expedite: the grid connection.

THE ISLAND CASE A behind-the-meter plant must catch the same millisecond step — without a grid to lean on**The POI changes source and protection duty**

Opening the POI removes the utility's fault current and inertia from the bus



Bars show relative size, not to scale.

Island: one source sets voltage and frequency; faults must still trip relays.

Inertia acts first; storage carries the load until governor and fuel catch up.

Schematic · Source: guide synthesis · as of September 2026

The online synchronous machine answers first through torque and inertia; storage bridges the governor and fuel recovery. An island needs a machine or inverter reference and its own protection.

ACT III

The grid now plans against you

NERC's May 4, 2026 alert makes computational-load behavior a planning input. Recipient acknowledgment/reporting is required; the Essential Actions differ from effective regional rules and site contracts.

THE SAME FAULT, THREE POSTURES

Take the operative curve before fixing settings

POSTURE	ACCEPTED PROJECT ENVELOPE	ENGINEERING CONSEQUENCE
Trip/transfer	Submit the actual pickup, delay and recovery settings with the dynamic model	Model the resulting grid-demand loss and retained IT/cooling duty
Hold within the curve	SPP HILL table (0.90–1.10 pu continuous) or ERCOT NOGRR 282 (≥90% of load back within 2 s)	Coordinate relay, UPS and rectifier settings; qualify reachable storage
Ride plus VAR support	Utility sets the POI range (EKPC: 0.95 lagging–0.98 leading); credit revenue only under an executed tariff	Qualify inverter headroom or STATCOM; reserve outage energy

THE ESCALATION

Level 3

NERC Level 3: separate reporting duties from recommended actions

The May 4, 2026 alert, p.2, requires recipient acknowledgment/reporting but makes the Essential Actions nonmandatory. Effective regional rules and the site agreement establish actual settings.

NERC, LEVEL 3 ALERT ON COMPUTATIONAL LOADS · AS-OF 2026-05-04

THE PROPOSED REGISTRATION THRESHOLD

≥ 50 MW @ ≥ 100 kV

CLO/CLOP: a ≥ 50 MW site fed at ≥ 100 kV registers its owner and operator

NERC's 2026-08-19 posting (Appendices 2/5A/5B) is a proposal until filed under FERC Order RD26-7-000 and approved; verify the operative appendices and site-aggregation rule before assigning duties.


NERC, PROPOSED RULES OF PROCEDURE APPENDICES 2/5A/5B • AS-OF 2026-10-08

BACK TO THE ARRESTOR

The strike walks in from outside — and each zone boundary takes a bite

Each zone entrance gets an SPD to the shared bond

Dashed lines: zone boundaries. Surge leaves through the bond, not the load.



SPDs clamp surges; they do not clear a sustained DC earth fault.

Study ground potential rise at remote-grid bonds; IEC SPD types are not UL types.

Schematic · Source: IEC + OCP · as of September 2026

Coordinate IEC impulse duties from entrance to equipment; pick UL installation types separately. Bond buildings and remote-grid ties from GPR and transferred potential; prove DC first-fault detection.

THE THROUGH-LINE

Grid asset or grid risk: the same building, two opposite reliability postures — and you choose between them in relay settings, not press releases.

TAKE IT HOME

What to do Monday

- 01 Order the transformer and switchgear at design freeze — the yard's lead times, not the building's, set your energization date.
- 02 Assign continuous load, peak/EDPp versus duration and recovery to each boundary; summed chip TDP is not a facility load case.
- 03 Choose AC or DC by the offered current, thermal, fault and service envelope plus installed cost; rack kW alone gives no crossover.
- 04 Size chip, rack and facility buffers by the residual each must cover; prove recharge without spending the outage reserve.
- 05 Write ride-through behavior into the interconnection agreement and the relay settings — trip-to-protect is now a flagged risk.
- 06 Commission with step loads and electronic load banks — a resistive test proves ampacity, not survival.

|| THE DEFINITIVE GUIDE TO
AI Data Centers

This deck is the map. The guide is the territory — every number here is dated, sourced, and kept current in the full chapters.

