DIRECT POWER ALLIANCE · TECHNICAL WHITEPAPER

An 800 VDC AI Factory

Native-DC power from generation to rack — a reference architecture for behind-the-meter AI campuses

Version 1.1 · June 2026

About this paper. This is a technical whitepaper of the Direct Power Alliance, a neutral consortium defining an open, native-DC power standard for AI data centers. It was seeded by a contribution from ISS Mining and is offered for open review and revision by members. It presents design considerations and a worked reference case; it is not a specification, and it complements — and will be reconciled with — the Alliance’s Reference Architecture v0.1. Figures for unreleased products are drawn from public vendor roadmaps and are cited; they should be treated as directional.

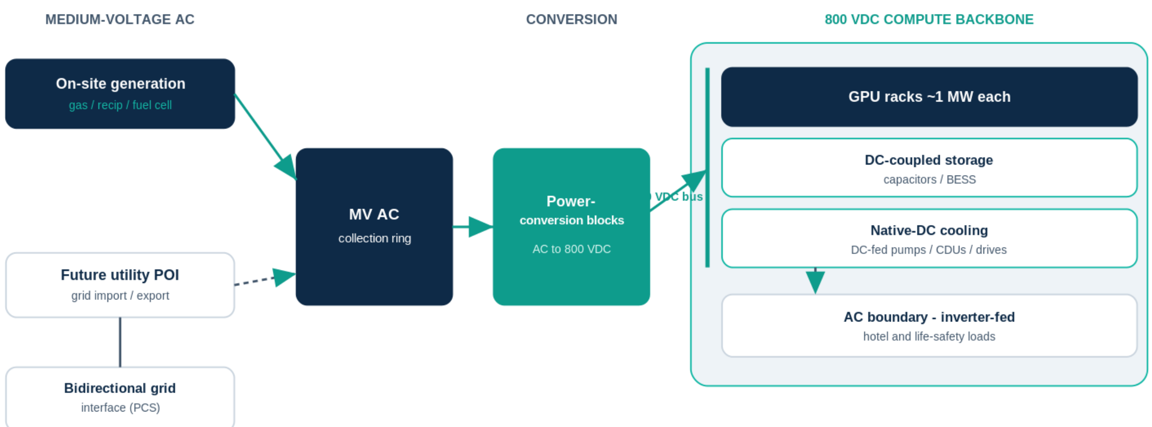

Figure 1. Conceptual campus architecture. 800 VDC is the compute backbone — racks, DC-coupled storage, and native-DC cooling all share the bus — while AC is preserved at the boundaries where it creates value: rotating generation, the utility interconnection, and inverter-fed hotel and life-safety loads.

Executive summary

NVIDIA-class AI racks are moving the power problem from a secondary data-center utility into the core production architecture. NVIDIA has described an industry transition toward 800 VDC data-center power infrastructure for 1 MW IT racks and beyond, starting in 2027, to reduce current, copper, conversion stages, and the rack space consumed by power equipment [1][2]. For a behind-the-meter AI campus, the implication is bigger than rack power: the project must decide where AC belongs, where DC belongs, and how the site will later connect to the grid without stranding early infrastructure.

The recommended pattern is a DC-first, AC-boundary architecture. Use medium-voltage AC where it is naturally strong — rotating generation, transformers, utility interconnection, and long campus feeders — but establish 800 VDC as the data-hall compute backbone. Convert as close to the data halls as practical, distribute 800 VDC over short, high-ampacity busways to rack power sidecars, and integrate storage directly on or near the DC bus. Then invert only where AC is actually needed: legacy hotel loads, life-safety systems, and utility export/import.

Two clarifications shape the rest of this paper. First, "800 VDC" is not yet a single settled convention. The ecosystem is still deciding between an 800 V monopolar bus and a ±400 V bipolar bus, and that choice propagates into connectors, protection, grounding, and touch-safety. Settling it is precisely the kind of work a neutral standard body exists to do. Second, 800 VDC is a distribution voltage, not a campus bulk-transmission voltage: at 1 MW per rack an 800 VDC feeder carries about 1,250 A, which is manageable over data-hall distances with busway and parallel conductors, but is the wrong voltage for moving tens or hundreds of megawatts across a campus. For those distances, use medium-voltage AC, or a qualified MVDC architecture once mature equipment exists.

