Life Is On | Schneider Electric

5 Principles for 800 VDC in AI Data Centers: Rack-level Architectures as the Immediate Enabler

White Paper 213

Version 1

Data Center Research & Strategy

by Wendy Torell
Maria A. Torres Arango
Jim Simonelli

Executive summary

The demand for increased AI workload performance is pushing rack power densities beyond the practical limits of conventional power distribution. The electrical infrastructure supporting these AI data centers must evolve to deliver stable power safely, efficiently, and at scale. The industry has increasingly aligned around 800 VDC for these next-generation AI rack densities. In this paper, we identify 800 VDC power racks – sometimes referred to as 'sidecars' – as the immediate enabler for this shift. To support this transition, we propose five guiding principles that allow data center designers and operators to evaluate and adopt 800 VDC architectures tailored to their specific facility and business needs.

RATE THIS PAPER ★★★★★

Key takeaways

1. **Rising rack density forces a fundamental shift in power architecture.** AI-driven rack densities beyond ~400 kW exceed the practical limits of traditional AC and 48 VDC designs, requiring a new approach to power delivery. Relocating power conversion *out of the IT rack* and increasing the voltage to 800 VDC reduces congestion, frees IT space, increases power efficiency, and overcomes the power limitations of lower voltage.

2. **Industry alignment around 800 VDC is driven by physics, standards, and supply chain readiness.** Experience from adjacent industries such as electric vehicle (EV) charging, combined with evolving codes and standards, has positioned 800 VDC as a globally-scalable voltage.

3. **800 VDC rack-level conversion is the most feasible near-term approach and represents lower risk.** Using a power rack (adjacent to the IT rack) limits disruption to existing AC-based facility systems. It also reduces the impact of early-lifecycle failures and enables faster qualification and commissioning, accelerating deployment.

4. **Growing industry maturity around 800 VDC will ultimately enable the safe and reliable deployment of more centralized architectures.** The location of AC to 800 VDC conversion will evolve beyond the rack to upstream locations, driven by three conditions: evolution of industry standards, broader supply chains, and increased operational experience. These future architectures may offer benefits in efficiency, space utilization, or system integration, but they depend on higher levels of ecosystem readiness.

5. **Protection, grounding, and energy storage must be designed as a holistic system**. DC fault behavior, grounding strategy, capacitor-driven fault currents, and energy-storage placement interact in ways that require coordinated, end-to-end design rather than isolated decisions. Organizations like Open Compute Project (OCP), Current/OS, and Open Direct Current Alliance (ODCA) are helping with holistic design through an open ecosystem.

6. **Successful 800 VDC deployment depends as much on execution and operations as on design.** Workforce readiness, supply-chain maturity, validated procedures, and clearly defined maintenance and safety practices are critical to deploying and operating 800 VDC systems at scale.

Introduction

The demand for increased AI workload performance is driving a steady increase in data center rack power density. A key to increasing this performance is to reduce the chip-to-chip communication latency by increasing the number of chips (e.g., GPUs) in a single rack. This trend concentrates more compute – and more power – into a single rack. For a more detailed explanation of this trend, see Schneider Electric White Paper 110, *How 6 AI Attributes Change Data Center Design*.

Power can be delivered to the rack using several distribution approaches, and rack power density is a primary factor in determining which approaches remain practical. Table 1 summarizes typical rack density thresholds for two commercially available server power distribution approaches.

Table 1

Practical rack density thresholds for traditional and 48 VDC distribution, driving adoption of 800 VDC for densities greater than 400 kW.

Note, North America voltages are shown, but the same applies to global voltages.

Server power distribution approaches for AI workloads	Typical density threshold (kW/rack)
Traditional server PSUs **415/240 VAC feeds *PSUs inside servers* that convert 240 VAC to 12 VDC**	**170***
Open rack[1] **480 VAC feeds consolidated *power shelves within the IT rack* that convert 480 VAC to 48 VDC busbar**	**400****

* Assumes (3) 100A rPDUs, de-rated by 20% per code, with no redundancy and no energy storage, supporting 100% liquid-cooled IT (rounded from 173 kW)

** Assumes (4) busbar sections, each with 110 kW rated power shelves, with no redundancy and no energy storage, supporting 100% liquid-cooled IT (rounded from 440 kW)

As rack density increases, power delivery hardware begins to compete directly with IT equipment for space, airflow, and serviceability. In traditional distribution, servers are natively AC-powered, with conversion from AC to 12 VDC occurring inside each server. Beyond ~170 kW, limitations arise not from a single factor, but from the combined effects of:

- congestion from large rack PDUs, power connectors, and liquid-cooling manifolds inside the IT rack
- space consumed by power supplies inside every server

Open-rack power designs address these two limitations by consolidating traditional server power supplies into shared power shelves, converting from 480 VAC to a low-voltage DC busbar within the rack. The rack PDUs are replaced with a busbar (historically 12 VDC, now 48 VDC) that feeds the individual servers. This approach improves volumetric efficiency and simplifies cabling relative to traditional AC server PSUs. However, as rack densities rise over ~400 kW, the following limitations emerge:

- congestion from feeds & connectors supplying the AC power to the power shelves inside the IT rack
- space consumed by power shelves (and batteries) constraining the available volume for IT equipment
- low voltage (48-54 VDC) limiting the practical power capability of the busbar

Beyond ~400 kW, incremental improvements such as higher-capacity power shelves or higher-ampacity liquid-cooled busbars yield diminishing returns due to a lack of space for more IT equipment. Meaningful increases in rack density therefore require two changes: (1) moving AC-to-DC conversion out of the IT rack and (2) delivering power at a higher voltage. Together, these shifts enable higher power levels while minimizing conductor size. In essence, this creates a new server power distribution approach. While the IT industry defines the trajectory of compute, the electrical system must evolve to deliver that power efficiently, safely, and at scale.

