

MarSum Solutions: High-Efficiency Data Center Power Systems: HVDC and 800V Distribution

Large cloud, colocation, and artificial intelligence (AI) data centers are pushing electrical distribution far beyond legacy server-room assumptions. Rack power that once measured in single-digit kilowatts is now moving toward hundreds of kilowatts, with future AI factory racks being discussed at the megawatt scale. At those levels, power architecture becomes a core part of compute density, efficiency, thermal management, and total cost of ownership.

High-voltage direct current (HVDC) distribution is becoming a practical path for reducing conversion stages and moving more power through the same physical space. Instead of distributing low-voltage 48 V or 54 V DC inside the rack, selected next-generation architectures are exploring or moving toward 380 V to 800 V DC distribution, with centralized AC/DC conversion and localized DC/DC conversion close to the compute load. This paper explains why that transition is happening and what engineering discipline is needed to make high-voltage data center power systems safe, reliable, and efficient.

Why AI data centers need a new power architecture

Traditional data center power systems rely on alternating current (AC) distribution, uninterruptible power supplies, power distribution units, rack power supplies, and multiple downstream conversion stages before power reaches processors, memory, networking, and accelerators. That architecture is mature and serviceable, but it becomes strained as rack-level power rises.

The basic constraint is current. For a fixed power level, lower distribution voltage requires higher current, which increases conduction losses, busbar mass, connector heating, cable bulk, and mechanical complexity. A 54 V rack bus can work well at moderate rack power, but it becomes difficult to scale into the hundreds-of-kilowatts range. Raising distribution voltage to 400 V, 800 V, or similar HVDC levels reduces current for the same power transfer and leaves more space for compute hardware instead of copper and conversion equipment.

Legacy 48V/54V distribution vs. high-voltage DC distribution

The shift from low-voltage rack distribution to HVDC changes where conversion happens, how protection is designed, how energy storage can be integrated, and how much electrical infrastructure must sit near the compute equipment.

Legacy 48V/54V rack distribution

Low-voltage DC distribution is familiar, relatively safe to service, and compatible with existing server power ecosystems. The drawback is current. As rack power rises, busbars, cable assemblies, power shelves, connectors, and thermal interfaces become larger and more difficult to package. Repeated conversion stages also add loss, heat, and potential failure points.

400V to 800V DC distribution

Higher-voltage DC distribution moves more power at lower current, reducing copper bulk and distribution losses. It also supports architectures where AC/DC conversion is centralized at the facility or row level, while high-ratio DC/DC conversion occurs close to the processors. The tradeoff is that protection, isolation, sensing, service procedures, and compliance requirements become more demanding.

What is inside a high-efficiency HVDC data center power system

A next-generation HVDC data center power architecture can be described as a coordinated chain of power conversion, distribution, protection, and control functions. The major engineering blocks are consistent even when implementation details differ.

Utility and medium-voltage interface: Medium-voltage service, grounding, switchgear, metering, and utility rules establish the front end of the data center power train.

Centralized AC/DC rectification: Large rectifier systems convert facility AC power into a regulated HVDC bus, reducing repeated conversion stages while raising redundancy, cooling, and serviceability requirements.

HVDC busways and rack distribution: HVDC distribution carries power through the data hall to rows, racks, or sidecar power assemblies. Busway design, connector selection, insulation coordination, creepage, clearance, and service disconnect strategy become central to safe operation.

Rack and board-level DC/DC conversion: Localized DC/DC converters step the HVDC bus down to intermediate and board-level voltages used by processors, memory, networking, and accelerators. High conversion ratio, efficiency, fast transient response, and thermal density are key targets.

Solid-state protection and sensing: HVDC systems need fast fault detection, current sensing, voltage sensing, isolation monitoring, arc-energy management, and interruption strategies coordinated with fuses, breakers, controls, and maintenance procedures.

Backup energy and controls integration: Battery backup units, central battery systems, generators, renewables, and site-level controls can be integrated more directly into a DC architecture when ride-through behavior, charge control, fault isolation, and priorities are coordinated.

NVIDIA 800 VDC architecture as a case study

NVIDIA has publicly described an 800 VDC architecture for next-generation AI factories, with the goal of supporting 1 MW IT racks and beyond. The concept is straightforward: reduce the number of conversion stages, distribute power at higher voltage, and perform the final high-ratio DC/DC conversion close to the compute load. In NVIDIA's described future architecture, facility-level AC/DC conversion establishes an 800 VDC backbone that can be distributed through the data hall and used directly by compute racks.

The engineering motivation is consistent with the broader HVDC trend. Higher voltage lowers current for the same power level, reducing copper mass and distribution loss. Centralized conversion can remove layers of intermediate AC switchgear, transformers, power distribution units, and rack-level power shelves. Late-stage DC/DC conversion then steps 800 VDC down to processor-adjacent voltages where tight regulation and fast transient response matter most.

The case study is important because it shows that HVDC data center power is moving from research and specialty deployments into serious next-generation AI infrastructure planning. The challenge is not simply choosing an 800 VDC bus. The system must include protection, sensing, isolation, standards alignment, thermal design, serviceability, and controls that are appropriate for high-density facilities operating as critical infrastructure.

Power electronics design requirements

HVDC data center architectures are enabled by power electronics, but they are also constrained by power electronics. The design must be efficient at high load, stable through transients, safe under fault conditions, and serviceable in facilities where downtime is expensive.

