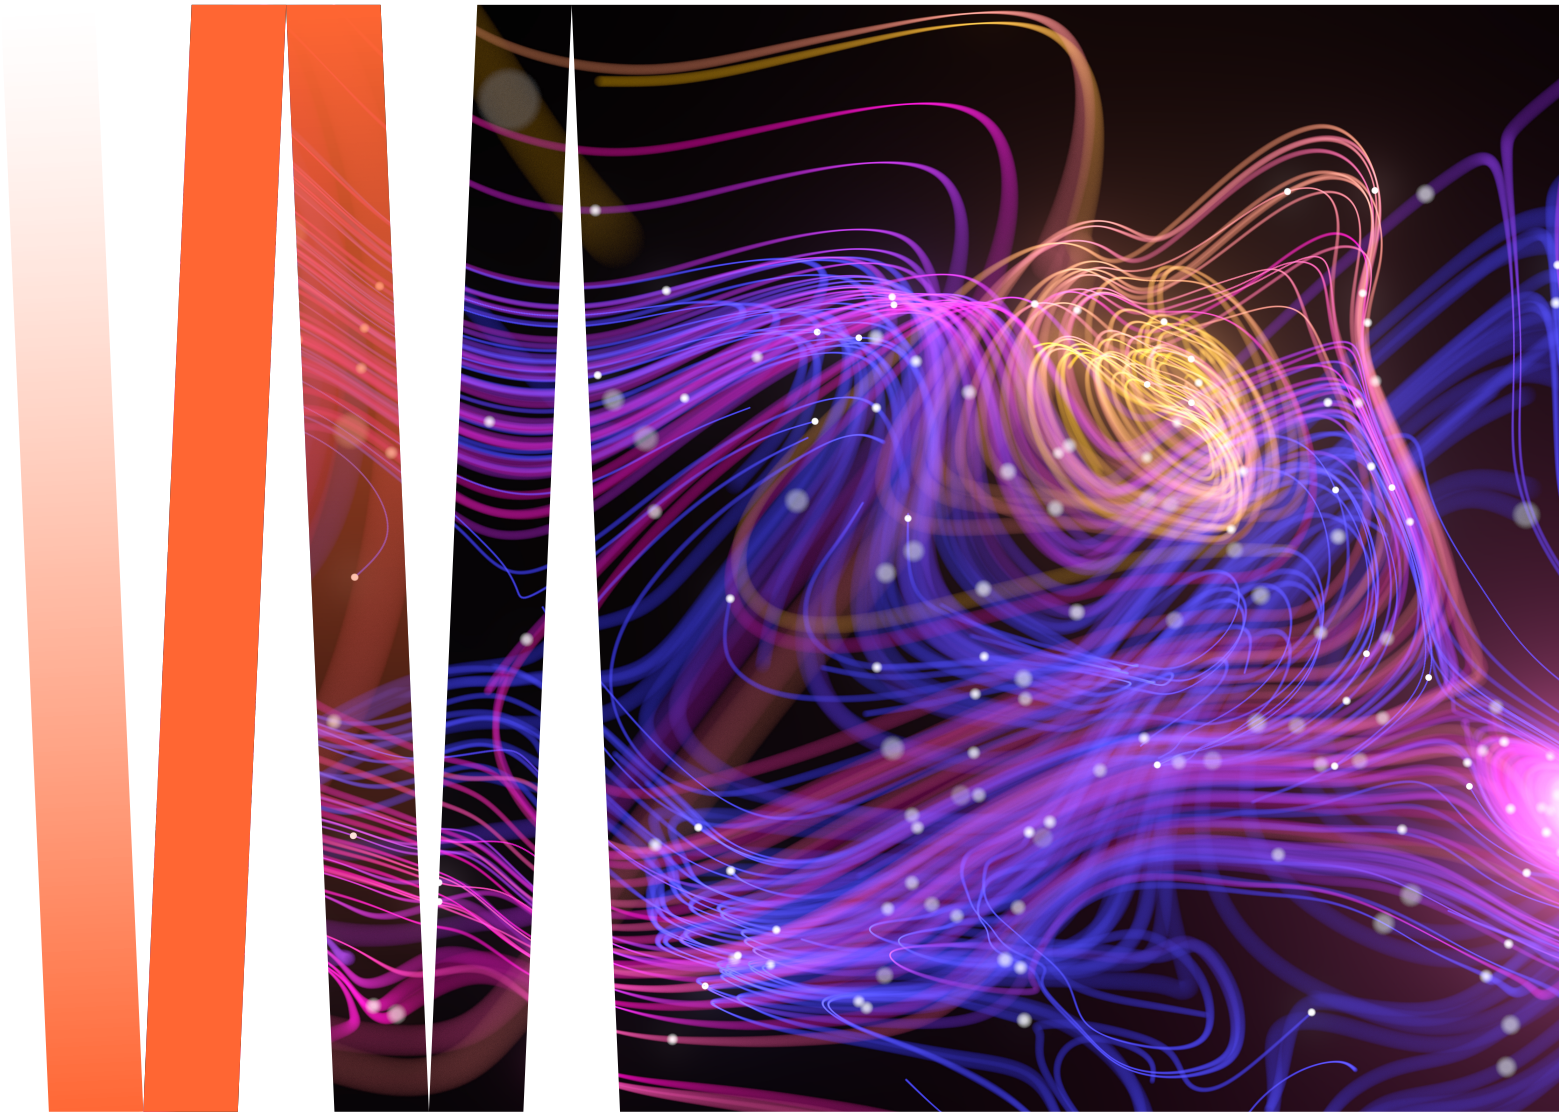
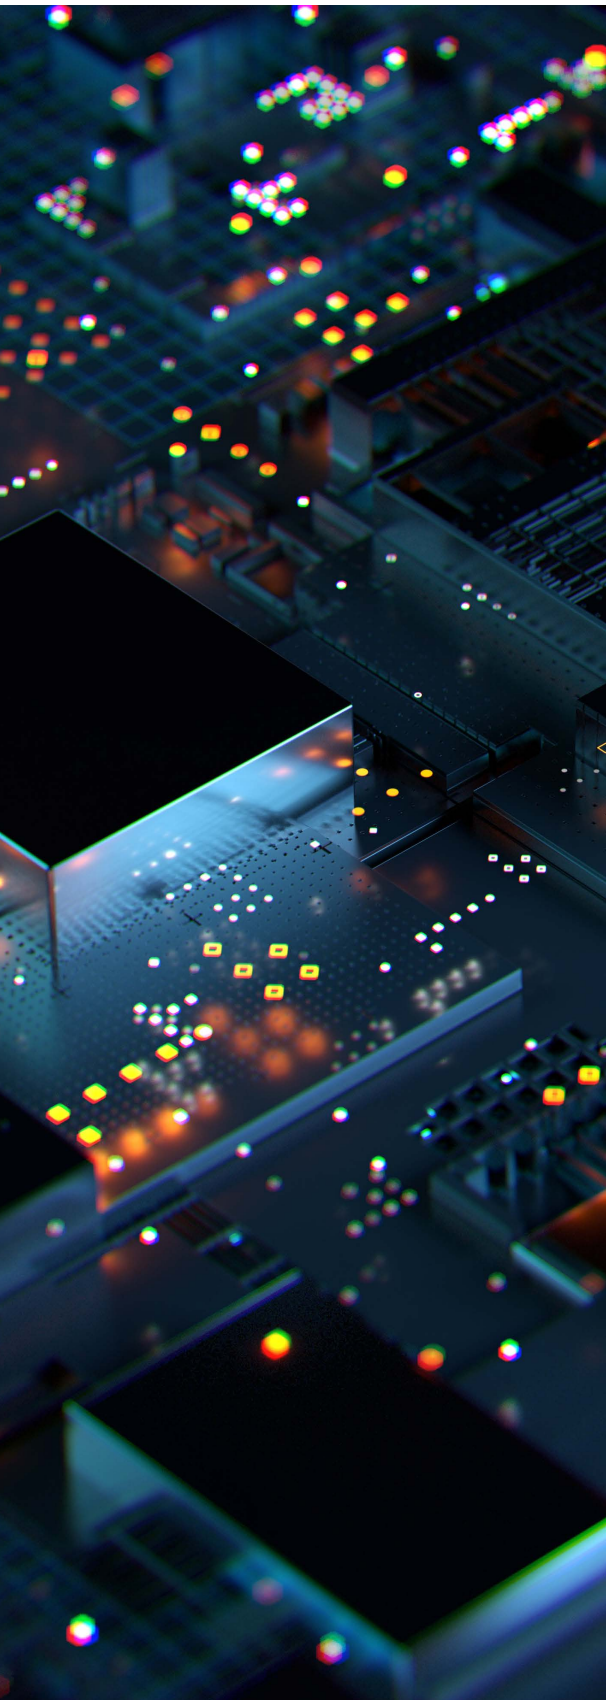
VERTIV.