The largest non-IT load — cooling — deserves the same DC-first scrutiny as compute. Today chillers, pumps, fans, and CDUs are AC equipment, which leaves an AC island inside an otherwise-DC campus. Where OEM support exists, DC-fed compressors, pumps, fans, and motor drives can sit natively on the 800 VDC bus, removing a conversion stage; AC remains a deliberate boundary for legacy and life-safety equipment, not the default for everything.

Finally, the future grid connection should be designed from day one, not bolted on later. U.S. large-load interconnection is evolving quickly: FERC opened Docket RM26-4-000 for large loads generally above 20 MW and, on June 18, 2026, directed the six RTOs/ISOs to justify or reform tariffs for large energy users, including co-located and behind-the-meter generation [8][9].

Primary recommendations.

1. Adopt 800 VDC as the IT-load backbone, not the whole-campus transmission voltage. Convert with local AC-to-DC or DC-to-DC power blocks near data halls and racks.
2. Use layered storage: rack-level capacitors for milliseconds-to-seconds, data-hall DC-coupled storage for seconds-to-minutes, and facility BESS/generation for ride-through, ramping, and grid services.
3. Treat cooling as a DC-bus load where practical. Favor DC-fed drives and DC-native cooling equipment; keep AC as a boundary for legacy and life-safety auxiliaries.
4. Keep AC at the boundaries — inverter-fed AC buses for hotel loads and a bidirectional grid interface for import/export — but do not route the majority IT load through unnecessary AC stages.
5. Engineer for interconnection now: specify the telemetry, models, protection, and controls that will satisfy utility and RTO/ISO review even while the first phase operates behind the meter.

1. Context and decision frame

The reference case used throughout this paper is a behind-the-meter AI data-center campus that initially relies on dedicated or co-located power and later expects to interconnect with the grid. The compute fleet consumes power at 800 VDC at or near the rack, following NVIDIA ecosystem direction. An operator must therefore choose a power architecture that works across three time horizons:

- **Initial energization:** operate from on-site generation and storage with limited or no utility capacity.

- **Expansion:** add racks, data halls, BESS, and renewable or fuel-cell resources without rebuilding the electrical backbone.
- **Grid integration:** connect as a controllable large load, and potentially export, while serving AC hotel loads and meeting interconnection requirements.

The key decision is not simply AC versus DC. The useful question is: at each point in the chain from fuel, sunlight, or grid energy to GPU silicon, what voltage and current minimize total cost, losses, failure modes, and schedule risk? NVIDIA states that traditional 54 VDC in-rack distribution begins to hit physical limits as racks exceed roughly 200 kW, citing space, copper, and conversion inefficiency, and that an 800 VDC architecture can support 1 MW IT racks and beyond, starting in 2027 [1]. In a later 2025 ecosystem update, NVIDIA described converting medium-voltage AC directly to 800 VDC with large power-conversion systems, then distributing 800 VDC through the data hall to compute racks [2].

2. Why 800 VDC changes the architecture

The historical Tesla-versus-Edison lesson is often summarized as AC defeating DC. The more useful engineering lesson is that high voltage defeated low voltage. Resistive loss is I^2R; for a fixed power P, current is approximately P/V, so conductor loss falls with the square of voltage. Early AC won because transformers made voltage conversion cheap and reliable. Modern power electronics reopens the DC option, especially where generation, storage, and end-use electronics are already DC internally.

For a two-wire DC feeder with copper conductors, a planning-level line-loss approximation is $f = (2 \cdot \rho \cdot L \cdot P) / (A \cdot V^2)$, where ρ is copper resistivity, L is one-way length, A is conductor area, P is delivered power, and V is the DC bus voltage. Using standard annealed copper at 20 °C, the 800 VDC shortcut is L (m) ≈ 18,560 · f · A(mm^2) / P(kW). For a 2% loss allowance, L ≈ 371 · A / P; for 5%, L ≈ 928 · A / P. These are planning tools only; real designs must include temperature, conduit grouping, busway construction, terminal losses, redundancy, overload, and code/AHJ requirements.

Table 1 shows the key insight: 800 VDC is excellent for rack, row, and data-hall distribution, but it is not a campus bulk-transfer voltage at tens or hundreds of megawatts. Moving 10 MW at 800 V requires 12,500 A; even very large parallel copper sets become distance-limited. Use medium voltage for long runs and convert near the data halls.