In response to these constraints, **the industry is coalescing around 800 VDC as a practical approach to support rack densities from 400 kW to well over 1 MW**. This alignment reflects a combination of electrical fundamentals, experience from the electric vehicle (EV) industry,[2] and the need for a solution that can be manufactured, deployed, and supported globally.

[1] Open Compute Project, *What is OCP and Who is it for?*

[2] N. Srivastava, H. Petty, *Building the 800 VDC Ecosystem for Efficient, Scalable AI Factories*

In this paper, we identify 800 VDC power racks – sometimes referred to as 'side-cars' – as the immediate enabler for this shift. To support this transition, we introduce five guiding principles (Figure 1) to help data center designers and operators evaluate and adopt 800 VDC architectures that align with their facility constraints, operational practices, and business objectives.

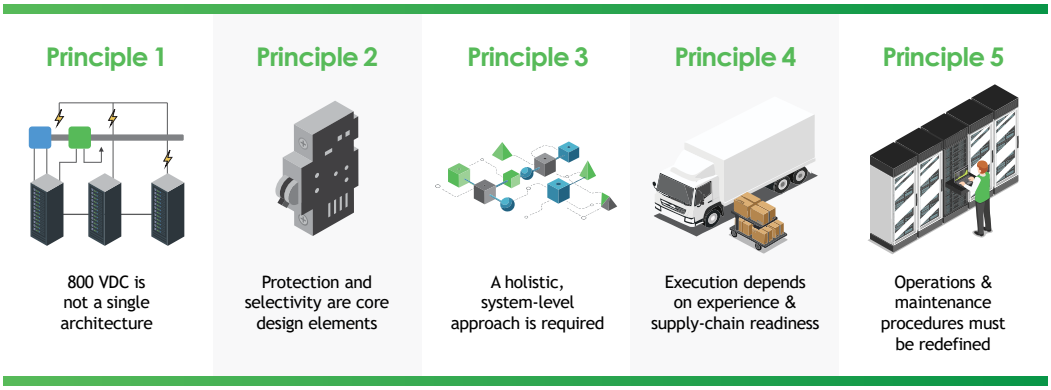

Figure 1

Five guiding principles for 800 VDC deployments

It's important to note that the topics covered in principles 2 through 5 can also apply to AC distribution data centers. They're included in this paper to provide comprehensive guidance on the shift to 800 VDC.

Principle 1: 800 VDC is not a single architecture

800 VDC should *not* be interpreted as a single, fixed architecture. Differences in workload composition, facility constraints, resiliency objectives, and regulatory environments naturally lead to different design choices. As a result, multiple 800 VDC architectures will coexist, based on the following **five key architectural attributes**:

- the location of conversion from AC to 800 VDC
- the isolation and grounding of the DC supply
- the form of 800 VDC
- the implementation of resiliency (redundancy, concurrent maintenance)
- the placement of energy storage, if present, within the power path

Together, these attributes create a range of valid architectures rather than a single reference design.

The location of conversion from AC to 800 VDC

The selection of the conversion location is pivotal to the architectural decision, as it influences the other four attributes. Conversion from AC to 800 VDC can take place at a spectrum of locations in the electrical system, as Figure 2 illustrates.

Given the space constraints, we are left with three logical locations for converting to 800 VDC to achieve higher density. Each option presents different tradeoffs between several factors: efficiency, flexibility, fault containment, readiness to scale, and operational complexity.

800 VDC power racks are the most feasible near-term approach today and represent lower risk to data center operators.

Figure 2

Spectrum of locations for AC to DC conversion

Disaggregating the power into a separate power rack is the most feasible near-term approach today and presents a lower risk to data center operators.