White paper

800 VDC distribution for AI data centers: Architecture selection criteria and design elements





Executive summary

Advances in AI computing have driven rack power densities beyond 100 kW/rack. As rack densities approach 600 kW, conventional data center IT distribution voltages, like 240/415 VAC, become impractical, even with 3-phase power shelves. Increasing the voltage provides more power with the same conductors and thus fewer power feeds and copper. AI chip manufacturers and power infrastructure vendors are designing future data centers to run on 800 VDC, in part due to its maturity in the electric vehicle (EV) industry. In this paper, we assess 800 VDC architectures across eight criteria and recommend a rack-based architecture as the sensible catalyst for 800 VDC deployments. We then provide a detailed explanation of eight key 800 VDC design elements, including arc flash mitigation, circuit protection and selectivity, impact of dynamic AI loads, distribution polarity, isolation and grounding, maintenance, resiliency, and workforce competency and training. Recommendations for these design elements are also provided for the rack-level architecture.

Key takeaways

1. **Six of the eight decision criteria favor a rack-level architecture for near-term deployment.** These criteria are: supply chain maturity, failure domain size, installation and commissioning expertise, operational maturity, regulatory readiness, and AC & DC load flexibility. The remaining decision criteria, efficiency and token throughput, favor facility-level architectures. The rack-level architecture provides lower execution risk than centralized architectures.

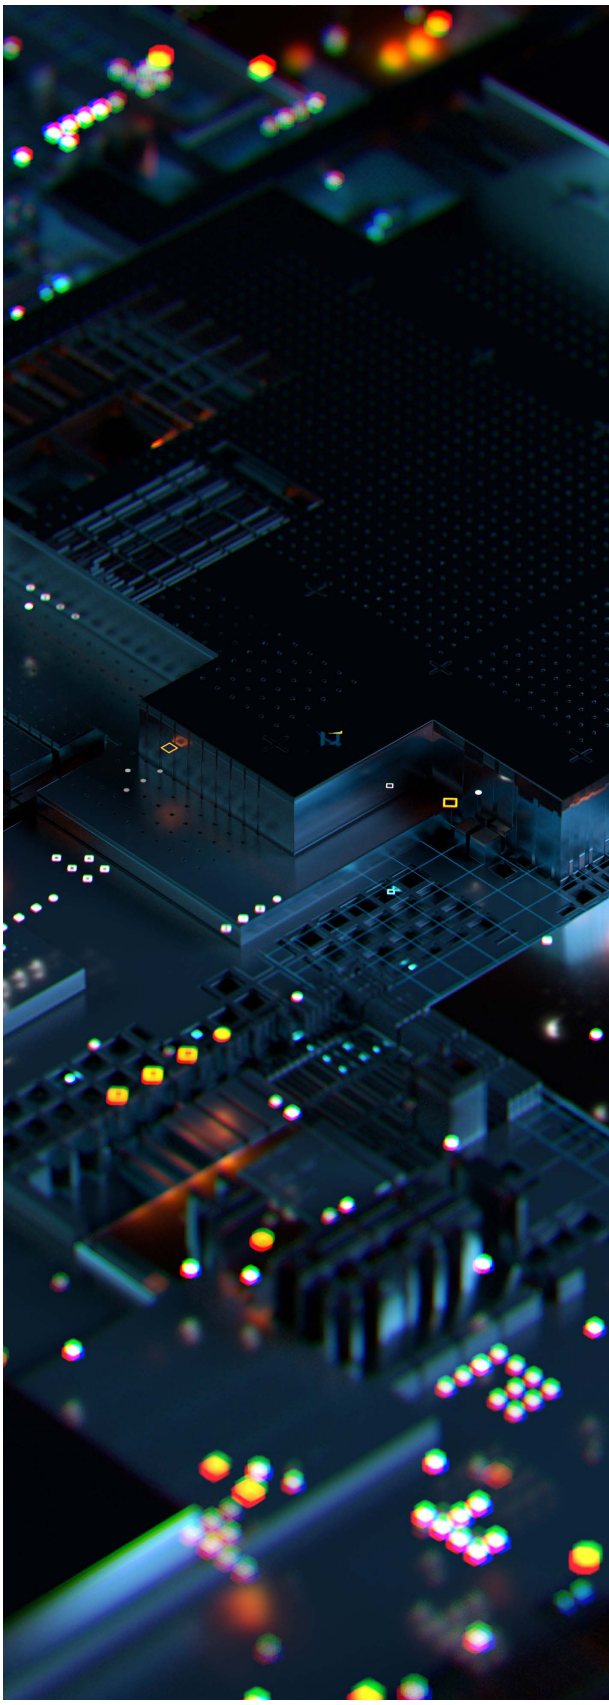
2. **The rack-level architecture moves the AC-to-DC conversion stage (i.e., rectifiers) out of the IT rack and places it in the power rack (a.k.a. sidecar).** These rectifiers convert 400-480 VAC to 800 VDC and distribute it directly to the adjacent IT rack(s).

3. **Installing and operating power racks benefit from standardization.** Power racks can be validated once with arc flash studies, protection coordination, and other engineering studies. Whereas a facility-level architecture differs for each data center and must be validated for every project.

4. **No published timeline exists for a DC-specific arc flash hazard calculation standard.** The rack-level architecture confines the 800 VDC arc-flash scope to a standardized power rack. The enclosure arc containment and incident energy at defined test points can be certified once by the manufacturer under an internal arc test methodology adapted from IEC 61641 and applied to every unit deployed.

5. **Synchronous AI training workloads can create load transients with negative impacts.** If not addressed with energy storage and other design practices, these load transients can induce voltages that interfere with critical IT load.

6. **The rack-level architecture is a stepping stone, not the final stop in this 800 VDC evolution.** Starting at the rack will allow the industry to accelerate the learning curve for installation, operation, and maintenance. Later, facility-level architectures will leverage this experience.



Introduction

The drive to maximize tokens per watt has resulted in AI chip roadmaps with increasing transistors, projecting rack densities up to 1 MW by 2027.[1] As IT rack densities increase to these levels, it becomes increasingly difficult to find the space for all the power circuits and associated connectors at 480 or 400 VAC. Space allocation is made even more challenging in a liquid-cooled AI rack with direct-to-chip manifolds and, where used, in-rack CDUs. Increasing voltage helps address space constraints. For example, doubling voltage allows the same conductor to carry twice the power at the same current. This helps explain AI data centers' support for 800 VDC. But why 800 VDC instead of 800 VAC? Primarily because 800 VDC was the most readily scalable thanks to the electric vehicle (EV) industry. Components like connectors, contactors, and protective devices certified to IEC , and voltage class respectively, can be adapted for data center applications with substantially less standards development EVs have also advanced the performance of power semiconductors.

800 VDC can be implemented in multiple data center architectures. The available options are largely distinguished by where the conversion from VAC to 800 VDC occurs. The conversion location can range from adjacent to the compute rack to the data center's utility service entrance. These architectures are described in the companion white paper, "Descriptions of 800 VDC distribution architectures for AI data centers."

While each architecture has merit, we show that the rack-level approach is the most sensible choice for near-term deployment. Early-stage deployments benefit from smaller failure containment zones, and standardizing on a smaller, repeatable system (the power rack) accelerates reliability improvements, installation experience, and operational learning.

We propose eight criteria for choosing which architecture to deploy and recommend a rack-based architecture as the sensible catalyst for 800 VDC deployments. We then provide a detailed explanation of eight key 800 VDC design elements, including

1. arc flash mitigation,
2. circuit protection and selectivity,
3. impact of dynamic AI loads,
4. distribution polarity,
5. isolation and grounding,
6. maintenance,
7. resilience,
8. workforce competency and training.

Recommendations for these design elements are also provided for the rack-level architecture. This paper provides some initial rack-level power-smoothing design guidance.

[1] https://developer.nvidia.com/blog/nvidia-800-v-hvdc-architecture-will-power-the-next-generation-of-ai-factories/

Assessing each architecture across eight criteria

How does a data center owner decide which architecture to deploy, either today or as part of a longer-term plan? We propose eight key criteria in Table 1 to help with this evaluation. The first six favor rack-level architectures and address execution risk, supply chain readiness, and operational maturity. The remaining two favor facility-level architectures and address efficiency and floor space utilization.

#	Criterion	Favors	Why
1	Scalability and supply chain maturity	Rack-level	EV-derived supply chain is mass-producing at rack-level voltage/power range today
2	Failure containment zones (failure domain)	Rack-level	Smallest failure domain – one power rack feeds one AI rack
3	Installation and commissioning expertise	Rack-level	Repeatable, self-contained scope shortens the learning curve
4	Operational maturity	Rack-level	Standardized MOPs/SOPs/EOPs apply across hundreds of identical units
5	Regulatory standards	Rack-level	Bounded, standardized scope is easier to certify ahead of a DC arc-flash standard
6	AC/DC load flexibility	Rack-level	Preserves existing AC distribution up to the cabinet; no parallel AC path needed
7	Total conversion efficiency (PUE optimization)	Facility-level	Fewer, larger conversion stages and shorter upstream runs cut distribution losses
8	Token throughput (decrease inter-rack latency)	Facility-level	Moves conversion out of the IT aisle, letting compute racks sit closer together

Table 1. Eight criteria for assessing applicable rack- and facility-level architectures for data centers. Source: Vertiv.