Load on 800 VDC feeder	Current	Copper area / pole	Distance @ 2% loss	Distance @ 5% loss
250 kW	313 A	500 mm^2	742 m	1,856 m
1 MW	1,250 A	500 mm^2	185 m	464 m
10 MW	12,500 A	500 mm^2	19 m	46 m
10 MW	12,500 A	2,000 mm^2	74 m	186 m

Table 1. Planning-level one-way distances at 800 VDC. 800 VDC suits the rack and the hall, not long campus transfer.

The comparison with 54 VDC is dramatic. A 1 MW rack at 54 V would draw about 18,500 A before conversion losses; at 800 V it draws about 1,250 A — a 14.8× current reduction and about a 219× reduction in I^2R loss for the same conductors. This is why NVIDIA frames 800 VDC as a way to reduce copper, conversion stages, and power-equipment volume in the compute space [1][2].

2.1 An unsettled convention: 800 V monopolar vs. ±400 V bipolar

"800 VDC" is shorthand for a family of conventions that the ecosystem has not yet standardized. The two leading candidates are an 800 V monopolar bus (a single 800 V rail referenced to return) and a ±400 V bipolar bus (two ±400 V rails about a midpoint, presenting 800 V rail-to-rail). The choice is not cosmetic: it changes connector and busbar design, the protection and grounding scheme, insulation coordination, and — importantly for safety — the touch-voltage a technician can encounter. Because every connector, converter, and protection device is built to whichever convention wins, leaving it unsettled fragments the supplier ecosystem. Converging on a convention, a protection scheme, and a certification path is the core reason a neutral standards body is needed, and it is the first item on the Alliance's agenda.

3. Architecture options

Three architectures are practical candidates. They are not mutually exclusive; most real campuses phase through them (Table 2). The recommended path is a hybrid of Options B and C: keep a medium-voltage AC collection ring for rotating generation, utility interconnection, and long feeders; at each data hall, use modular power-conversion blocks to create an 800 VDC compute backbone; and design those blocks so DC-coupled batteries, capacitors, and future native-DC sources connect without forcing energy through unnecessary AC stages. This is constructible near-term while preserving a path to a native-DC AI factory.

Option	Description	Advantages	Limitations / best use
A. Conventional AC	MV utility/generation to LV AC; AC UPS/PDU/busway; rack PSUs convert to DC.	Mature equipment, familiar protection, simple hotel-load integration, utility-friendly.	Multiple conversion stages; rack-space pressure; high LV copper. Use for legacy halls or a first phase before the 800 VDC supply chain is ready.
B. 800 VDC overlay	Conventional MV/LV feeds centralized rectifiers; 800 VDC distributed inside the hall to rack sidecars.	Good transition path; removes repeated in-rack AC/DC conversion while preserving utility/generation AC.	Still duplicates some AC and DC equipment; storage integration depends on converter topology. Use for near-term NVIDIA-class racks in partly-conventional campuses.
C. Native-DC hall	Facility power blocks convert MV AC or native-DC sources to 800 VDC; BESS is DC-coupled; AC mainly for hotel loads and the grid POI.	Highest compute-path efficiency; direct storage integration; fewest conversions; most future-proof.	HVDC protection, controls, qualification, and AHJ familiarity are more demanding. Use for greenfield AI factories designed around MW-class racks.

Table 2. Three architecture options; most campuses phase A → B → C.

4. Generation-to-compute power chain

Most behind-the-meter generation produces AC at medium voltage: gas turbines, reciprocating engines, and synchronous generator systems. Forcing the entire plant into low-voltage DC is not beneficial. The economical pattern is to generate at MV AC, collect and protect at MV AC, then convert near the compute halls to 800 VDC. Where sources are native DC — PV arrays, batteries, fuel cells — use DC/DC interfaces into the 800 VDC (or higher-voltage DC) collection bus, provided protection and isolation requirements are met.

A practical first-phase chain: on-site generators feed an MV AC ring with conventional breakers, relays, synchronization, and black-start provisions; power-conversion skids near each data hall use active front-end rectifiers or solid-state transformers to produce regulated 800 VDC; 800 VDC busway distributes over short distances to rack sidecars; rack-level converters step down to local board voltages close to the load; and DC-coupled storage connects at rack, row, hall, and plant levels. For long feeders, medium voltage remains preferred because loss falls with V^2 — moving the same power at 13.8 kV instead of 800 V cuts conductor loss by roughly $(13{,}800/800)^2$, about 298×, before converter and transformer losses. The conclusion is not that AC must always be used; it is that 800 VDC should be the consumption voltage at the hall and rack, while long-distance movement uses a higher voltage.