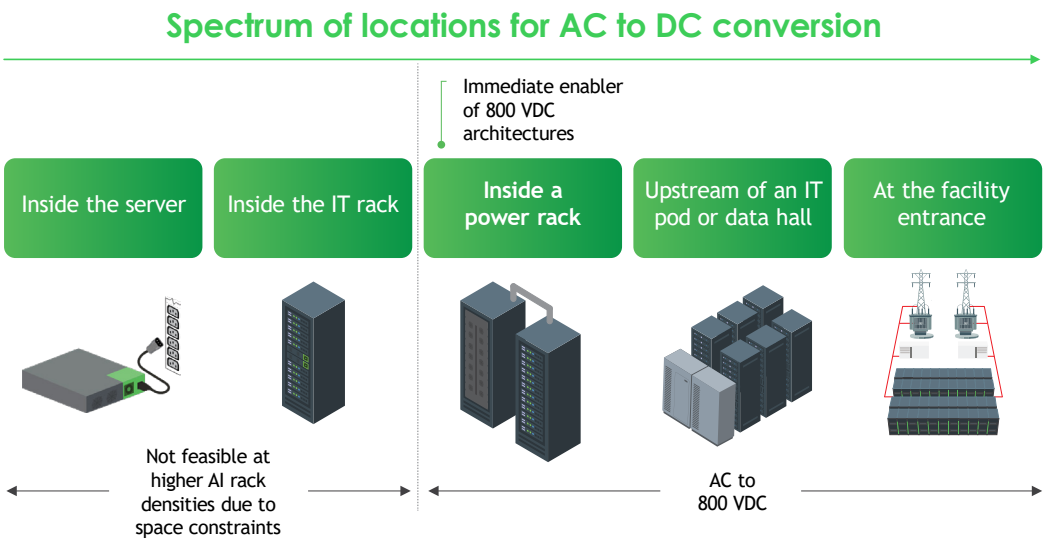

A power rack[3], is a dedicated power module located adjacent to, but outside of, the IT rack. The power rack houses power conversion, protection, and optional energy storage that are impractical to place within the IT rack at very high-power densities. This approach reduces risk because:

- **It minimizes disruption to the existing AC-based data center ecosystem.** Power racks introduce the smallest architectural change relative to today's AC-based data centers. They avoid re-designing upstream switchgear, protection schemes, grounding strategies, and operational workflows that are well understood. For example, a standard centralized AC UPS can still be used rather than changing the electrical design for a DC UPS.
- **There is a more established supply chain.** The availability of key components required for rack-level 800 VDC systems (such as power electronics, connectors, protection devices, and controls), benefits from the existing EV market. This market offers component availability and manufacturing scale at the power and voltage levels required for rack-scale systems, reducing procurement uncertainty.
- **It reduces the failure zone to a single IT rack.** Introducing a new type of distribution system may lead to failures early in its lifecycle. Specific examples might include installation wiring errors, protection mis-coordination, errors in procedures, and human error. Converting at the rack level significantly reduces the risk of downtime at the pod, data hall (IT room), or facility level.
- **It shortens qualification and commissioning time.** Standardized power racks can be tested, validated, and commissioned as discrete units. This replicable approach allows for accelerated learning, enabling faster installation, which lowers the risk of project delays.

While centralized conversion at the pod, data hall, and facility level offers promise, it does not satisfy one or more of these immediate requirements, particularly around supply chain readiness and risk containment. **In contrast, conversion at the rack level provides a clear, actionable path for 800 VDC implementation today.**

Because the industry currently lacks a critical mass of electricians and operators experienced in higher voltage DC distribution systems, the power rack's modularity is its greatest asset. Power racks enable a "step-and-repeat" learning process that standardizes installation, safety training, and maintenance procedures on a small,

[3] J. Huntington, M. Tu, *800 VDC Architecture for Next-Generation AI Infrastructure*

repeatable scale. This modularity not only accelerates deployment but also creates a logical integration point for power-smoothing capacitors and back-up batteries.

Selecting a rack-based architecture today is a pragmatic choice, not a critique of more centralized designs. We expect pod, data hall, and facility-level architectures to become viable as the industry matures. This includes evolved standards, proven protection strategies, improved component availability, and growth in operational experience. These more centralized architectures avoid occupying valuable IT space used by power racks. Future papers will explore this in greater detail.

The isolation and grounding of the DC supply

Isolation is a way to electrically and physically separate two circuits such that current on one circuit cannot flow to the other circuit (e.g., fault current, electrical noise). For example, a transformer performs this function with two physically separated coils. A circuit is either isolated or not. Grounding is the connection of a conductor to a ground reference (e.g., earth ground) for safety (primarily) and equipment operation. A data center's electric supply grounding is governed by local electrical codes (for safety) and design preferences (for critical load reliability). These two concepts are no different for 800 VDC than for AC distribution in data centers today.

Both AC and DC distribution systems use several common grounding philosophies. These influence fault behavior, protection requirements, and operational reliability. Electrical codes prioritize human safety in *all cases*, but the choice of grounding method does affect how the system responds to faults and the risk of downtime.

- Solidly grounded – One of the conductors is grounded. In the event of a fault to ground, a circuit protection device (e.g., circuit breaker, fuse) opens immediately, potentially causing downtime of the critical load. While this is a simpler and lower-cost approach, it doesn't offer the operational reliability of the next two.
- High-resistance grounded (HRG) – One of the conductors is connected to ground via a resistor. In case of a ground fault, the high-resistance path allows a small amount of current to flow to ground, but not enough that it opens the circuit protection device. While the critical load continues to operate under a fault condition, it must be resolved to avoid the possibility of downtime due to a second fault. Regulations always require a ground-fault detection system[4] for this approach so that the user can be alerted and perform maintenance to find and fix the fault. The ground fault detection system itself does not tell you where the fault is, but you can add optional sensors at every feeder or branch to identify the fault location. Operational reliability can be high with this approach but is more costly than solidly grounded.
- Floating (i.e., ungrounded) – The return conductor is not grounded. In the event of a fault, the protection device doesn't open because there is no return path for the current. Like HRG, operational reliability can be high, and a ground-fault detection system[4] is always required, since the fault cannot be detected otherwise. While the critical load continues to operate under a fault condition, it must be resolved to avoid the possibility of downtime due to a second fault. As with HRG, you can add optional sensors at every feeder or branch to identify the fault location. This is the least common approach primarily because the first fault risks harming the IT equipment.

[4] For HRG systems, ground-fault detection is done through monitoring leakage current; For floating systems, it is done through monitoring insulation degradation.