1. Scalability and supply chain maturity

Scaling any new technology to hundreds of megawatts requires a mature and disciplined supply chain. Supply chain isn't only about transforming components into final products. It also means that the product's quality meets customer expectations, at scale. 800 VDC products used by other industries will help to establish a supply chain for products listed for use in AI factory applications[2] and will be easier to scale. This favors rack-level architectures, where conversion equipment is smaller, more numerous, and built around components already in mass production. Rack-level architectures allow operators to build 800 VDC AI factories in the short term while allowing time for the other architectures to develop (a phased approach).

For facility-level architectures, the supply chain is less mature. Multi-megawatt solid-state transformers (SST), 800 VDC switchgear, and high-ampacity DC distribution equipment exist but are produced in smaller volumes by fewer suppliers. Solid-state distribution products, such as solid-state circuit breakers (SSCBs), are advancing but are not yet in mass production at the ratings facility-level deployment requires.

2. Failure containment zones (failure domain)

As with many emerging technologies, early deployments may experience higher rates of equipment issues, design refinements, commissioning challenges, and operational errors as products and operating practices mature. Therefore, in the context of a new 800 VDC architecture, an important question is: how many AI racks are at risk of downtime due to a common cause failure? In other words, what's its failure domain?

The failure domain is smallest in rack-level architectures. Each power rack typically feeds one AI rack. A pod-level conversion system serving an entire row of AI racks has a larger failure domain. A facility-level conversion system serving an entire data hall has the largest. Unless specifically designed for concurrent maintenance, planned work on a system can affect all racks within its failure domain.

3. Installation and commissioning expertise

Most electrical contractors have extensive experience installing and commissioning AC systems in data centers; fewer have direct experience with DC systems, and even fewer with 800 VDC. Perhaps the closest experience comes from installing and commissioning large uninterruptible power supply (UPS) battery plants, but installation and commissioning tasks do differ for DC plants. For example, polarity must be verified at every termination before energization, since reversed polarity immediately damages equipment, unlike in AC systems, where a phase rotation error is recoverable. This polarity verification requirement illustrates why 800 VDC installations require procedural training beyond what AC installation and commissioning experience provides. While several industry certifications address general data center operations and electrical safety, there are limited certification programs today that specifically address the hazards, procedures, and protection technologies unique to 800 VDC data center systems.

A facility-level architecture requires installing and commissioning 800 VDC from the main electrical room to the AI rack. The conductors, switchgear, protection coordination, and grounding schemes are unfamiliar to most installation teams. Furthermore, facility-level design elements such as arc flash studies and protection coordination must each be reworked for every site's unique layout and electrical service. This limits the rate at which install teams can learn and increases the probability of errors.

The rack-level architecture concentrates the 800 VDC scope within a self-contained piece of equipment that is delivered, set in place, and connected to existing AC infrastructure. The same installation and commissioning procedure is repeated hundreds of times within a single data center, then again at the next data center. This repeatability shortens the learning curve and reduces the risk of installation error. The industry experience built from installing thousands of rack-level systems becomes the foundation for installing facility-level systems when the supply chain and standards mature enough to support them.

[2] https://www.powersystemsdesign.com/articles/how-800v-ev-innovations-are-redefining-ai-data-centers/97/23694

4. Operational maturity

Facilities teams have limited experience operating DC systems beyond UPS batteries. An example of how DC operation differs is the need to de-energize equipment with energy-storage components such as capacitors. The lockout/tagout (LOTO) procedure must include capacitor discharge time and verification.

The repeatability of a rack-level solution that benefits installation also benefits operations. Methods of procedure (MOPs), standard operating procedures (SOPs), and emergency operating procedures (EOPs) are developed for a single equipment type, then applied across hundreds of identical units. These procedures also benefit spares inventory, diagnostic tools, and maintenance training. A technician who has been properly trained, and practiced the appropriate safety procedures, executes these procedures faster, more consistently, and with fewer errors. Rack-level architecture amplifies this benefit. A well-designed training program can build proficiency on one unit that transfers directly to identical units in the field. Organizations that invest in DC-specific training before deployment will see the operational maturity benefits of rack-level architecture compound.

A facility-level 800 VDC architecture changes more of the data center's electrical operation at once. The gray space, distribution runs, switchgear, and protection systems all transition from AC to DC, expanding the scope of new training, new procedures, and new safety boundaries. For example, personal protective equipment (PPE) requirements, arc flash boundaries, and lockout/tagout procedures would need to be updated for these new areas. This is achievable but represents a larger operational change than confining 800 VDC to the rack level.

5. Regulatory standards

Several electrical and mechanical interface standards developed for EVs are relevant to data center 800 VDC architectures, including IEC 62196-3 (connectors), the IEC 61851 series (conductive charging), and the ISO 26262 series (functional safety). Components certified to these standards, such as connectors, contactors, and protective devices, can be adapted to data center applications with significantly less new standards development than a clean-sheet design. Safety mechanisms developed for EVs, including high voltage interlock loop and pyrofuse, have direct application to data center 800 VDC equipment. Standards specific to data center 800 VDC distribution, however, are still being developed.

Arc flash hazard calculations, coordination studies, and protection schemes do not yet have the depth of established practice as they do for AC systems. Performing these analyses is more straightforward at the rack-level, where the system is bounded, standardized, and largely independent of upstream variations. This allows data center designers to apply the same study to every power rack in the facility. Arc flash standards under active development are summarized in the key design elements section.

Facility-level architectures distribute 800 VDC over much longer runs and through more equipment, each variation of which requires its own analysis. Standards development will eventually catch up, but in the near term, the highest probability of regulatory success lies with rack-level architectures.

A related point applies to the surrounding electrical system. Because all upstream infrastructure remains conventional AC, no changes are required to the facility transformers, switchgear, UPS systems, or PDUs. The 800 VDC scope is inserted at the boundary of the existing power system rather than replacing any part of it. This is particularly valuable for retrofitting existing data centers, where replacing the upstream electrical infrastructure would trigger re-certification of the entire system under standards that are still being defined.

6. AC/DC load flexibility

Cooling compressors, pumps, fans, lighting, building management systems, security systems, and many networking devices continue to operate on AC voltage. As a result, most data centers must support a mix of AC and DC loads. A rack-level conversion architecture preserves the traditional AC facility distribution up to the server cabinet. This allows AC and DC loads to coexist and draw from the same upstream AC power system without requiring parallel distribution paths. Even AI factories with homogeneous AI loads still have AC-powered equipment.

With a facility-level 800 VDC architecture the primary distribution operates at 800 VDC. Therefore, AC loads either require a parallel AC distribution system running alongside the DC system, or they require a dedicated 800 VDC to AC inverter stage. Either path adds equipment, cost, and complexity.

7. Total conversion efficiency (PUE optimization)

The criteria above have favored rack-level architectures. The remaining two criteria favor facility-level architectures. Total conversion efficiency depends on three factors: the number of conversion stages from utility to chip, the distribution losses between stages, and the efficiency of each conversion. The facility-level architecture can improve on all three compared to the rack-level architecture. Figure 1 illustrates a range of 800 VDC architecture efficiencies based on our detailed analysis. Facility-level designs occupy the upper efficiency range for the following reasons:

- While more conversion steps do not automatically make a system inefficient, they raise the performance bar of each step. For example, to match a single-step system efficiency of 97.0%, a five-step system requires every single step to average a near-perfect 99.393% efficiency (99.393% x 99.393% x 99.393% x 99.393% x 99.393% = 97.0%). For this reason, converting from MV AC directly to 800 VDC favors efficiency for centralized architectures.
- The further upstream the AC to 800 VDC conversion occurs, the fewer distribution losses we have compared to 480 VAC. For the same power, 800 VDC carries slightly more current than 480 VAC three-phase (1,250 A vs 1,203 A per MW), but does so on two conductors rather than three. Also, given a fixed (non-dynamic) load, DC current is immune to the skin effect and inductive reactance expected with AC current. As a result, architectures that convert to 800 VDC further upstream generally experience lower losses in conductors, busbars, and circuit breakers. Circuit monitoring, metering, and protection (trip units) add a small amount of fixed losses to these square losses. Our analysis showed 800 VDC distribution architectures reduced conductor losses by 20-40% relative to the power rack architecture.[3]
- Every component that electricity runs through generates losses. These losses are either fixed, proportional, square, or a combination thereof, which determine the efficiency curve of converters.[4] An efficiency curve typically peaks at some load % and falls significantly at low and high percent loads. The most important converter efficiency is the one your data center operates at. The IT load, converter redundancy, and converter oversizing determine the operating efficiency.
- With diversified workloads (i.e., asynchronous), centralizing converters is more efficient because it removes the oversizing that occurs with distributed converters. Each small converter must be sized for its server's individual peak. However, with synchronous workloads like AI training, this advantage largely disappears since all the servers peak simultaneously. What remains is a more modest efficiency benefit: centralizing converters consolidates fixed losses. The overhead losses from electronics controls, sensing, communication, and monitoring are roughly constant per converter unit, regardless of its rated capacity. A 1 MW converter does not have 100 times the controller power losses of a 10 kW converter. Fixed losses as a percentage of rated capacity can be lower in larger units, though this effect is design-dependent and should not be assumed across all products.