5. Storage as an active part of the architecture

AI power is not just large; it is dynamic. NVIDIA notes that AI factories behave differently from conventional data centers because synchronized workloads create rapid load swings, and its 2026 BESS guidance frames the BESS as an integrated system of batteries, power conversion, telemetry, and controls that shapes how power is absorbed, injected, and regulated in real time [4]. Treat storage as a control asset, not just backup runtime, layered by the timescale it must serve (Table 3).

Layer	Timescale	Location	Mission / notes
Rack capacitors / CBU	ms to s	In rack, sidecar, or power shelf	Absorb GPU transient spikes; reduce peak demand. NVIDIA describes GB300 smoothing via caps, storage, and burn, cutting peak grid demand up to ~30% in its example [5].
Data-hall DC storage	s to min	On/near the 800 VDC hall bus	Smooth residual hall ramps; bridge source transfers; cut converter oversizing. DC-coupling avoids AC double-conversion but needs mature HVDC protection.
Facility BESS	min to hours	Generation yard, POI, or MV bus	Grid-friendly load profile, ride-through, generator support, black start. Useful across grid-connected, generator-coordinated, and islanded modes [4].
Fuel / thermal reserve	hours to days	Generation plant, fuel supply	Sustained operation during outage or islanded first phase; coordinate emissions, fuel logistics, redundancy.

Table 3. A layered storage stack, matched to the timescale each layer serves.

The control objective is to make the AI factory look electrically boring upstream while the GPUs remain computationally aggressive downstream. That requires state-of-charge policies, ramp-rate limits, active/reactive priorities, telemetry, and event recording — and site models that include IT and non-IT load behavior, UPS modes, protection settings, reconnect logic, generation behavior, and BESS controls [4].

6. Cooling and hotel loads: move the largest auxiliary onto the bus

Cooling is the largest non-IT load in an AI campus — chillers, pumps, CDUs, cooling towers, and fans — and today nearly all of it is AC equipment driven by AC motors and AC-fed variable-speed drives. In an otherwise-DC campus, that leaves an AC island whose drives must be fed by inverting the DC bus back to AC, re-introducing exactly the conversion stages 800 VDC was meant to remove. The design question is whether AC should be the primary auxiliary backbone or a deliberate boundary layer around the DC compute backbone.

The recommendation is to move cooling onto the DC bus wherever OEM support and maintainability allow, and to keep AC as a boundary for legacy and life-safety equipment. Three patterns are available (Table 4). The most native is DC-fed motor drives and DC-native cooling equipment — compressors, pumps, and fans whose drives take the 800 VDC bus directly, avoiding a rectifier front end. Today this is a thin product category: essentially no one builds 800 VDC-native cooling plants, because cooling equipment evolved on AC. That gap is itself an opportunity — and the reason the Alliance convenes the cooling-equipment and motor-drive supply chains alongside the power-equipment makers, so the largest auxiliary load can be built to live on the bus rather than beside it. (The Alliance addresses the power interface of cooling; thermal and fluid standards remain with ASHRAE and the OCP cooling community, with whom it liaises.)

Strategy	How it works	Trade-offs	Recommended role
DC-fed drives & DC-native equipment	Pumps, fans, compressors, and CDUs use drives or DC links fed from the 800 VDC bus.	Avoids rectifier front ends and improves controllability; vendor support and code familiarity still maturing.	Preferred where OEM support exists — moves the largest auxiliary load onto the bus.
Inverter-fed AC from the DC bus	Grid-forming inverters on the 800 VDC bus supply 480/415 VAC panels.	Lets storage and compute share one DC backbone and makes AC loads controllable; the inverter becomes a critical path needing redundancy.	For flexible/critical AC loads that must remain AC.
Dedicated AC auxiliary bus	Generators/utility feed conventional AC transformers for hotel loads on a separate path.	Mature, simple for chillers and motors, easy AHJ familiarity; may duplicate UPS/BESS and is least integrated.	Fallback and life-safety baseline; minimize its share over time.

Table 4. Cooling and hotel-load strategies, ordered from most-native-DC to AC fallback.