The optimal isolation and grounding approach depends on system objectives, protection and selectivity strategy, operational practices, and applicable codes and standards. **As with AC distribution, these approaches should be evaluated as part of the overall system architecture rather than as an isolated design decision**. More detail on DC grounding methods and their implications is provided in Schneider Electric White Paper 47, *Necessary Considerations for Designing DC Architectures in Data Centers*.

The form of 800 VDC

In the context of 800 VDC, high-density AI equipment can be specified to support one of two voltages: 800 VDC or 400 VDC. This specification can drive the form of 800 VDC supply used – either **differential (+800 VDC)** or **bipolar (±400 VDC)**. The choice affects insulation requirements, protection strategy, grounding approach, and component selection.[5]

- **Differential (+800 VDC, 2-wire) –** This form is suitable for loads requiring a full 800 VDC input. Power is supplied to the load via two-wires, which simplifies cabling and protection. Because the load draws its energy differentially – measuring only the potential difference between the two supply conductors – the earthing scheme does not impact power delivery. However, different earthing choices can be implemented: a mid-point earth scheme limits the potential to 400 VDC relative to ground, whereas earthing one of the current carrying conductors requires insulation and clearances to be sized for 800 VDC.
- **Bipolar (±400 VDC, 3 wire) –** This form offers increased flexibility by supporting both 400 VDC and 800 VDC loads. It utilizes a positive (+400 VDC) conductor, a negative (-400 VDC) conductor, and a midpoint common current-carrying conductor (which typically sees zero current). However, in the case of load imbalances, it carries the differential current, analogous to a neutral conductor in an AC system. This approach has the drawback of adding more current-carrying conductors where space is premium.

The data center industry is deliberately keeping both approaches on the table for rack-level 800 VDC systems. Selection is driven by ecosystem alignment, safety philosophy, available standards and certifications, and the degree of architectural complexity an operator can manage. Examples of both implementations exist in the industry today, including differential designs used in NVIDIA-led platforms[6] and bipolar implementations, such as Open Compute Project's (OCP) high power rack V4 'Diablo' design[7]. We anticipate that experience from both will eventually lead to a unified approach, along with data center-specific standards and more readily available components.

The implementation of resiliency

Just like AC systems, resiliency in 800 VDC architectures can be implemented at multiple levels, including redundancy in conversion equipment, distribution paths, and energy storage. Choices range from centralized redundancy serving large load blocks to more localized redundancy closer to the rack. Each has different implications for availability and fault containment.

Higher resiliency increases equipment count, system complexity, and operational coordination. The appropriate level and placement of redundancy depend on

[5] Open DC Alliance, *VDE SPEC 90037*

[6] M. Blake, et al., *800 VDC Architecture Will Power the Next Generation of AI Factories*

[7] OCP, *Diablo 400 Project: Rack and Power*

availability targets, acceptable fault domains, and the ability to maintain or service components without disrupting IT operations.

The placement of energy storage

Energy storage serves four main functions in an AI data center:

- Back-up power – to maintain operation during long-duration upstream power failures. This is done to either keep the critical loads running (hours/days) or to provide enough time to properly shut down servers and store data (seconds/minutes).
- Smoothing dynamic load profiles of large AI factories – to prevent multi-MW load swings from adversely affecting power sources. AI workloads, like training, perform synchronized operations (i.e., orchestration) that can produce rapid power spikes (up to ~150% of nominal load) as well as deeper sustained drops (down to ~40%).[8] Different operations call upon different IT components (e.g., GPUs, memory chips, network cards) at different times, thereby creating a dynamic power consumption profile. This profile varies in magnitude (peak height), duration (how long a peak is held), and frequency (peaks per second). The smoothing function stabilizes the power profile seen by generators, UPS systems, and the grid. It reduces mechanical and electrical stress and avoids frequency and voltage instability caused by sudden load changes.[9] Smoothing requires a two-tiered approach in which each energy storage technology addresses peaks and valleys of varying durations within the power profile. For example, one type of energy storage absorbs the short-duration spikes above 100% and fills in rapid dips, while another addresses the longer duration swings.
- Grid fault ride through (FRT) – to keep the data center electrically connected to the grid during short-duration disturbances (e.g., fault). FRT prevents an instantaneous drop in grid voltage from cascading to a loss of data center load which can destabilize the grid for the entire region. Without FRT, a short grid disturbance (<150 ms) can cause a large data center (e.g., 1 GW) to transfer immediately to onsite sources. If the facility does not return to utility supply shortly after the disturbance restores, the grid will see a loss of 1 GW of load. This rapid decrease in significant load can drive temporary overvoltage and over-frequency conditions. For example, over-frequency relays can trip grid generators offline. These interactions can cascade into grid blackouts. FRT energy storage temporarily "buffers" the power imbalance during the disturbance, maintaining electrical coupling to the grid and enabling a controlled transition back to utility supply once the fault clears. This function becomes increasingly important as AI deployments scale into the hundreds of MWs.[10]
- Ramp up / ramp down – to control the sudden start and stop of large multi-MW AI workloads, preventing step changes in grid demand. Similar to FRT, the energy storage gradually supplies or absorbs energy during these transitions, reducing stress on the grid and upstream electrical infrastructure. These controlled ramps occur over seconds and require sizable storage capacity (MWh range). This is specific to AI training workloads and often not required for other workloads (e.g., web hosting, AI inference).