[3] 10 MW data center, 95% load, N+1 redundant (4 to make 3), 240 kW/rack, 55-230 ft. 800 VDC conductor length depending on architecture

[4] Examples of losses: Fixed (transformer core losses), proportional (semiconductor switching losses), and square (transformer winding losses)

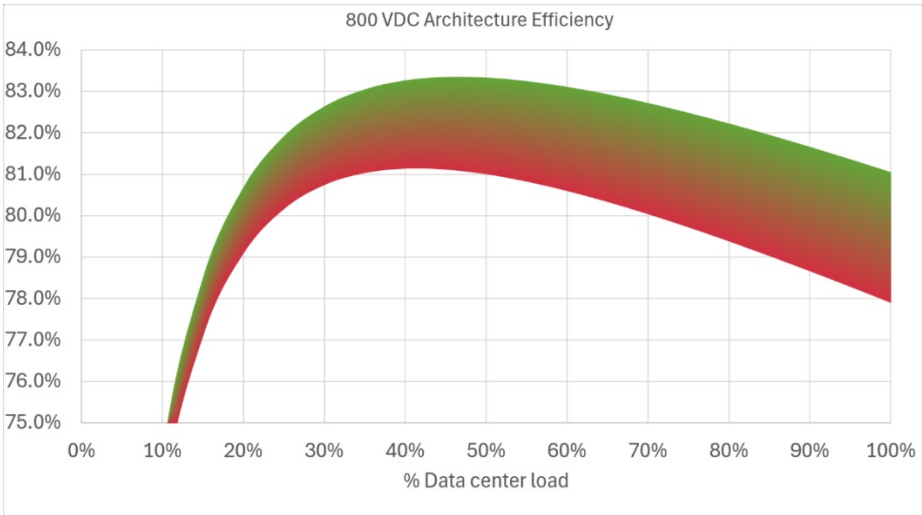


Figure 1. Range of efficiencies for 800 VDC architectures based on Vertiv's efficiency analysis. Source: Vertiv.

Rack-level architecture is the optimal choice today

Weighing the eight criteria, the rack-level architecture is the optimal choice for 800 VDC deployments today. Six criteria favor it. The two that favor facility-level architectures (total conversion efficiency and token throughput) do not yet outweigh the execution risk advantages of rack-level deployment. For operators retrofitting to support high-density AI workloads without rebuilding their upstream electrical infrastructure, the power rack is the most practical and lowest-disruption path to 800 VDC.

This balance shifts as compute rack densities approach 1 MW. The electrical and cooling infrastructure, and some of the associated standards will evolve to support these higher densities given the timeline for 1 MW compute racks.[5] The criteria that favor rack-level architectures today will remain valid, but their relative weight will decline as the supply chain matures and rack densities push past the practical limits of rack-level conversion. Using rack-level architectures as a stepping stone moves operators up the learning curve and buys time for all IT equipment and some physical infrastructure equipment to transition to 800 VDC input, thereby maximizing the efficiencies granted by 800 VDC.

[5] Rob Campbell, 1 MW Racks and Supply Chain Resilience: Planning for the Data Centers of the Future, Data Center Knowledge, 7/25/2025

Key design elements of 800 VDC for the rack-level architecture

Designing a rack-level architecture for reliability, safety, and maintainability requires deliberate choices across eight elements. The rack-level architecture allows a manufacturer to certify design choices once, as part of a power rack product. While these design elements apply to all 800 VDC architectures, the **implementation guidance in this section targets the rack-level architecture**, as this represents the most sensible near-term deployment path for most operators.

Before discussing the design elements, it's helpful to understand some important electrical characteristics of direct current. These characteristics are all consequences of current flowing continuously in one direction.

- **No zero crossing:** Alternating current rises and falls from a positive peak to a negative peak crossing the zero volts point 100 to 120 times per second. Direct current flows at a fixed voltage and doesn't cross the zero-volt point. This is illustrated in Figure 2. The absence of a zero crossing is the root cause of the arc sustainment and contact material transfer behaviors described below.

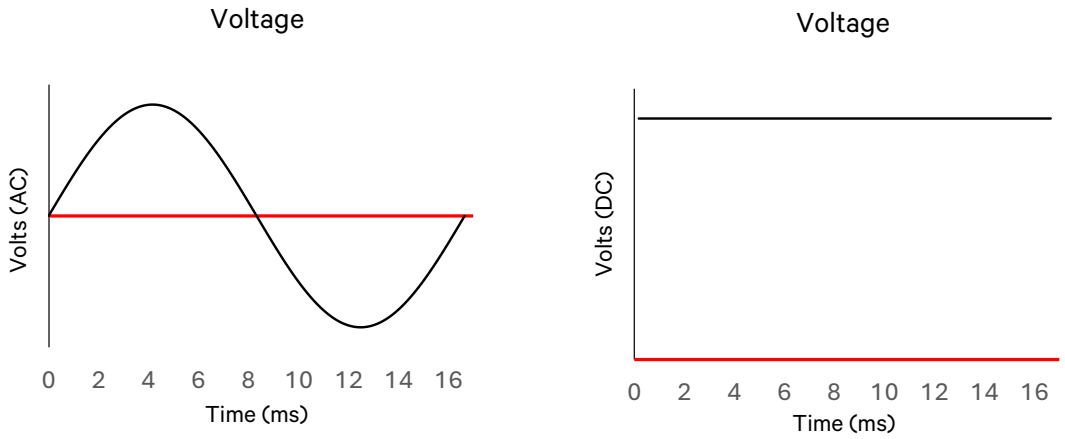


Figure 2. AC sine wave illustrating the zero crossing

- **Arc sustainment:** Both AC and DC systems produce arcs during faults and when closing or opening circuits under load. In a 480 VAC system, at each zero crossing, current drops to zero momentarily, giving the ionized gas between the contacts time to cool before the next half-cycle. This is typically sufficient to prevent the arc from re-establishing. In an 800 VDC system, there's no zero crossing to help extinguish the arc. Therefore, DC systems must use DC-rated disconnection devices. These devices use different contact geometry, arc chute design, and sometimes magnetic blow-out coils to extinguish the DC arc. As discussed later in the "Circuit protection and selectivity" subsection, this is one of the reasons why SSCBs are beneficial in 800 VDC systems.
- **Contact material transfer:** When the contacts in an 800 VDC system open under load, the arc that forms travels in the same direction: from positive to negative. The positive contact (anode) loses metal through arc erosion, and the negative contact (cathode) accumulates that metal as a deposit. This transfer happens every time the contacts open under load. In an AC system, the arc reverses direction with each half-cycle, distributing erosion roughly evenly between the contacts. After enough DC switching operations, the anode develops craters, and the cathode develops mounds where the transferred metal accumulates. These geometric changes have two consequences. First, contact resistance increases – the crater surface on the anode has less effective contact area and higher surface resistance than the original flat face. Second, there's an increased risk of welding, which prevents opening during the next operation.

We now discuss the eight key design elements:

- **Arc flash mitigation:** Assess the hazard and limit the amount of energy during a fault.
- **Circuit protection and selectivity:** A combination of human safety and system availability.
- **Impact of dynamic AI loads:** Issues caused by AI training workloads that generate transients.
- **Distribution polarity:** How 800 VDC (+/- 400 VDC) or (800 VDC) is referenced and distributed.
- **Isolation and grounding:** Safety and how the system behaves during a fault.
- **Maintenance:** Work practices specific to 800 VDC.
- **Resiliency:** Prevent downtime during an outage in a DC system.
- **Workforce competency and training:** Service qualifications and readiness for 800 VDC-specific hazards and procedures.

1. Arc flash mitigation

Arc flash[6] and the resultant incident energy[7] are dangerous and require an assessment to understand the hazard, which informs the required personal protective equipment (PPE) and approach boundary distances. The two key types of standards governing arc flash are electrical safety standards and hazard calculations. For example, NFPA 70E is the electrical safety standard for the U.S., which specifies safety requirements like PPE.[8] The standard requires PPE if incident energy is 1.2 calories/cm^2 or higher, which applies to both AC and DC arc flash. The IEEE 1584-2018 standard is used globally for calculating the incident energy (i.e., hazard calculations);[9] however, this applies to AC systems. This is because the arc in DC systems behaves differently. Not only does DC lack a zero crossing, but its source energy is limited by converter control loops rather than by transformer impedance in AC systems. Also, DC systems have significant capacitive energy stored downstream of the source (e.g., bus capacitance, PSU input capacitors, etc.) that can discharge into a fault independent of the source's ability to current-limit. Reverse current blocking devices prevent stored energy from feeding upstream faults, as is common in microgrids and telecom.