The efficiency stakes are real. If a hall delivers 250 MW to compute and a conventional path is 90% efficient while an 800 VDC path is 95% efficient, input power falls from about 278 MW to about 263 MW — roughly 15 MW saved before cooling effects, and every megawatt not lost is a megawatt not rejected as heat. Exported power is a separate matter: a utility grid is AC, so any DC-coupled generation or storage that exports must pass through a grid interface — bidirectional PCS, transformer, relaying, metering, controls. That defines the AC boundary at the point of interconnection, not inside every rack.

7. Grid interconnection and export readiness

A project that expects later grid connection should design the interconnection package in Phase 1, even if it energizes without export. Retrofits are expensive because the interconnection package touches protection, grounding, metering, SCADA, inverter settings, transformer ratings, harmonic studies, and operating procedures.

Two frameworks matter for U.S. projects, and they govern the AC boundary — not the DC equipment inside the hall. First, interconnection of distributed energy resources at distribution voltages generally invokes IEEE 1547-2018, which defines technical requirements and testing for the interface between DER and the electric power system, including performance, operation, testing, safety, power quality, and islanding [6]. UL 1741 covers inverters, converters, controllers, and interconnection equipment; for interactive equipment, UL states its requirements supplement IEEE 1547 and IEEE 1547.1 [7]. These are interconnection standards for the AC interface; they do not, by themselves, certify 800 VDC distribution equipment inside the hall (see Section 8). Second, very large loads and co-located generation can be subject to RTO/ISO, transmission-owner, and FERC processes: FERC describes RM26-4-000 as focused on large loads generally above 20 MW [8], and its June 18, 2026 action directed the six RTOs/ISOs to address co-location agreements, behind-the-meter generation, flexible large-load transmission service, and studies of generators serving proximate large loads [9].

Specify the future grid interface as a controllable energy port, not a simple breaker — with bidirectional active power (import, export, zero-export limits); reactive/voltage support; grid-forming capability for islanding or black start and grid-following where required; ride-through, anti-islanding, reconnection, and ramp-rate behavior; revenue-grade metering and high-speed disturbance recording; validated EMT and positive-sequence models for utility studies; and cybersecure dispatch/curtailment interfaces. Write the operating modes into the control philosophy before procurement (Table 5).

Mode	Utility tie	Generation	BESS role	Load behavior	Export
Islanded start	Open	Grid-forming gen or BESS	Black start, V/f support	Staged ramp of hotel then IT load	None
Generator-coordinated	Open / limited	Gen supplies average load	Smooth ramps; hold reserve	Follows internal ramp limits	None / local
Grid-parallel, import-only	Closed	On-site optional	Peak shave, ride-through	Utility sees controlled net load	Zero-export enforced
Grid-parallel, export-limited	Closed	Gen / BESS dispatchable	Hold SOC and export ramp	Compute flex may support grid	Up to contract limit
Emergency support	Closed / islanded	As available	Fast response; hold critical loads	Curtail noncritical compute	As directed

Table 5. Operating modes to define before procurement.

8. Protection, controls, safety — and the certification gap

800 VDC distribution is not a drop-in replacement for 480 VAC distribution. DC faults do not naturally pass through a current zero-crossing, and constant-power electronic loads can present negative incremental impedance. Treat protection and controls as first-order design elements:

- **DC protection:** qualified DC breakers, fuses, contactors, or solid-state interruption; selective coordination and safe isolation for maintenance.
- **Grounding & insulation monitoring:** choose grounded, impedance-grounded, or ungrounded DC topology deliberately; include insulation monitoring and ground-fault response. (This choice is tied to the monopolar-vs-bipolar convention of Section 2.1.)
- **Arc-flash & touch safety:** 800 VDC is hazardous; specify enclosures, interlocks, labeling, PPE boundaries, emergency stop, and service procedures with the AHJ.
- **Converter stability:** model active front ends, DC/DC converters, BESS PCS, and rack loads together; validate with hardware-in-the-loop where practical.
- **Cybersecurity:** BESS, PCS, and workload controllers influence megawatts of load; treat control systems as critical infrastructure.
- **Maintainability:** modularize power blocks so a failed converter, string, or bus segment can be isolated without dropping a training cluster.