Energy storage technologies for an 800 VDC architecture can be centralized at the facility level, distributed at the data hall or pod level, located adjacent to the rack, or within the server. The optimum strategy is driven by the IT workload, the required function(s) for energy storage, availability targets, and operational preferences.

[8] This series of IT operations is repeated until an LLM is trained, lasting from days to months.

[9] Schneider Electric White Paper 126, *Retrofitting Existing Power Systems for AI Clusters*

[10] A recent example of this issue occurred in Northen Virginia nearly resulting in a blackout.

Table 2 provides typical locations and technologies used for energy storage depending on function.

Table 2

Descriptions of energy storage location and technologies

Energy storage function	Location and technologies
Back-up power	If runtime and availability requirements are uniform for all loads, it makes sense to keep energy storage at the **facility**, **data hall**, or **pod level** to minimize complexity and maintenance. It also avoids occupying IT space. Batteries are typically used for this function, but the technology varies depending on the runtime requirements. Runtime could range from seconds to hours. However, if availability and/or runtime requirements are not uniform, the logical placement is at the **rack level**, if space permits. Schneider Electric White Paper 229, *Battery Technology for Data Centers: VRLA vs. Li-ion*, provides more detail.
Smoothing dynamic load profiles of large AI factories	Two energy storage technologies and locations combine to filter the varying peaks and valleys. The first addresses the higher-frequency, shorter-duration (<50 ms) fluctuations that can dip to ~80% or spike to ~150% of nominal load. These fluctuations are best managed at the **rack level** (e.g., integrated into the power shelf). Electrical double-layer capacitors (EDLC)[11], or other supercapacitors[12], are typically used due to their very high power density and rapid response. After these capacitors remove the spikes above 100%, a second tier manages longer swings from 40% to 100% of nominal load lasting hundreds of milliseconds to seconds. This function is typically handled at the **pod or data hall level** using high-cycling batteries[13] or larger capacitor systems. Managing this layer of fluctuations centrally smooths the aggregate AI load profile seen at the grid interconnection point.
Grid fault ride through (FRT)	The technology for this function depends on whether you have critical load that requires continuous power (no break during power events). If you *don't* have critical load – and the entire data center can go down for a few seconds – a **facility-level** battery energy storage system (BESS) can meet the grid's FRT requirement. However, it must be sized to support the majority[14] of the total data center load, including IT, chillers, etc. If you *do* have critical load, at minimum, you need a power back-up system that is voltage and frequency independent[15] (VFI). This meets the grid's FRT requirements while maintaining continuous voltage to the critical load. VFI is what allows the energy storage solution to simultaneously support the critical load while also drawing power from the grid. Examples of a VFI back-up system are double-conversion UPSs and rack-level power shelves with battery back-up units (BBUs). As noted above, the FRT system(s) must be designed to support the majority of the facility loads (IT and auxiliary). This can be achieved by supporting all loads on the same **facility-level** VFI UPS system. Alternatively, you could use a VFI UPS or **rack-level** FRT solution (e.g., in the power racks) to support critical IT, and a separate **facility-level** FRT-compliant storage system (BESS or UPS) to support auxiliary loads.
Ramp up / ramp down	Generally leveraging the same energy storage system as FRT, the ramp-up / ramp down function is typically implemented at the **facility level** to manage the power profile at the grid-connection point during the sudden starts or stops of large multi-MW AI workloads.

[11] TDK, *Capacitors, Part 8 "Electric Double Layer Capacitors (EDLC)"*

[12] Wikipedia, Supercapacitor explanation

[13] Schneider Electric White Paper 284, *Considerations for Selecting a Lithium-ion Battery System for UPSs and Energy Storage Systems*

[14] Exact amount depends on local utility requirements. For example, ERCOT specifies "at least 90% of its pre-disturbance consumption level from the grid".

[15] VFI means that the back-up system can supply voltage and current independent of the grid supply. See Schneider Electric White Paper 1, *The Different Types of UPS Systems*, for more on VFI.

Principle 1 summary: It's not just about technology, but how preferences & constraints drive the appropriate architecture

As with AC data centers, the appropriate 800 VDC architecture for a given data center depends on several interrelated preferences and constraints, including:

- **Availability and resiliency level**, which affects component redundancy, concurrent maintainability, and acceptable fault domains (i.e., failure zone)
- **Cost**, which is often a tradeoff between capital expenses and operational expenses (e.g., maintenance, energy, downtime, etc.)
- **Greenfield versus brownfield**, which influences available space, pathways, and protection schemes
- **Operational and maintenance practices**, such as energized-work policies, staff training, and safety procedures
- **Physical constraints**, such as space, weight, routing, and cooling capacity
- **Protection, selectivity, and fault-containment**, which shape how power and protection are distributed
- **Regional codes, standards, and authorities having jurisdiction (AHJ) expectations**, which influence acceptable configurations and maintenance approaches
- **Sustainability targets**, such as energy efficiency, embodied carbon, choice of onsite energy generation, etc., that align with corporate commitments
- **Target IT rack power density and growth trajectory**, which drives conversion location
- **Time-to-deploy**, driven by schedule sensitivity and supply-chain readiness
- **Workload characteristics**, such as whether loads are homogeneous or mixed-use environments

Recognizing that 800 VDC is not a single architecture – but a range of architectures built around specific choices – is critical to meeting deployment objectives. Today, rack-level conversion is the most immediate and widely accessible enabler of 800 VDC, allowing adoption without fundamentally re-designing upstream power systems. However, conversion location and energy storage placement have cascading impacts on protection strategy, operations, and scalability. By understanding these tradeoffs, designers and operators can align 800 VDC architectures with their business requirements rather than forcing a one-size-fits-all solution.