Several IEEE papers are commonly referenced for DC arc flash calculation methods:

- Arc Flash Calculations for Exposures to DC Systems
- DC-Arc Models and Incident-Energy Calculations
- Modeling the Dynamic Behavior of DC Arcs (this lists current DC arc flash models)

These calculation methods can help inform protection levels while the industry works toward codes and standards. The following organizations are actively supporting this effort:

- Open Compute Project (OCP): Developing white papers[10] and curating a list of codes & standards[11] related to DC power distribution. Vertiv is an active partner for research and development.
- Underwriters Laboratories (UL): Started the Direct Current Safety Research Consortium (DCSRC)[12]
- National Fire Protection Association (NFPA): Developing a model for direct current arc flash.[13]
- International Electrotechnical Commission (IEC): TC 73 Short-circuit current in DC systems – Part 1 Calculation of currents – LVDC. [14]

[6] "short circuit conditions when electricity passes through the air and releases tremendous amounts of heat and light"

[7] "the amount of thermal energy impressed on a surface, a certain distance from the source, generated during an electrical arc event.", p 1.

[8] NFPA 70E, Standard for Electrical Safety in the Workplace, 2027 Edition

[9] IEEE 1584-2018, IEEE Guide for Performing Arc-Flash Hazard Calculations

[10] OCP, Data Center Facility – Low Voltage Direct Current Power Distribution, p 100, March 2026

[11] OCP has curated a large inventory of LVDC Power Distribution Codes & Standards

[12] Direct Current Safety Research Consortium (DCSRC)

[13] NFPA, Development of a Model for Direct Current (DC) Arc Flash Phenomena, 2/2026

[14] IEC, TC 73: Short-circuit currents

These organizations are all advancing, but no published timeline exists for a DC-specific arc flash hazard calculation standard equivalent to IEEE 1584-2018.

The rack-level architecture bounds the 800 VDC arc flash scope within a standardized power rack. The enclosure arc containment and incident energy at defined test points can be certified once by the manufacturer under internal arc test methodology adapted from IEC 61641[15] and applied to every unit deployed. Site-specific arc flash studies are still required under electrical safety standards like NFPA 70E, but the studies focus on the upstream AC distribution (well-established methods) and the 800 VDC output connections at the IT rack boundary (a bounded, standardized interface).

Design guidance for rack-level arc flash mitigation: We explain key mitigation techniques below and how to apply them to rack-level designs.

- **Fast interruption:** Reduces incident energy at the fault by orders of magnitude compared to mechanical DC circuit breakers, reducing calculated incident energy to thresholds below that which drive high-category PPE and expanded arc flash boundaries. There are various technologies available that interrupt DC faults faster than traditional mechanical circuit breakers. The total clearing time includes the sensing time required to trigger the device. The total time is the most important for specifying protection devices but nearly all the vendor information fails to clarify whether “interruption” or “clearing” time accounts for all delays. Table 1 lists the total interruption time from fastest to slowest. Because incident energy scales linearly[16] with arc duration, large reductions in total interruption time produce proportional reductions in incident energy. SSCBs currently offer the fastest clearing times for 800 VDC applications and are discussed further in the circuit protection subsection.

Device	Total interruption time[17]	Method	Reusable
Solid-state circuit breaker	tens to hundreds of µs	Semiconductor switching, no arc	Yes
Hybrid SSCB	2 to 4 ms	Mechanical + semiconductor commutation	Yes
Pyrofuse	2 to 5 ms	Explosive severance	No
Current-limiting fuse	2 to 8 ms	Element vaporization + arc quench	No
Circuit breaker	20 to 100 ms	Electromagnetic and mechanical contact separation Faster CBs use electronic fault sensing and high-speed opening mechanisms (e.g., Thomson-coil), enabling both earlier fault detection and faster contact separation.	Yes

Table 1. List of interrupting technologies and rough estimate of total interruption time.

[15] IEC, Technical Specification 61641

[16] Michael D. Fontaine, Stephen W. McCluer, Arc-in-a-Box: DC Arc Flash Calculations Using a Simplified Approach, p 15-5

[17] Estimates based on various sources including, A Review of Solid-State Circuit Breakers and Hybrid DC circuit breaker with coupled inductor for automatic current commutation – contact DCRI@vertiv.com for more sources

In a rack-level architecture, the manufacturer certifies the SSCB coordination and clearing time as part of the power rack. This means that every deployment benefits from that certification.

- **Current limiting at the source:** Fuses, current-limiting reactors, isolation transformers, reverse current blocking (e.g., diode), and converter output current limits restrict the fault current available to feed an arc, reducing the arc power for any given duration. Note that limiting the current for traditional circuit breakers may actually increase incident energy by increasing the interruption time. Thermal magnetic circuit breakers typically open faster with a higher fault current. Therefore, all these solutions must be designed holistically. In a rack-level architecture, the fuse coordination, current-limiting element sizing, and breaker interaction are designed and tested once as part of the power rack.
- **Physical separation and arc-resistant enclosures:** IEC 61641 internal arc testing applied to 800 VDC switchgear validates that if an arc does occur, the enclosure directs the energy away from the worker. Rack-level designs use a common enclosure and power rating across deployments.
- **Touch-proof connector design:** Removes the possibility of accidental contact with energized 800 VDC conductors during connect/disconnect operations. In a rack-level architecture, touch-safe connector performance is certified once at the product level.
- **Conservative working distance and PPE selection:** Until DC-specific calculation methods are fully validated, apply working distances and PPE selection according to the higher of the calculated incident energy values rather than the most precise. This practice benefits from the consistency of rack-based designs.

2. Circuit protection and selectivity

Circuit protection means clearing faults before they damage equipment or injure personnel.[18] Overcurrent protection requirements are addressed in IEC 60364-4-43[19] and in Articles 110 and 240 of NFPA 70.[20] Device-level standards are further along. Electromechanical DC circuit breakers are covered by IEC 60947-2,[21] UL 489,[22] and UL 1066.[23] Solid-state circuit breakers (SSCBs) and semiconductor hybrid[24] circuit breakers are covered by UL 489I[25] and by IEC 60947-10.[26] Fuses are covered by IEC 60269[27] and UL 248.[28]

In a DC system, there's a very real chance that a fault may not trip the circuit breaker. One of the challenges in an 800 VDC power rack comes from the behavior of its converters. The converters are either fault-limiting (i.e., current limiting) or fault-feeding. Converters are typically fault-limiting. So, during a fault, their control loops limit output current to roughly nominal levels. Typically, fault current is also supplied by capacitors in every converter connected to the common bus. The peak fault current is high in magnitude but only lasts a few hundred microseconds.

[18] See Data Center Facility – Low Voltage Direct Current Power Distribution, for more detailed information.

[19] IEC 60364-4-43, Low-voltage electrical installations - Part 4-43: Protection for safety - Protection against overcurrent. Memberships and subscriptions apply

[20] NFPA 70, National Electrical Code. Memberships and subscriptions apply.

[21] IEC 60947-2, Low-voltage switchgear and controlgear - Part 2: Circuit-breakers. Memberships and subscriptions apply.

[22] UL 489, Molded-Case Circuit Breakers, Molded-Case Switches and Circuit-Breaker Enclosures. Memberships and subscriptions apply.

[23] UL 1066, Power Circuit Breakers up to 1000 V AC and 1500 V DC Used in Enclosures. Memberships and subscriptions apply.

[24] Hybrid SSCBs use a mechanical disconnect in parallel with the semiconductor to lower losses and to provide physical isolation.

[25] UL, The Path to Safety and Quality for Solid-State Circuit Breakers.

[26] IEC 60947-10, Low-voltage switchgear and controlgear - Part 10: Semiconductor circuit-breakers. Memberships and subscriptions apply.

[27] IEC 60269-1, Low-voltage fuses - Part 1: General requirements. Memberships and subscriptions apply.

[28] UL 248-1, Low-Voltage Fuses - Part 1: General Requirements. Memberships and subscriptions apply.

Now consider a conventional thermal-magnetic breaker protecting a branch circuit on this bus. The capacitor-discharge peak tends to be too short to trigger the magnetic trip element. This mechanism has mechanical inertia, making it difficult to respond to a sub-millisecond pulse. The steady-state current that follows may be too low to trigger the magnetic element and may look like high load current to the thermal element. The result is a very long trip time or no trip at all. While the rectifiers feed the fault at their current limit, the bus voltage decreases. If the fault isn't cleared before the bus voltage falls below the minimum operating voltage of the downstream converters, the IT loads dependent on that bus brown out or drop, assuming no redundant buses. Depending on the system design, this voltage drop can occur in a few hundred microseconds. Protection speed in an 800 VDC system is therefore driven by at least three requirements: limiting arc energy, tripping reliably against a current-limited fault, and clearing before the healthy loads lose their bus voltage.

Design guidance for rack-level protection: Designers have two choices on how to address faults: clear them fast enough to avoid downtime, or allow the bus voltage to collapse and rely on redundant buses to maintain the critical load. Trying to clear the fault fast enough before the bus voltage collapses is challenging without readily available SSCBs listed for the application. These solutions are also more expensive. A lower-cost option is to temporarily collapse the bus voltage and clear the fault using the collective current from all rectifiers to open a traditional mechanical DC breaker. Fuses whose I2t melt rating is verified against the available capacitive discharge energy may also work. The bus capacitance, fault-loop inductance (resistance to the change in current), and converter current limit, together determine how long the loads survive a fault without bus redundancy.