A distinct gap deserves emphasis. The interconnection standards of Section 7 (IEEE 1547, UL 1741) govern how the site connects to the grid; they do not provide a complete listing and certification path for 800 VDC distribution equipment inside the data hall — DC switchgear, busway, connectors, and converters. UL and IEC product listings for this class of native-DC equipment, and the code language and AHJ pathway to deploy it, do not yet fully exist. Closing that gap — alongside the unsettled voltage convention — is exactly the work of the Alliance’s Codes & Certification activity.

A critical operational risk is synchronized AI load behavior. The NVIDIA/Microsoft/OpenAI power-stabilization paper explains that large AI training workloads create large power swings because GPUs work in lockstep, and that at scale these swings are visible at rack, data-center, and grid levels, with potential resonance concerns if they harmonize with utility critical frequencies [10]. Mitigate in layers: GPU/rack power smoothing, rack or row storage, BESS buffering, and site-level dispatch.

9. Economic implications

The business case for 800 VDC is a stack of avoided costs and operational benefits, not a single efficiency number:

- **Conversion reduction:** fewer AC/DC stages between source, storage, and silicon reduce losses and heat; NVIDIA states 800 VDC can improve end-to-end efficiency by up to 5% [1].
- **Copper & space:** higher voltage means lower current, reducing conductor area, busbar bulk, terminal size, and conductor cooling.
- **Compute density:** moving conversion out of the IT rack frees rack space for GPUs. Flex states its 800 VDC power rack uses a disaggregated architecture that, for a future GPU platform, raises available compute per rack from about 125 kW to up to 880 kW [3].

- **Grid readiness:** controllable BESS and PCS help the site look like a flexible large load, potentially improving interconnection negotiations [4][8][9].
- **Cooling:** every megawatt not lost in conversion is a megawatt not rejected as heat, reducing both electrical loss and cooling burden.

The principal counter-costs are the immaturity premium for 800 VDC protection, specialized converters, qualified busway, spares, training, AHJ review, and vendor bankability. Compare these against conventional AC infrastructure scaled to MW-class rack density — not against a legacy 10–30 kW/rack data center.

10. Implementation roadmap

The roadmap below phases the campus from basis-of-design through native-DC expansion, so that each phase reaches time-to-power without stranding later decisions (Table 6).

Phase	Electrical scope	Controls & studies	Commercial objective
0. Basis of design	Define rack roadmap, 800 VDC voltage window/convention, redundancy, generation, storage duration, hotel-load split, future POI.	Load-flow, short-circuit, arc-flash, dynamic model, AI ramp profile, interconnection pre-application.	Avoid stranded design assumptions.
1. Behind-the-meter energization	MV generator ring, first conversion blocks, 800 VDC hall bus, AC/auxiliary path, rack storage/CBU.	Islanded controls, black-start sequence, generator ramp protection, no-export logic if a tie exists.	Reach time-to-power without full grid capacity.
2. Interconnection-ready buildout	Install space, conduits, transformers, switchgear, metering, PCS bays for the future POI.	IEEE 1547 / UL 1741 evidence where applicable; RTO/utility models; protection coordination.	Preserve schedule flexibility and bankability.
3. Grid-parallel operation	Close POI; operate import-only or export-limited per agreement.	Ride-through, reactive support, ramp control, SOC policy, dispatch telemetry.	Use the grid as capacity, reliability, and market option.
4. Native-DC expansion	Add DC-coupled PV/fuel cells/BESS and DC-fed cooling; more 800 VDC halls; MVDC links if mature.	Campus optimizer coordinates compute, storage, generation, cooling, and grid constraints.	Maximize compute per constrained MW; reduce energy cost.

Table 6. A phased roadmap from basis-of-design to native-DC expansion.

11. Risk register and mitigations

The principal risks are technical maturity, fault interruption, controls interaction, interconnection uncertainty, auxiliary dependency, storage safety, and operational complexity (Table 7).