The remaining four principles build on this foundation by examining how protection, holistic design approach, execution readiness, and operations influence the success of 800 VDC adoption.

Principle 2: Protection and selectivity are core design elements

Protection and selectivity strategies commonly used in conventional AC data center architectures are also applicable for 800 VDC. **Protection** refers to the methods used to reduce the risk of personnel injury or system damage due to electrical hazards. These hazards include arc flash, electric shock, ground faults, and overcurrent conditions. Protection is primarily achieved through overcurrent, short-circuit, and ground-fault protection functions. **Selectivity**, also referred to as selective coordination, is the ability of a power system to isolate only the portion of the system affected by a fault while keeping the remainder of the system energized.[16]

[16] If a short circuit occurs inside an IT rack, the circuit breaker closest to the rack should open first, as opposed to opening the circuit breaker further upstream and taking down a larger number of racks.

Zero crossing

In AC systems, the sinewave periodically passes through a zero crossing, where voltage and current momentarily reach zero. While there are many other factors at play, these zero crossings assist fault interruption because the arc that forms between a circuit breaker's open contacts, loses the source energy required to remain conductive once the current reaches zero.

DC systems deliver a constant voltage and therefore do not have a zero crossing to assist with fault interruption.

Ultimately, the longer it takes to interrupt the fault, the more energy is allowed to feed the fault, increasing the hazard.

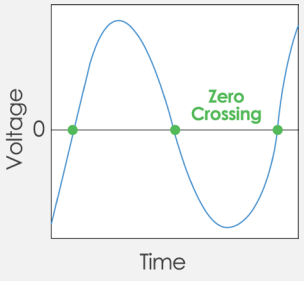

800 VDC architectures typically include bulk stored energy sources like batteries and capacitors. These represent a significant amount of energy that could be released into the data center infrastructure (e.g., arc flash). However, 800 VDC converters are typically placed in between bulk energy storage and the DC bus. These converters provide the necessary voltage regulation while also preventing this stored energy from feeding downstream faults. Specifically, the converter actively limits output current and eventually shuts down under fault conditions.[17]

Despite their protection benefits, electrical codes do not consider converters to be branch circuit protection. Therefore, most, if not all, electrical codes require the use of DC protection devices external to converters. Fuses can provide very fast, current-limiting interruption that reduces let-through energy, while circuit breakers offer reclosability, operational flexibility, and integration with monitoring and control systems. Importantly, DC circuit breakers require a **higher arc-extinction performance** than AC circuit breakers, due in part to the lack of a zero crossing (see **Sidebar "Zero crossing"**). This means they must extinguish longer-lasting, higher-energy DC arcs. This is achieved by connecting poles in series and through features like greater contact spacing, faster contact separation, higher insulation strength, magnetic blowout,[18] and enhanced arc chutes that lengthen and cool the arc.

While selecting and coordinating protection devices is a critical aspect of 800 VDC system design, **rack-level architectures simplify selectivity by shortening fault current paths and limiting the number of series-connected protective devices**. Effective protection in 800 VDC systems is typically achieved through coordinated use of multiple device types rather than reliance on a single technology.

In a closed ecosystem[19], standard over-current protection devices such as air circuit breakers (ACBs) and molded case circuit breakers (MCCBs) are "tuned" against a defined rectifier fault current profile. In contrast, in an open ecosystem, enhanced flexibility is necessary because custom configuration of individual components is not always available or allowed. Emerging solid-state circuit breakers and hybrid circuit breakers provide additional choices by enabling faster fault interruption than traditional circuit breakers or fuses. However, these devices introduce trade-offs, including higher cost, lower product availability, and integration challenges.

Industry collaboration is helping advance best practices in this area. Organizations such as UL, IEC, OCP, Open Direct Current Alliance (ODCA), and Current/OS are key forums for these discussions. Schneider Electric is also contributing to the development of reference designs, testing approaches, and guidance that support improved protection, coordination, and safety in high-power DC environments.

Principle 3: A holistic, system-level approach is required

Successfully deploying 800 VDC, including rack-level architectures, requires more than selecting individual components or subsystems. Power distribution, conversion, protection, cooling, monitoring, and controls must be designed and validated as a holistic system. Decisions made in one part of the architecture can have cascading effects on performance, safety, and operability elsewhere.

We provide some key examples showing how a holistic, system-level design avoids risks in operations, efficiency, and cost, that are otherwise hidden until deployment.

[17] Converters can contribute to faults through internal components like capacitors. While this energy is significantly less than bulk energy storage, it must be addressed depending on your system design.

[18] Schneider Electric Cahier Technique 154, *Low Voltage Breaking Techniques*

[19] The system is entirely designed and "tuned" by a single vendor (e.g., the rectifier is designed and built for a specific 800 VDC power rack vs. systems that could have a wide range of variables).