Selectivity means the protection device closest to a fault clears and isolates only the affected part of the system, while the healthy portions continue operating. In a data center, selectivity is an availability requirement as much as a safety one. In AC systems, selectivity is achieved through time-current coordination: devices are set to trip at different time delays for the same overcurrent. Downstream devices open first with each upstream device waiting longer. In 800 VDC systems, there are two challenges with this approach:

First, no standardized time-current curves exist across vendors. DC breaker manufacturers publish curves for their own products, but no industry standard defines the test conditions or reporting format, and no recognized practice explains how to account for the capacitor-discharge peak in a coordination study.[13] Coordination design is currently vendor-specific engineering work.

Second, the intentional delays set in upstream DC circuit breakers increase the likelihood that the DC bus voltage will drop below the minimum operating voltage. Having redundant converters does no good if a single converter fault takes down the entire 800 VDC bus. An additional complication is that all energy storage devices on the bus, including the input capacitors of load converters, may discharge into the fault. The protection devices in the healthy parts of the system must not trip on this outgoing discharge current.[29] These are challenging problems for power rack designers due to the short timescales and the interactions of many components on the same bus. Power rack designers should test their systems under varying failure modes to validate that the system can maintain operation. This is another reason why it's best for the industry to lead with a rack-level architecture.

Design guidance for rack-level selectivity: As fast as semiconductor devices are, it's challenging to clear faults before the bus voltage sags and before upstream devices see a sustained overcurrent. When no intentional delay is used, the arcing time is limited to the instantaneous interruption time, thereby reducing arc flash energy. When slower protection devices must be used, increasing the bus capacitance for voltage hold-up helps but remains difficult to achieve within these time ranges. Designers will also need to balance this additional stored energy with an increase in fault energy.

Though an evolving solution in DC systems, designers can also use directional sensing where multiple sources share the bus. Current-direction detection is common for larger (e.g., 1600 amp) AC circuit breakers with electronic trip units, but not for DC circuit breakers. However, this can be done using an external system that uses current sensors to measure the current and a controller. The controller determines whether a fault exists, identifies its direction and location, and signals the correct overcurrent devices to trip open and isolate the fault. The sensing system can therefore prevent tripping healthy zones. This solution is most useful when BBUs and converters are paralleled on the same bus. Because no standard coordination methodology exists for these fault profiles, manufacturers must validate this coordination with the specific devices and revalidate if a device is substituted.

[29] OCP, Data Center Facility – Low Voltage Direct Current Power Distribution, p 90 & 91, March 2026

3. Impact of dynamic AI loads

AI training workloads draw power synchronously across the entire compute cluster, creating rapid power transients (very fast peaks and valleys). "Rapid" means that the current rises or falls in a short amount of time (i.e., di/dt or slew rate), but the amount of current and time depends on where you measure. You could measure 10-100 amps per nanosecond[30] at the GPU die, but this rate decreases the further upstream you measure. The reason for this starts with the multi-gigahertz clock speed frequencies of GPUs and other AI accelerators[31]. These chips perform their computations so fast with such high power density that the di/dt demands can only be met by placing capacitors adjacent to the transistors on the silicon die. These decoupling capacitors (aka decaps) serve as tiny energy-storage reservoirs, feeding the transistors and preventing induced voltages that interfere with tight voltage regulation at the die level. If the decaps were placed further away, the path inductance[32] would be so large that voltage regulation would be impossible. Once the on-die decaps discharge, they must be recharged from an upstream reservoir, and the cycle repeats.

This same phenomenon continues upstream at the chip package, motherboard, intermediate bus converter (IBC), and eventually to the service switchgear. However, there are two significant changes with each cascading energy replenishment:

- The energy capacity is larger for each upstream reservoir: There are progressively fewer but larger energy reservoirs as power moves upstream from the die. The additional physical space available outside the processor package allows the use of larger capacitors and batteries with greater energy storage capacity. These larger reservoirs can replenish downstream reservoirs for longer periods before requiring replenishment themselves. In effect, each reservoir acts as a time buffer, allowing the upstream energy source to respond at a lower slew rate.

- The distance between reservoirs increases: As slew rates decrease the energy reservoirs can be located farther from the load, reducing the need to occupy valuable IT space. Note that increasing conductor length also increases the overall inductance of the circuit (i.e., loop). However, the lower slew rates reduce the associated voltage drop as shown by this formula:
  $V = L \cdot di/dt$ V=voltage drop L=inductance di/dt=rate of current change

As a result of these two changes, as the transient progresses upstream through the electrical system, it becomes flatter and wider. We can say that its "frequency" and "amplitude" decrease as illustrated in Figure 3, while carrying the same amount of energy.

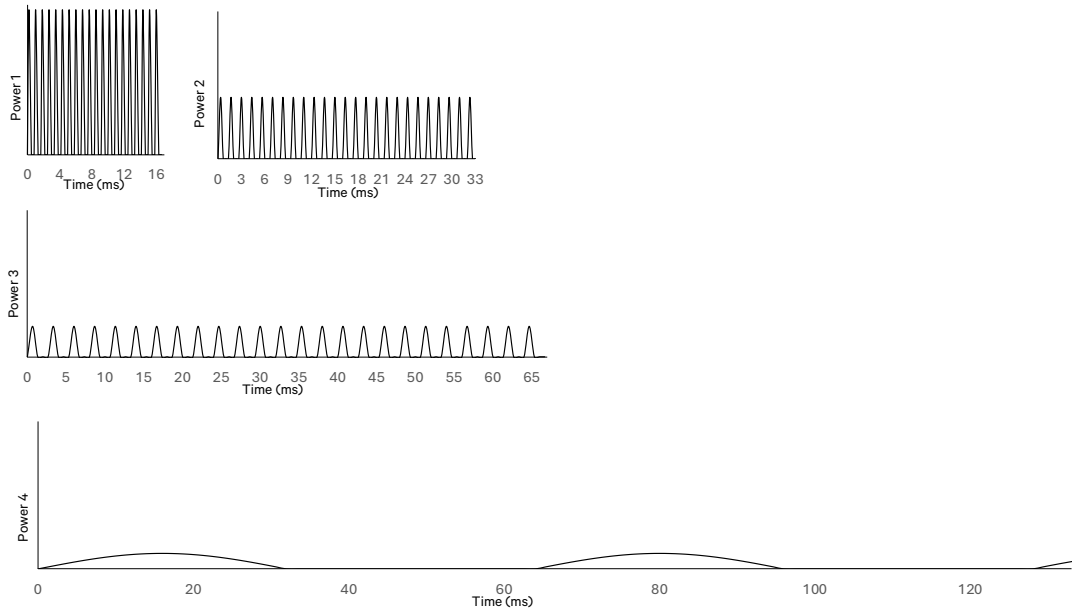


Figure 3. Illustrative transient load wave form as it progresses up through the electrical system toward utility with equal energy (area under the curve). Source: Vertiv.

[30] Asav Gandhi, Cumulative Voltage Droop: Preventing Silent Data Corruption in Modern AI Chips, Anasim, 4/2026

[31] Bryon Moyer, New Approaches To Power Decoupling, 10/24/2024

[32] Keysight, What Is Inductance?

What we just described above is known as power smoothing, the process of absorbing or supplying energy along the power path to mitigate transients so that the power drawn from the utility changes more slowly than the instantaneous power demanded at the AI workload.

By the time the power transient reaches the power rack or the busway's tap-off box, its frequency is on the order of kilohertz. And this is the point where we provide guidance on addressing the transient with a rack-level architecture.

Design guidance for rack-level dynamic load impact:

- Size capacitance at the power rack to address the kHz transients at that point in the power train. If the DC rectifiers feed a DC busway, the capacitance should be placed at the row or pod level.
- Verify busway inductance is minimized by design. Laminated busway (positive and return conductors sandwiched tightly with thin dielectric) mitigates the magnetic field and minimizes inductance. Non-laminated busway or cable-based distribution with wide conductor separation has much higher inductance.
- Provide local snubber capacitance in tap-off boxes. Snubber capacitors integrated in the tap-off box, placed immediately adjacent to the plug-in stabs, absorb transient load energy.
- If using cable conductors, lay them tightly together in the tray with alternating positive and negative conductors. This will reduce the loop inductance.
- Account for skin effect when specifying conductor sizing. DC ampacity tables assume uniform current distribution. If the distribution will carry meaningful kHz content, effective resistance at those frequencies is higher than the DC value, and the conductor sizing should reflect this. Where practical, use multiple smaller parallel conductors rather than fewer large ones - smaller conductors have less severe skin effect penalty because their radius is closer to the skin depth.