- Ch 4.1 Voltage selection and the conversion-stage discipline
- Ch 4.5 UPS and storage as transient shock-absorbers
- Ch 4.7 The DC power revolution: bipolar ± 400 V vs unipolar 800 V rail topology
- Ch 4.10 Ride-through, NERC, and grid-interactive behavior

READ THE FULL PART, WITH EVERY SOURCE AND DERIVATION, AT AIDATACENTERGUIDE.COM

SOURCES

Every figure in this deck traces to the guide's dated register (1/3)

NERC, Incident Review: Large Load Loss (2025-01-08), pp.7-8;
NERC comments to FERC, Docket EL25-49 (2025-04-23), Attachment
A (April 17 presentation), PDF p.17 · as-of 2024-07 ·
https://www.nerc.com/globalassets/our-work/reports/event-reports/incident_review_large_load_loss.pdf

Guide illustrative functional ledger dated September 8, 2026;
OCP deployment-path discussion published August 11, 2026,
opened September 8. The count is a guide derivation, not an OCP
efficiency measurement. · as-of 2026-09 ·
<https://www.opencompute.org/index.php/blog/powering-the-next-era-of-ai-how-google-microsoft-and-nvidia-are-standardizing-and-accelerating-the-industry-transition-to-lvdc>

Guide derivation, September 8, 2026; Chapter 4.1 assumed four-
package AC ledger. · as-of 2026-09-08

Kolar et al., Efficient SSTs for Future Datacenters, INTELEC
2025 keynote (ETH Zurich, 2025-10-13) · as-of 2025-10-13 ·
https://www.ams-publications.ee.ethz.ch/uploads/tx_ethpublications/INTELEC_2025_-_2_30PM_Kolar_99___Efficient_SSTs_for_Future_Datacenters_Keynote_231025.pdf

NVIDIA, Building the 800 VDC Ecosystem (Oct 2025) · as-of 2025-
10-13 · <https://developer.nvidia.com/blog/building-the-800-vdc-ecosystem-for-efficient-scalable-ai-factories/>

HPE, NVIDIA GB200 NVL72 by HPE data sheet, October 2026
edition. · as-of 2026-10 ·
<https://www.hpe.com/psnow/generateDDS/NVIDIA%20GB200%20NVL72%20by%20HPE%20data%20sheet-PSN1014890104MYEN.pdf>

SOURCES

Every figure in this deck traces to the guide's dated register (2/3)

Pegatron RA4803-72N3 datasheet · as-of 2026-08 ·
https://svr.pegatroncorp.com/SpecFile/PEGATRON_RA4803-72N3_DS_2026v1.pdf

NVIDIA GTC (Jensen Huang); DCD, Tom's Hardware · as-of H2 2027 roadmap (GTC 2025); 2028 reported July 2026 ·
<https://www.tomshardware.com/pc-components/gpus/nvidia-shows-off-rubin-ultra-with-600-000-watt-kyber-racks-and-infrastructure-coming-in-2027>

Wood Mackenzie (Ben Boucher), reported by Data Center Knowledge, 2026-05-04 · as-of 2026-05 ·
<https://www.datacenterknowledge.com/build-design/ai-data-center-boom-rewires-us-power-supply-chain>

Wood Mackenzie / American Clean Power, Making the Connection, September 2025 · as-of 2025-Q2 · <https://www.powermag.com/wp-content/uploads/2026/01/making-the-connection-executive-summary-september-2025.pdf>

SemiAnalysis · as-of 2025-12 ·
<https://newsletter.semianalysis.com/p/how-ai-labs-are-solving-the-power>

NVIDIA Technical Blog, NVIDIA Vera Rubin POD: Seven Chips, Five Rack-Scale Systems, One AI Supercomputer (2026-03-16) · as-of 2026-03 · <https://developer.nvidia.com/blog/>

NVIDIA Developer Blog / LITEON · as-of 2025-07 ·
<https://developer.nvidia.com/blog/how-new-gb300-nvl72-features-provide-steady-power-for-ai/>

NVIDIA Developer (GB300 steady power) · as-of 2025-07 ·
<https://developer.nvidia.com/blog/how-new-gb300-nvl72-features-provide-steady-power-for-ai/>

SOURCES

Every figure in this deck traces to the guide's dated register (3/3)

SPP, updated High Impact Large Load Ride-Through Requirements, revision 2025-09-23, RR696 package (filename V2.0) · as-of 2025-09-23 · <https://spp.org/Documents/74204/RR696.zip>

NERC, Level 3 Alert on computational loads (initial distribution May 4, 2026) · as-of 2026-05-04 · <https://www.nerc.com/globalassets/programs/bpsa/alerts/level-3-computational-load-alert.pdf>

NERC, proposed Rules of Procedure Appendices 2/5A/5B (Computational Load Owner/Operator registry criteria), second-round posting 2026-08-19; Steptoe analysis · as-of 2026-10-08 · <https://www.nerc.com/initiatives/large-loads-action-plan>

SOURCES

About these figures

Figures are resolved by reference from the guide's claim register at build time (Second edition · September 2026 (first published July 2026)) — a number appears here only as it appears, dated and sourced, at aidatacenterguide.com. Newer data may exist: check the live register.