Risk	Why it matters	Mitigation
800 VDC equipment & convention maturity	Roadmaps, listings, service models, and the monopolar/bipolar convention are still developing.	Procure from ecosystem vendors; require qualification evidence, spares, service commitments, interoperability testing; track the standard.
HVDC fault interruption	DC arcs and faults are harder to interrupt than AC.	Listed DC breakers/fuses/solid-state devices; selective coordination and arc-flash

		studies; engage AHJ early.
Controls interaction	PCS, BESS, generators, and constant-power racks can oscillate if poorly coordinated.	Site-level dynamic model; require vendor models; HIL/EMT studies; staged commissioning.
Grid interconnection uncertainty	Rules for large loads, co-location, and export are evolving.	File early; design export-limited modes; track FERC/RTO/ISO changes; keep zero-export capability.
Cooling / auxiliary dependency	Native-DC cooling is immature, and AC auxiliaries fed by DC-bus inverters add a critical path.	Qualify DC-fed cooling carefully; use N+ redundancy, bypass feeds, segmented AC buses; keep life-safety on conservative AC.
BESS safety & degradation	Storage is mission-critical and can create fire, thermal, and warranty risks.	UL/NFPA-compliant systems, thermal management, fire detection/suppression, duty-cycle warranties, SOC reserve.
Operational complexity	A DC-first AI factory behaves like a power plant and data center combined.	Staff for power-plant operations, SCADA, relay settings, cybersecurity, and workload-aware dispatch.

Table 7. Risk register and mitigations.

12. Conclusion

The right answer is not to choose AC or DC universally, but to assign each to the voltage level and system boundary where it has the advantage. AC remains compelling for rotating generation, transformers, legacy hotel-load equipment, and the utility grid. DC becomes compelling inside the AI factory because GPUs, batteries, PV, fuel cells, and power electronics are all fundamentally DC-linked, and because 800 VDC lowers current enough to make MW-class racks physically plausible.

The strategic design is therefore: medium-voltage collection and interconnection; modular conversion near data halls; an 800 VDC compute backbone; layered DC-coupled storage; cooling moved onto the bus where practical; inverter-fed AC only for legacy and life-safety loads and grid export/import; and a control system that makes the site behave as a stable, flexible, utility-compatible resource. Two things must still be settled by the industry rather than by any single operator — the 800 VDC convention and the certification path — which is the work the Direct Power Alliance exists to convene.

Appendix A. Planning formulas

Current at 800 VDC: $I = P / V$. Examples: 250 kW / 800 V = 313 A; 1 MW / 800 V = 1,250 A; 10 MW / 800 V = 12,500 A.

Two-wire DC conductor loss: $P\_loss = I^2 \cdot R\_loop$, where $R\_loop = 2 \cdot \rho \cdot L / A$. With $P = V \cdot I$, the fractional loss is $f = 2 \cdot \rho \cdot L \cdot P / (A \cdot V^2)$. Solving for one-way length: $L = f \cdot A \cdot V^2 / (2 \cdot \rho \cdot P)$. At 800 VDC with copper at 20 °C, $L(m) = 18{,}560 \cdot f \cdot A(mm^2) / P(kW)$. Implication: a 1 MW rack can be served at 800 VDC over data-hall distances, but a 10 MW block should not be moved long distances at 800 V unless conductor and busway costs are explicitly justified.

References

[1] NVIDIA Technical Blog, “NVIDIA 800 VDC Architecture Will Power the Next Generation of AI Factories,” May 20, 2025. developer.nvidia.com/blog/nvidia-800-v-hvdc-architecture-will-power-the-next-generation-of-ai-factories/

[2] NVIDIA Technical Blog, “Building the 800 VDC Ecosystem for Efficient, Scalable AI Factories,” Oct. 13, 2025.

[3] Flex, “Flex Launches 800 VDC Power Rack for Next-Generation NVIDIA AI Infrastructure,” Mar. 16, 2026.

[4] NVIDIA Technical Blog, “Designing Production-Ready Battery Energy Storage Systems for AI Factories,” June 10, 2026.

[5] NVIDIA Technical Blog, “How New GB300 NVL72 Features Provide Steady Power for AI,” July 28, 2025.

[6] IEEE 1547-2018, “Standard for Interconnection and Interoperability of Distributed Energy Resources with Associated Electric Power Systems Interfaces.”

[7] UL 1741, “Inverters, Converters, Controllers and Interconnection System Equipment for Use With Distributed Energy Resources.”

[8] FERC, “Interconnection of Large Loads to the Interstate Transmission System,” Docket No. RM26-4-000.

[9] FERC, “FERC Launches Aggressive Targeted Action to Speed Large Load Integration,” June 18, 2026 (directing the six RTOs/ISOs).

[10] E. Choukse et al., “Power Stabilization for AI Training Datacenters,” arXiv:2508.14318 (NVIDIA, Microsoft, OpenAI), 2025.

[11] National Bureau of Standards, “Copper Wire Tables,” NBS Handbook 100.