4. Distribution polarity

Distribution polarity is the choice of how many conductors supply the 800 VDC and how their voltages are referenced to each other. The question of the reference point is one reason it's difficult to discuss polarity without grounding approaches. But for simplicity, we explain polarity in this subsection before introducing isolation and grounding in the next subsection. Two polarity configurations dominate current data center design discussions.

Unipolar (800 VDC across two conductors). A unipolar system uses two active conductors: a positive (supply) pole and a return. All loads connect across the full 800 VDC.

Bipolar (+/- 400 VDC with a mid-point). A bipolar system typically uses three conductors: a positive pole at +400 VDC, a negative pole at −400 VDC, and a mid-point conductor. Loads can connect across the full 800 VDC (pole-to-pole) or across 400 V (pole-to-midpoint). The system delivers two voltage levels from the same distribution. Table 2 provides a comparison between the two polarities.

Attribute	Unipolar (800 VDC)	Bipolar (+/- 400 VDC)
Active conductors	2	3 (2 without midpoint)
Voltage between any conductor and ground	Up to 800 VDC	Up to 400 VDC
Ground-fault current from a single pole (solidly grounded)	Full fault current at 800 VDC	Approximately half, at 400 VDC
Insulation and clearance requirements	Rated for 800 VDC throughout	Rated for 400 VDC pole-to-midpoint & 800 VDC pole-to-pole
Load flexibility	800 VDC loads only	400 & 800 VDC loads
Grounding	Simple	Complex
Distribution copper (for the same power)	Baseline	~ 50% more copper[33]
Risk of imbalanced loads	No	Yes

Table 2. Comparison of unipolar and bipolar distribution polarity

Bipolar is beneficial for supporting both 400 VDC and 800 VDC IT equipment and provides three safety benefits, highlighted in Table 2. This bipolar-voltage option provides flexibility for data center operators designing their own AI compute servers. In contrast, a unipolar system requires one fewer conductor, avoids load-balancing complexity, and simplifies grounding schemes. Unipolar provides four benefits highlighted in BRAND COLOR in Table 2. Depending on priorities, unipolar could be the preferred polarity, but may not be compelling enough to override engineering preferences and constraints of existing data center operators. Grounding schemes, discussed next, also influence the polarity choice, but we assume the use of unipolar for that discussion.

5.Isolation and grounding

Isolation and grounding are two separate design decisions that are often discussed together because both concern how the 800 VDC system relates to earth and to surrounding systems. This subsection defines each, then covers the grounding choices relevant to a unipolar rack-level architecture.

Isolation is a means of separating two parts of a circuit such that dangerous currents (e.g., AC-side ground fault) on one side are not directly conducted to the other side, mainly for human safety. In a data center, this is usually accomplished with a transformer.[34] The windings transfer power magnetically rather than through a direct conductive path, so a fault potential on one side is not imposed on the other. Isolation also creates a separately derived source, which allows the designer to choose the grounding system on the DC side rather than inherit it from the primary.

[33] This assumes a bipolar application that requires a midpoint conductor for a total of 3. This represents one extra wire compared to unipolar's two or ~50% more copper.

[34] 800 VDC systems typically use high-frequency isolation transformers to lower the losses across the transformer.

Grounding (known as earthing in IEC terminology) is the intentional connection of the electrical system to earth, mainly to protect humans from shock hazards. Grounding is a separate decision from isolation: isolation determines whether the DC system can be grounded independently, and if so, grounding determines how it is referenced to earth.

For DC systems, only two grounding schemes are used: TN-S (solidly grounded) and IT (ungrounded or high-resistance grounded - HRG).[35] Other schemes common in AC systems, such as TT and TN-C, are not used in DC because their current paths can run through grounding electrodes into soil or concrete, causing electrolytic corrosion of the electrodes. This corrosion mechanism is specific to DC current (moving in the same direction) as opposed to AC current, which switches direction 100 to 120 times per second.

TN-S (solidly grounded) systems (Figure 4, left) connect one point of the DC system directly to earth, with separate conductors for the current-carrying return and the protective earth. A ground fault in this scenario provides a low-impedance path and opens the overcurrent protection device (e.g., fuse). This system is simple to design, and its fault behavior is well understood, it provides lower operational availability. This system also increases shock hazard from leakage current. When used with unipolar 800 VDC, the TN-S system requires only one overcurrent protection device. In addition to lowering cost this could potentially free up space for IT compute.

IT grounding systems either have no intentional connection to earth (ungrounded, middle of Figure 4) or connect to earth through a resistor (HRG, rightmost of Figure 4). A first ground fault produces either no fault current or a small current flow, so the system continues operating while the fault is located and cleared. This continuity is valuable in a data center. These systems require an insulation monitoring device to detect the first fault, and components must be rated for the full pole-to-pole voltage because that voltage appears to ground when one pole faults. The HRG's grounding resistor stabilizes the system's reference to earth, making insulation monitoring more stable. The floating ground system is susceptible to stray capacitance.[36] Also, clearance distances must be larger than in a grounded system, and the operator must locate and clear the first fault quickly; otherwise, a second ground fault on the opposite pole creates a pole-to-pole fault at full energy. NFPA 70 makes ground-fault detection mandatory in these two DC grounding systems.

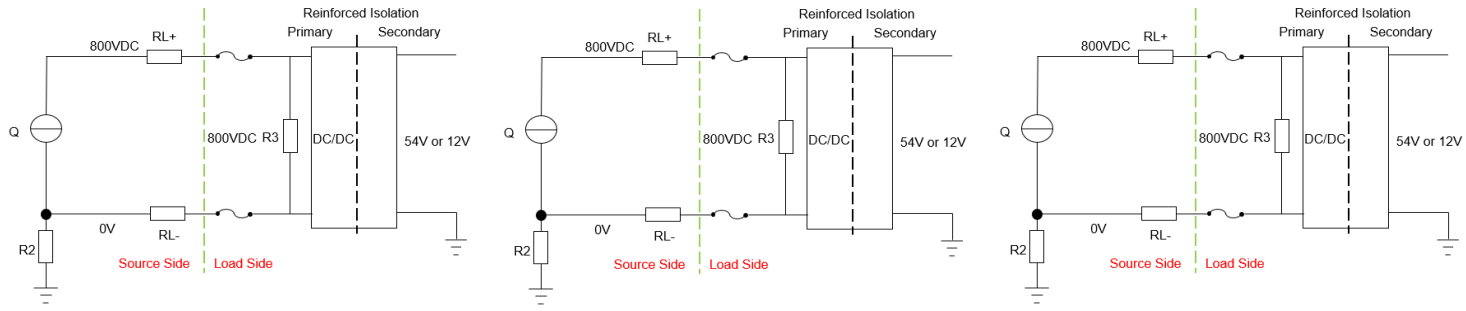


Figure 4. From left: TN-S (solidly grounded), IT-ungrounded (middle), and IT-high-resistance grounded (right).

Design guidance for the unipolar rack-level architecture: The rack-level system should be isolated at 800 VDC to create a separately derived source, which allows the designer to choose the DC grounding scheme rather than inherit it from the AC supply. While the grounding choice should consider the data center operator's constraints, the HRG system provides higher availability because it reduces the likelihood of a ground fault escalating into an arc flash event. The data center operator must implement fault-location procedures, safety procedures, and training. The HRG scheme requires an insulation monitoring device, changes the touch voltage present during service, and drives specific lockout/tagout steps. These consequences are covered in the maintenance subsection. The HRG system comes at a higher cost than the more simple solidly-grounded system.

[35] For more detailed information on grounding see, Data Center Facility – Low Voltage Direct Current Power Distribution, p 63

[36] Stray capacitance is the small, unintended capacitance that exists between conductors and ground due to the physical proximity of cables, equipment, and structures. In ungrounded systems, it can cause the system voltage to shift relative to ground.

6. Maintenance

Maintaining an 800 VDC system requires procedures that differ from AC practices, and even from those of UPS battery plants. This subsection covers key maintenance insights that follow from the electrical characteristics described earlier in this section.

- **Verification of de-energization:** In an AC system, opening an upstream breaker generally de-energizes the load side within one cycle, and a voltage check confirms it. In an 800 VDC system, opening a breaker doesn't guarantee the conductors downstream are safe to touch. Bus capacitance, PSU input capacitors, and BBU modules can hold hazardous voltage for seconds to minutes after the source is isolated.[37] This is why de-energization procedures and training are important.
- **Pre-charge:** The same capacitance that must be discharged before work must be charged in a controlled way when the system returns to service. Connecting a discharged capacitive load directly to an energized 800 VDC bus draws an inrush current, which can weld contacts, damage the load, or trip upstream protection. A pre-charge circuit brings the load capacitance up to bus voltage gradually. Every re-energization of an 800 VDC circuit should follow the pre-charge sequence.
- **Switching device wear:** Contact material transfer, described at the start of this section, changes how switching devices are maintained. A contactor installed with consistent polarity develops predictable wear: the anode contact erodes, the cathode accumulates. If polarity is inadvertently reversed during maintenance, the eroded contact becomes the cathode and the deposited contact becomes the anode. The wear pattern reverses on already-degraded surfaces, accelerating the path to failure.
- **Demarcation between facilities and IT staff:** In a conventional data center, the boundary between facilities responsibility and IT responsibility generally falls at the rack PDU, and everything on the IT side of that boundary operates at voltages IT staff are trained to work around. An 800 VDC architecture moves this higher voltage across that boundary and into the IT rack. Data center organizations should articulate demarcation points that require facilities-level qualification versus IT staff-level qualification.