Conversion efficiency and power density

Every conversion stage produces heat. Centralized rectifiers, rack converters, and board-level regulators must be optimized as a system so gains in one stage do not create thermal or packaging problems in another.

Isolation and insulation coordination

Higher voltage requires intentional creepage, clearance, isolation monitoring, touch-safe design, connector selection, and maintenance procedures. Mechanical layout and contamination control become part of the electrical design.

EMI/EMC and power quality

Fast-switching converters, high-current transients, and dense rack layouts can create electromagnetic interference and conducted noise. Filter topology, grounding, shielding, switching strategy, and layout discipline all influence compliance and field reliability.

Protection and fault interruption

DC faults do not naturally pass through a current zero crossing like AC faults. Protection must detect abnormal current quickly, interrupt fault energy, manage arc risk, and leave the system in a safe diagnostic state.

Thermal management

Rectifiers, converters, busbars, protection devices, and server trays all add heat to an already dense thermal environment. Electrical architecture must be coordinated with air or liquid cooling strategy from the beginning.

Controls, monitoring, and diagnostics

High-density power systems need accurate telemetry and coordinated controls across facility rectifiers, rack converters, backup energy systems, and compute loads. Diagnostics must support fast fault isolation without unnecessary shutdown of healthy capacity.

Deployment patterns and market direction

Interest in HVDC data center power is rising because AI workloads, high-performance computing, and cloud expansion are increasing rack density faster than traditional low-voltage architectures can comfortably support. Data center operators are looking for power systems that reduce energy loss, conserve copper, simplify distribution, and preserve space for revenue-generating compute equipment.

AI factories and hyperscale campuses: The strongest near-term pull is from AI training and inference facilities where rack density, facility power, and cooling demand are increasing together. These projects benefit from fewer conversion stages, high-density busways, and power architectures that can scale across rows or halls rather than treating each rack as a standalone island.

Retrofitted cloud and colocation facilities: Existing facilities may not move directly to full 800 VDC distribution, but they can adopt intermediate architectures, higher-voltage rows, sidecar power assemblies, or centralized conversion for specific high-density zones. Retrofit planning must account for available utility service, switchgear space, floor loading, cooling capacity, and service practices.

Europe and energy-constrained regions: Regions with expensive energy, constrained grid interconnections, or strong sustainability targets have additional motivation to reduce losses and improve utilization of electrical infrastructure. HVDC can support those goals, but site-specific code compliance and operating procedures must be addressed early.

Integration challenges that determine project success

The most difficult HVDC data center problems usually appear at the interfaces: between facility power and rack power, between protection and serviceability, between converter efficiency and thermal design, and between emerging architectures and existing codes or standards.

Architecture selection

Teams must decide whether the site uses conventional AC distribution, 380 VDC zones, 800 VDC rack distribution, or a phased hybrid architecture. That decision affects equipment selection, floor planning, safety procedures, and future expansion.

Protection coordination

Fuses, breakers, solid-state interrupters, hot-swap devices, sensors, and software controls must work together. Protection must be fast enough to limit DC fault energy but selective enough to avoid unnecessary facility-wide disruption.

Reliability and maintainability

AI facilities need high availability. Modular rectifiers, replaceable converter blocks, redundant controls, thermal monitoring, and diagnostic data all help maintain service while isolating failures.

Rack-level conversion strategy

The location of DC/DC conversion determines copper use, conversion loss, thermal density, and service access. Designers must balance centralized efficiency against localized regulation and transient performance near GPUs and processors.

Standards and safety alignment

High-voltage DC systems require clear labeling, touch-safe design, lockout procedures, connector standards, service training, and code review. Ambiguous requirements can slow deployment even when the electrical concept is sound.

Validation and commissioning

Commissioning should verify normal operation, load steps, ride-through behavior, protection response, thermal margins, and failure modes. Validation planning is especially important when multiple vendors supply rectifiers, busways, converters, protection devices, and rack hardware.

How MarSum supports high-efficiency data center power programs

MarSum Solutions supports high-power conversion, power distribution, controls, and system-integration programs with an engineering-first approach focused on performance, robustness, and manufacturability. HVDC data center power sits at the intersection of facility power, converter design, protection, thermal management, and controls, which is where early engineering decisions have the largest effect on deployment risk.

Our work can include architecture review, AC/DC and DC/DC converter support, power-stage and filter design input, HVDC protection and sensing strategy, EMI/EMC mitigation planning, thermal and reliability assessment, and validation planning. We also help teams translate broad goals such as higher rack density, lower distribution loss, reduced copper, backup energy integration, and serviceability into testable engineering criteria.

For data center operators, equipment suppliers, and project teams moving toward 380 VDC, 800 VDC, or hybrid AC/DC architectures, power distribution should be treated as a core system architecture decision rather than a downstream facilities detail. A disciplined system-level approach can reduce commissioning risk, improve reliability, and support a faster path from concept to deployable infrastructure.

Engagement models

We support customer teams through focused technical consulting, design reviews, power-conversion support, controls and modeling assistance, and deeper co-development efforts depending on program phase. Typical engagements range from early feasibility and requirements definition to prototype validation, field-debug support, and production-readiness planning for high-density data center power systems.

Selected sources: NVIDIA 800 VDC architecture and ecosystem materials; Texas Instruments 800 VDC protection and sensing materials; public 380 VDC data center distribution architecture references.

Contact us today to scope your HVDC data center power, 800 VDC distribution, or high-density power electronics project!