- **Protection and safety performance must be evaluated across the entire power path to avoid miscoordination and unintended fault behavior.** Selectivity, arc-flash mitigation, and fault containment are influenced by interactions among utility sources, energy storage, grounding schemes, and upstream and downstream protection devices. If these interactions are not evaluated end-to-end, actual fault currents and clearing times can differ from design assumptions. This leads to loss of selectivity or protection devices operating outside their intended fault domains. Even in rack-level DC architectures, upstream coordination mistakes could cut power to an entire pod of AI workloads.
- **Adding non-IT equipment such as emergent 800 VDC CDUs and fans to the IT feeder requires caution**. These mechanical loads could cause issues for IT loads when powered from the same feeder. For example, high DC inrush currents and transient step loads can cause voltage sags on the shared 800 VDC feeder. Just as with traditional AC systems, consider splitting the DC bus into non-IT loads and IT loads.
- **Fault path and grounding interaction studies must be conducted to prevent unintended tripping of upstream circuit breakers and loss of selectivity.** In mixed AC/DC architectures, fault current can return through grounding systems, shared neutrals, or power-electronic backfeed paths that are not obvious from single-line diagrams. If these paths are not intentionally designed and validated, downstream faults can cause upstream breakers to trip or expand the fault domain beyond what was intended, i.e., a wider failure zone.
- **Monitoring and control must account for operating mode changes to avoid miscoordination and delayed fault clearing.** Transitions between grid-connected operation, battery-supported operation, and grid-interactive modes change available fault current and acceptable protection behavior. For example, a downstream DC fault that is correctly isolated during grid operation may clear more slowly or trigger upstream devices when the system is operating on batteries. This could happen if protection settings and logic do not account for the reduced fault current. Without monitoring the system state (e.g., source of supply, circuit breaker status) and coordinated control of protection behavior, faults can result in unnecessary upstream trips or extended recovery times during normal operating transitions.

A system-level approach enables 800 VDC architectures to meet performance, availability, and safety objectives simultaneously. Treating the infrastructure as a coordinated whole – rather than a collection of independent elements – reduces late-stage redesign and operational risk. Organizations like OCP, Current/OS, and ODCA are helping to enable holistic design through an open ecosystem. With these system-level considerations addressed, the focus can shift to execution.

Principle 4: Execution depends on experience & supply-chain readiness

Deploying 800 VDC is not only a design challenge, but an execution challenge. Successful implementation depends on practical experience with DC systems and access to a supply chain capable of delivering, installing, commissioning, and supporting higher-voltage DC infrastructure at scale.

Compared to conventional AC systems, 800 VDC systems are not well understood by many stakeholders involved in data center deployment and operations. Fewer electricians, technicians, and service organizations have direct experience with installing and commissioning higher-voltage DC systems.

In parallel, industry standards, tools, and accepted methodologies for assessing DC-specific hazards are less mature and less consistently applied than their AC counterparts. These factors increase the importance of validated reference

designs, proven installation practices, and clearly defined operating procedures to reduce the risk of injury and downtime.

Supply-chain readiness is equally critical. While the availability of individual DC-rated components like circuit breakers continues to improve, standardization across complete 800 VDC systems remains limited. Lead times, regional availability, and qualification status can vary significantly, particularly as deployments move beyond pilot scale. Specifying power equipment, protection devices, controls, and services across multiple suppliers adds complexity, increases coordination effort, and can introduce schedule and commissioning risk if not managed well.

Execution experience helps mitigate these challenges. Organizations with prior DC deployment experience are better positioned to anticipate integration issues, align designs with evolving standards and regulatory expectations, and support safe installation and commissioning. As adoption accelerates, a well-designed system and practical experience become the determining factors in whether 800 VDC architectures can be deployed predictably and at scale.

Principle 5: Operations & maintenance procedures must be redefined

Operating and maintaining 800 VDC systems differs fundamentally from working with conventional AC-powered racks. Procedures, training, and safety practices that are appropriate for lower-voltage AC environments do not directly translate to higher-voltage DC systems. While higher-voltage DC maintenance practices exist for the EV and solar industries, they are still evolving for the data center industry. Staff must be qualified with the skills and knowledge required for the operation and maintenance of the equipment. Work practices must be re-evaluated to identify electrical hazards and reduce associated risks.

At higher voltage levels, physical access, approach distance boundaries for arc flash, and serviceability become primary operational design considerations. For example, maintenance activities that require access behind or adjacent to 800 VDC power racks must be planned to avoid unnecessary power interruptions and exposure to energized components.

In some 800 VDC configurations, power racks or subassemblies may need to be fully de-energized for specific maintenance tasks. In the future, the design may support controlled **live-swap** or live-service procedures, allowing tasks associated with specific components – such as power shelves or conversion modules – to be completed without shutting down the entire rack. Whether energized work is permitted, and under what conditions, has direct implications on availability, maintenance windows, staffing requirements, and operational risk.

Live-swap capability[20] can be a meaningful operational differentiator, but it is not inherent to 800 VDC environments. Achieving live replacement of power components depends on system-level design choices, including modular power conversion, defined electrical boundaries, coordinated protection, interlocks, and validated service procedures. Designing these elements together reduces the maintenance impact on availability.

Safety practices and training requirements also change with 800 VDC. Higher-voltage DC may introduce shock and arc flash hazards with higher severities or likelihood of occurrence. Like AC systems, 800 VDC systems will require workers to assess the risk of electrical shock and arc flash, and implement risk control measures. These energy levels may require the use of personal protective equipment (PPE),

[20] An example of live swap is discussed in Schneider Electric White Paper 13, *Mitigating Electrical Risk While Swapping Energized Equipment*.

revised lockout/tagout procedures, and specialized training for both facilities personnel and IT administrators. Also, as higher voltage moves closer to the IT rack (where IT personnel work), clear delineation of responsibilities between facilities and IT teams becomes critical for safe and reliable operation. Schneider Electric White Paper 194, *Arc Flash Considerations for Data Center IT Space*, addresses AC voltage systems but offers valuable context for what to expect with 800 VDC systems.