Design guidance for rack-level maintenance:

- Specify discharge time and expose voltage verification points without having to open the enclosure (e.g., power rack).
- Require operation counting on switching devices. Maintenance driven by operation count requires the system to count. Specify this as a monitoring requirement rather than estimating switching frequency.
- Connectors are the most frequently disturbed part of the 800 VDC path. Touch-safe designs improve safety during connection or disconnection. A high-voltage interlock loop (HVIL), a low-voltage sense circuit routed through the connector, de-energizes the bus before the power contacts separate. Both of these practices are borrowed directly from EV industry.
- Define the facilities/IT boundary in the design, and train to it. State where the 800 VDC hazard ends, enforce it physically where possible, and align qualification requirements and PPE with that boundary.
- Build procedures around the standardized power rack. Methods of procedure (MOP), standard operating procedures (SOP), and emergency operating procedures (EOP) written once for a standard power rack apply across most if not every deployment.

[37] Fluke, About absence of voltage testing

7. Resiliency

Resiliency is the system's ability to continue powering the IT load through component failures, utility disturbances, and planned maintenance. A unique property of AI training is that if any single compute rack fails, the training job must resume from its last saved training state (i.e., a checkpoint). Therefore, it makes sense to apply the same resiliency to all racks in the same AI training cluster. Resiliency should be commensurate with the cost of a restart, which comes back to the value placed on token throughput. Most of the practices that achieve this in AC systems apply directly to 800 VDC: redundant sources, N+1 module counts, redundant feeders, and stored energy to ride through outages. However, one AC capability has no direct DC equivalent.

No static bypass, and what replaces it: In case of a critical internal failure, a double-conversion AC UPS can transfer the critical load to a static bypass, connecting it directly to utility. An 800 VDC rectifier plant has no equivalent bypass path around the rectification stage. If the entire rectifier plant fails, it can't bypass to its AC input source and feed a DC load. A design approach to address this drawback for all 800 VDC architectures is striping.[38] This approach is used in AC architectures, where IT loads are distributed across redundant power paths to limit the impact of any single power path failure. You can further increase system availability by adding redundant rectifiers and DC-to-DC modules to power shelves. In other words, N+1 power feeds to N+1 power shelves, but the power shelves themselves have N+1 redundant converter modules. Energy storage redundancy may also be required in 800 VDC systems for battery backup and power-smoothing such as capacitors.

Paralleling is easier in DC: Paralleling AC sources requires synchronizing voltage, frequency, and phase angle, and maintaining that synchronization under varying load. Ineffective paralleling produces large circulating currents between sources. DC has no frequency or phase. Paralleling requires only that the source voltage be the same. Modern rectifiers accomplish this with droop control: each unit reduces its output voltage slightly as its current increases, so units naturally share the load. The practical result is that N+x redundancy is straightforward to build at the power rack. N+x redundancy is also possible with 800 VDC busways (i.e., feeders), assuming a larger rack-level DC plant that supports a row of racks.

Hot-swap capability: If a failed module can be replaced without de-energizing the bus, a module failure becomes a maintenance event rather than an outage. Hot-swap is therefore central to the availability case for a modular power rack. Hot-swap is feasible when the module's power path is switched off electronically before the physical contacts separate, such that no current flows at the moment of separation. Hot-swap capability must be designed into the module.

Concurrent maintainability: Concurrent maintainability means any component can be taken out of service for maintenance without dropping the IT load. Achieving it in an 800 VDC system requires:

- Enough redundant capacity that the remaining modules carry the full load with at least one unit out of service. N+1 at the power rack achieves this for module-level maintenance.
- A way to isolate the unit under maintenance without disturbing the bus.

Design guidance for rack-level resiliency:

- Specify redundancy in all modules (rectifiers, DC/DC converters, and BBUs).
- Specify rectifiers with load sharing control.
- Specify that modules are designed to be swapped in case of failure.
- Specify at least N+1 redundant AC input feeds.

[38] Striping is the practice of supplying a system with independent, redundant power sources. If one power path fails, the other keeps the hardware running. 2N is 2 redundant paths whereas N+1 could be 3+1 paths, also known as 4 to make 3.

8. Workforce competency and training

The eighth element addresses the people who will install, commission, operate, and maintain the electrical architecture. An 800 VDC architecture that is well-designed but poorly serviced will have a shorter operational life and lower availability than its design intended. The design elements in this paper, including arc flash mitigation, circuit protection, isolation and grounding, and maintenance procedures, all assume a well-trained technician. Training is what bridges the gap between a well-designed system and a safely operated and maintained system.

The preceding design elements describe several hazards that have limited direct AC equivalent. For example, arcs with no zero crossing, pre-charge sequencing requirements on re-energization, and directional contact erosion. These are procedural realities that require DC-specific training. Additionally, the first design element, arc flash mitigation, recommends conservative working distance and PPE selection as a key arc flash mitigation strategy until DC-specific codes and standards are published. While established guidance is still being developed, a technician's knowledge and training become critical to maintaining proper PPE selection and working distances are observed.

That means proper training of the people servicing this equipment is a critical way organizations can control risks associated with the 800 VDC architectures. Organizations building 800 VDC service competency should treat it that way by investing in DC-specific hazard awareness, structured isolation procedures, and a service culture where safe decision-making is the expectation.

The protection technologies described throughout this paper – solid-state circuit breakers (SSCBs), insulation monitoring, arc flash sensing, and fuse coordination – each require specific operational and maintenance competencies. Training programs should be mapped to the specific protection technologies present at each node in the power train. This approach also enables modular credentialing. A technician can be authorized for specific equipment categories based on demonstrated competency rather than holding a single certification that implies equal readiness across all equipment types.

Design guidance for rack-level workforce competency and training:

- Require demonstrated competency. Practical evaluation on representative equipment should be required prior to service authorization.
- Evaluate behavioral readiness alongside technical skill. Informed judgment matters when codified standards are in development
- Build service competency into equipment lifecycle decisions. Evaluate vendor ability to provide competent service support alongside hardware.

Next steps

- Ask power rack vendors for verification test results showing how their solutions perform with respect to arc flash, protection & coordination, isolation, and shock hazards.
- Determine the required level of resiliency for your IT applications. For example, AI inference is typically a revenue-producing workload, which tends to warrant redundancy throughout the power and cooling train.
- Confirm that the rack-level solution allows for maintenance practices like capacitor pre-charge and discharge sequencing and touch-safe design.
- Begin developing standardized documentation in anticipation of 800 VDC architectures. Specifically, methods of procedure (MOP), standard operating procedures (SOP), and emergency operating procedures (EOP).
- When evaluating equipment vendors, confirm their ability to provide competent service support through documented maintenance procedures and ongoing workforce development alongside the hardware. Begin developing or sourcing 800 VDC-specific service competencies that address DC arc flash analysis, PPE selection, capacitor discharge and pre-charge procedures, polarity verification, and protection technology-specific maintenance.

References

Vertiv white paper, Descriptions of 800 VDC distribution architectures for AI data centers.

Appendix:

This white paper was initially generated with the assistance of AI technology, but reviewed, referenced, and edited by professionals in Vertiv using the following platforms:

- Google. (2026). Gemini 3.1 Pro (Feb 19 2026 version) [Large language model]. https://gemini.google.com/.
- Anthropic. (2026). Claude Opus 4.8 (May 28, 2026 version) [Large language model]. https://claude.com/.
- Anthropic. (2026). Claude Sonnet 5 (June 30, 2026 version) [Large language model]. https://claude.com/.
- Anthropic. (2026). Claude Opus 5 (July 24, 2026 version) [Large language model]. https://claude.com/.

The platform was used for the following topics and purpose prompts:

- Outline organization and story flow
- Organization, drafting, and editorial refinement
- 800 VDC development for AI factories (i.e., AI data centers)
- Explanation of technical topics
- Finding sources for specific technical values, such as circuit breaker disconnection times
- Review of the paper for clarity, conciseness, grammar, logic, and factual integrity

AI-generated content was reviewed, fact-checked, and approved by Vertiv subject matter experts prior to the final draft and publication.



Vertiv.com | **Vertiv Headquarters,** 505 N Cleveland Ave, Westerville, OH 43082, USA

© 2026 Vertiv Group Corp. All rights reserved. Vertiv™ and the Vertiv logo are trademarks or registered trademarks of Vertiv Group Corp. All other names and logos referred to are trade names, trademarks, or registered trademarks of their respective owners. While every precaution has been taken to aim for accuracy and completeness here, Vertiv Group Corp. assumes no responsibility and disclaims all liability for damages resulting from use of this information or for any errors or omissions. Specifications, rebates, and other promotional offers are subject to change at Vertiv's sole discretion upon notice.

SL-80689 (10/26)