Defining operations and maintenance strategies early in the design process helps ensure that 800 VDC systems can be serviced safely, efficiently, and predictably over their full lifecycle. Failing to address these considerations can constrain maintenance practices, reduce availability, or result in safety risks that are difficult or expensive to mitigate after commissioning.

Next steps

800 VDC represents a significant evolution in data center power delivery, driven by the need to support next-generation AI rack densities that can no longer be supported with traditional server power distribution. As we have shown in this paper, near-term adoption of 800 VDC is enabled with rack-level architectures, while broader deployment options will emerge as the industry matures.

These next steps provide a practical framework for adopting 800 VDC when and where AI workloads warrant the shift:

1. Clarify business and workload requirements early. Assess workload homogeneity, rack densities, availability objectives, timelines, cost constraints, and whether the project is greenfield or retrofit. These factors quickly focus appropriate 800 VDC designs.

2. Evaluate vendors for DC readiness and system-level expertise. Select partners with demonstrated experience in DC system design, installation, commissioning, and maintenance. Prioritize partners with proven supply-chain readiness and the ability to deliver integrated, system-focused solutions rather than isolated components. This can significantly reduce risk and facilitate a successful deployment.

3. Evaluate live-swap capability as an operational and availability advantage. Determine whether the ability to service or replace power components without fully de-energizing IT racks provides meaningful benefits for availability, maintenance windows, and operational flexibility.

4. Validate the architecture through end-to-end system studies. Assess power conversion, protection and selectivity, grounding, energy storage interaction, cooling integration, and operational modes. This will confirm that the chosen 800 VDC power-rack design meets performance, safety, and availability targets before deployment. Factory acceptance tests (FAT) are one way of validating that a sub-system works as intended.

5. Stay current on rapidly evolving power technologies in this space. Continuously monitor developments such as more centralized AC-to-DC conversion approaches at the pod, data hall, or facility level. Also look for potential shifts beyond 800 VDC (e.g., 1,500 VDC), solid-state transformers, and advanced DC protection technologies. These innovations can change viable power architectures and design tradeoffs.

About the authors

Wendy Torell is a Senior Research Analyst in Schneider Electric's Data Center Research & Strategy group with 30 years of data center experience. She analyzes the value of emerging technologies and translates industry trends into practical design and operational guidance. Wendy leads development of Schneider Electric's interactive, web-based TradeOff Tools, helping clients quantify decisions related to availability, sustainability, and cost. She holds a BS in Mechanical Engineering from Union College, an MBA from the University of Rhode Island, and is an ASQ Certified Reliability Engineer.

Dr. Maria A. Torres Arango is a Research Analyst in Schneider Electric's Data Center Research & Strategy group. She investigates technologies, materials, and infrastructure to guide data center strategies. She contributes to developing analytical TradeOff Tools to evaluate the impact of technology decisions in data centers. Maria's expertise includes materials design and optimization, as well as fundamental studies on materials synthesis processes using X-ray characterization at the National Synchrotron Light Source II, Brookhaven National Laboratory. Maria holds a BS in Aeronautical Engineering from Universidad Pontificia Bolivariana, Colombia; and a MSc in Aerospace Engineering, and a PhD in Materials Science and Engineering from West Virginia University.

Jim Simonelli is Senior Vice President & Chief Technology Officer for Schneider Electric's Data Center and Networks division and segment. In this role, he is responsible for deployment of the R&D strategy for hardware, software and service offers and for forward-looking technology exploration activities. Jim is a long-time veteran with Schneider Electric holding many leadership positions relating to technology and innovation in the Data Center critical power and cooling business arena. Jim holds a BS and MS in Power Electronics from Worcester Polytechnic Institute.

Acknowledgements

Special thanks to Victor Avelar for his significant contribution, and to Tony Landry, Fabricio Alves Borges, Stuart Sheehan, Gia Wiryawan, Himanshu Prasad, and Rodney West for their expertise and technical review of this white paper.

RATE THIS PAPER ★★★★★

Resources

The Different Types of UPS Systems
White Paper 1

Mitigating Electrical Risk While Swapping Energized Equipment
White Paper 13

Necessary Considerations for Designing DC Architectures in Data Centers
White Paper 47

How 6 AI Attributes Change Data Center Design
White Paper 110

Retrofitting Existing Power Systems for AI Clusters
White Paper 126

Arc Flash Considerations for Data Center IT Space
White Paper 194

Analysis of Data Center Architectures Supporting Open Compute Project (OCP)
White Paper 228

Battery Technology for Data Centers: VRLA vs. Li-ion
White Paper 229

Considerations for Selecting a Lithium-ion Battery System for UPSs and Energy Storage Systems
White Paper 284

Low Voltage Breaking Techniques
Cahier Technique 154

Browse all white papers and executive reports
Data Center Research & Strategy Papers

Browse all TradeOff Tools™
Data Center Research & Strategy Tools

Note: Internet links can become obsolete over time. The referenced links were available at the time this paper was written but may no longer be available now.

Contact us

For feedback and comments about the content of this white paper:

Schneider Electric Data Center Research & Strategy
dcsc@se.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm