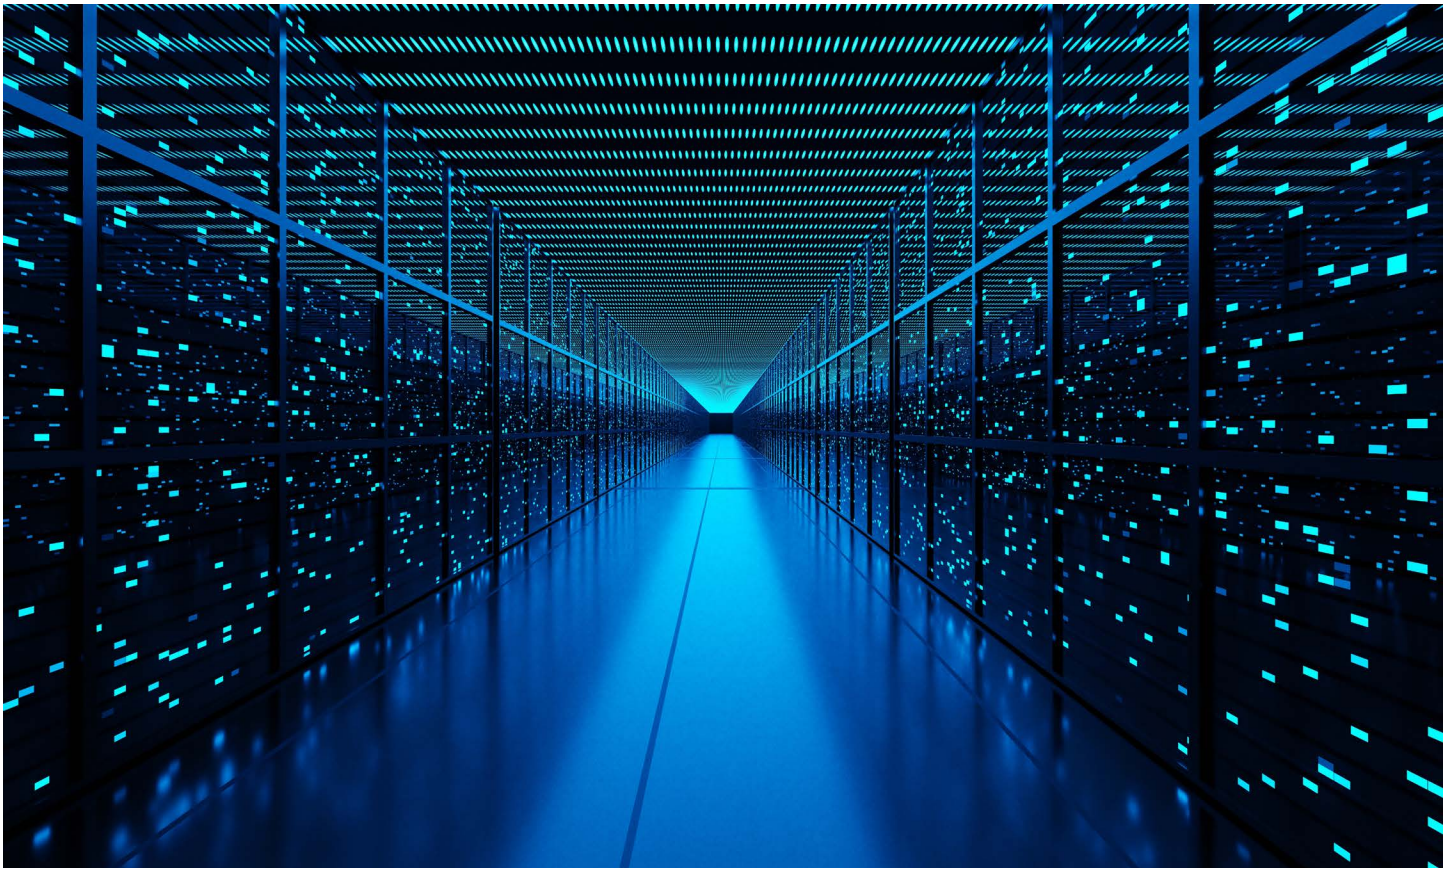
TREND PAPER

800VDC DATA CENTER POWER PATH: FROM UTILITY GRID TO GPU

Examine escalating data center rack power requirements, unique advantages of 800VDC architectures, and how TE is helping enable next-generation data center power systems.



By Sajjad Ahmed

Director of R&D and Engineering
Digital Data Networks
TE Connectivity

ABOUT THE AUTHOR



Sajjad Ahmed

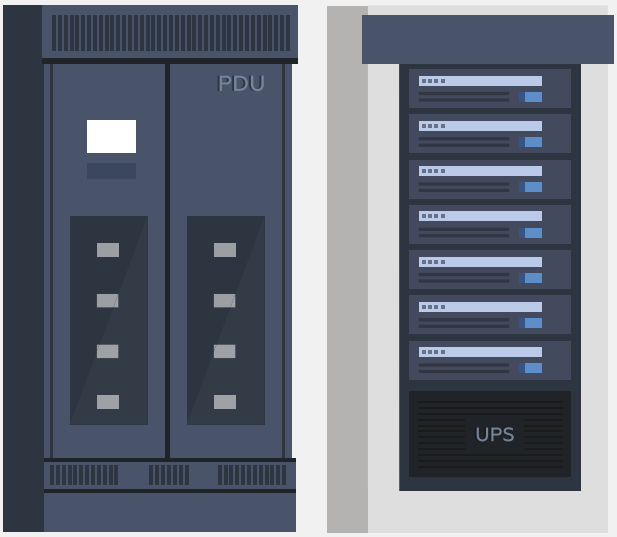
Sajjad Ahmed is the Director of R&D and Engineering for the Digital Data Networks business, where he leads the Advanced Engineering and Solutions team in developing next-generation interconnect architectures in close partnership with global customers. With over two decades of experience in architecting and scaling engineering ecosystems, he combines a deep understanding of multidisciplinary engineering with a proven ability to overcome mass-production challenges. Sajjad has consistently introduced innovations that have advanced the compute industry and continues to drive technologies that shape the future of data centers worldwide.

ABSTRACT

As AI and high-performance computing continue to raise rack power requirements, data center power architecture is entering a period of meaningful change. The traditional model of distributing AC deep into the facility and converting power repeatedly at multiple levels is increasingly challenged by higher current, greater thermal burden, and tighter efficiency targets.

In this environment, 800VDC has emerged as one of the key architectures under consideration. By converting to DC earlier in the power chain and distributing high voltage DC closer to the load, this model offers a path to lower losses, improved scalability, and better alignment with the needs of next generation digital infrastructure.

In this trend paper, we will illustrate both the immediate value of 800VDC and the direction the industry may take in the future. We take a step-by-step view of the complete power path, beginning at the utility grid and traversing through transformation, switching, power conversion, rectification, distribution, protection, and regulation to delivering the power to the GPUs processing the AI workloads.



WHY 800VDC MATTERS

800VDC is more than just a higher voltage architecture - it also creates significant system-level advantages by redefining where conversion occurs and how power is distributed throughout the data center.

By establishing a high voltage DC distribution layer earlier in the power path, the architecture can reduce conversion stages, lower current, and improve distribution efficiency. It also aligns more naturally with battery energy storage, renewable sources, and other DC native infrastructure.

For high density AI environments, these benefits are especially attractive. As facilities are required to deliver more power in less space, the conventional AC centered architecture becomes increasingly difficult to scale without incurring higher losses, larger conductors, more complex cooling burdens, and greater system inefficiency.

Industry Momentum Toward 800VDC

Industry momentum behind 800VDC continues to build as AI drives higher rack power and greater demands on power delivery efficiency. Market forecasts project strong growth through the early 2030s, with early deployments expected in 2026–2027 and broader adoption beginning around 2028. While ±400VDC is seen as a short-term transitional architecture, 800VDC is more widely viewed as a longer-term architecture.

800VDC MOMENTUM

Market growth is accelerating.

Forecasts estimate the 800V HVDC market will reach $5B–$12B by 2031–2032.

Sources: QY Research: Mobility Foresights

Adoption begins this decade.

Initial deployments are expected in 2026–2027, with broader implementation from 2028 onward.

Source: MarketScreener

Efficiency is a key driver.

800VDC architectures are projected to improve system efficiency by 5–10% while reducing power losses by 20–30%.

Sources: Abilkeen; MarketScreener

The industry is converging.

Public industry discussions increasingly position ±400VDC as transitional, with 800VDC as the long-term direction. 4

Source: Data Centre Dynamics

CHOOSING BETWEEN ±400VDC AND 800VDC

As 800VDC architectures move from concept to implementation, one of the most important design decisions is how the DC bus is configured. The two most commonly discussed approaches are a bipolar ±400VDC architecture and a unipolar 0–800VDC architecture. Both can deliver the same total distribution voltage, but they differ in grounding philosophy, protection behavior, insulation requirements, and physical distribution design.

A ±400VDC architecture uses positive and negative rails balanced around ground. One of its key advantages is that each active rail operates at a lower voltage relative to ground than a 0–800VDC system. That can simplify some aspects of insulation coordination, creepage and clearance management, and equipment packaging. At the same time, this approach often creates a more complex physical distribution structure. In practical terms, it can become a three conductor or three busbar problem when the positive rail, negative rail, and ground or reference scheme are considered together. That can mean more copper, more routing complexity, more terminations, and a more demanding interconnect architecture.

A 0–800VDC architecture, by contrast, uses a single high voltage rail referenced to return. One of its key practical advantages is a simpler physical distribution structure. Instead of managing both positive and negative high voltage rails, designers can often reduce conductor count and overall bus complexity. In practical terms, this can mean less copper, fewer busbars, simpler routing, and potentially fewer interconnect elements across the system. The tradeoff is that the active rail sits at a higher voltage relative to return or ground, which can increase demands on insulation, connector design, protection coordination, and serviceability.

The choice between these architectures is therefore not simply a matter of voltage. It is a system-level tradeoff between more copper and more physical distribution complexity on one side, and less copper but higher voltage stress to ground on the other. The preferred choice will depend on the operator's grounding strategy, protection philosophy, converter topology, packaging constraints, and service model.

From a market perspective, the current public signal suggests that ±400VDC is the more clearly defined hyperscaler implementation path today, while 800VDC and potentially 0–800V architectures are gaining momentum as part of the longer term industry direction. The strongest publicly visible hyperscaler activity to date has centered on ±400VDC through open rack power efforts, while the broader ecosystem is moving aggressively toward 800VDC for future AI factories and megawatt class racks. In practical terms, this suggests that ±400VDC may lead in early open hyperscaler implementations, while 800VDC and 0–800V architectures continue to gain traction as the industry evaluates longer term standardization, protection strategies, and system integration models.

From a converter efficiency standpoint, the available evidence does not yet suggest a clear inherent advantage for 0–800VDC over ±400VDC. Modern high voltage DC/DC platforms are being developed to support both architectures at very high efficiency. If the power path is modeled as high voltage DC to 12V, then 12V to 5V, then 5V to 3.3V, the total cascaded efficiency is expected to be broadly similar for both bus configurations when comparable converter topologies are used. In other words, the more important difference is not the downstream DC/DC chain itself, but the upstream architectural tradeoff between bus complexity, copper content, protection approach, and integration strategy.

For TE, this distinction is highly relevant. Connector design, busbar layout, cable assemblies, sensing architecture, creepage and clearance performance, and service interfaces are all influenced by whether the system is built around ±400VDC or 0–800VDC. As a result, the interconnect layer will play a central role in making either approach practical, safe, and scalable.

800VDC DATA CENTER POWER FLOW

An 800VDC architecture should be understood not as a single component change, but as a full system-level shift in where conversion occurs and how power is delivered to the compute load. The overall power path begins at the utility grid, moves through conventional AC transformation and switching, then transitions to high voltage DC through rectification. From there, power is distributed across the facility at 800VDC, handed off at the rack-level, and converted inside the IT equipment into intermediate and low voltage rails required by processors and accelerators.

What makes this architecture important is that each stage has a distinct role in efficiency, protection, and scalability. The upstream portions of the system determine how effectively power can be converted and distributed. The middle of the chain defines how high voltage DC is protected, monitored, and routed to the rack. The final stages inside the IT equipment determine how safely and efficiently that energy is converted into the tightly regulated rails required by GPUs, CPUs, and memory.

In the near term, sidecar architectures are likely to remain an important part of this transition because they provide a practical boundary for AC to DC conversion, protection, energy storage integration, and serviceability as the ecosystem matures. This is especially relevant while standards, protection strategies, and high voltage interconnect approaches are still evolving. In this transitional model, AC busway can continue to feed sidecars or row-level power systems, which then perform the AC to DC conversion before supplying power into the IT rack.

Over time, however, as 800VDC distribution becomes more prevalent, the architecture is likely to shift toward direct DC feed from facility busway into the IT rack, reducing the need for a separate sidecar power rack. In that future model, high voltage DC becomes a more native part of the rack power architecture rather than a function isolated in an adjacent conversion cabinet.

This is where solid state transformers could become especially important. If SST technology matures sufficiently for large scale deployment, it offers a path to convert AC directly into 800VDC, creating a more direct route from the upstream power system into the rack-level DC bus. In that scenario, the role of the sidecar could diminish significantly because the conversion from AC to high voltage DC would move further upstream into the power distribution network itself.

The deep dive that follows traces this path step by step, from utility entry to the processor, to show where the technical advantages of 800VDC arise and where the key engineering challenges still remain.

800 VDC DATA CENTER POWER PATH: FROM UTILITY GRID TO GPU



Utility Grid

- The power journey begins at the utility grid, where electrical power is delivered to the data center at high voltage AC, typically through the transmission network.
- This is the source of bulk facility power. At this stage, the focus is on receiving reliable utility service and establishing the entry point for downstream transformation, switching, and power conversion.

Step Down Transformer

- The incoming high voltage AC is then routed through a step-down transformer, reducing the voltage to a medium voltage level suitable for facility distribution and conversion equipment.
- This stage remains foundational in large-scale power infrastructure. Transmission voltages are efficient for moving power over long distances, but they must be reduced before the energy can be safely and effectively managed within the data center.

Medium Voltage Switchgear

- Once the voltage is reduced, power enters the medium voltage switchgear.
- This equipment provides switching, protection, and isolation for the AC side of the electrical system. It supports system control, fault management, maintenance isolation, and operational continuity. In any large facility, this is a critical layer for safe and reliable power handling before conversion to downstream distribution voltage.

Medium Voltage to 800VDC Rectification

- The next major step is rectification, where medium voltage AC is converted into 800 VDC, often implemented as a bipolar distribution bus.
- This is the defining shift in the architecture. Rather than maintaining AC distribution deeper into the facility and relying on repeated local conversion, the system establishes a high voltage DC backbone much earlier in the power path.
- That change has important implications. It reduces the number of conversion stages between source and load, lowers current for a given power level, and creates a more direct electrical path to high-density compute equipment.

800VDC Switchgear and Protection

- Following rectification, power is routed through 800 VDC switchgear and protection equipment.
- This stage is especially important because DC fault interruption is more complex than AC fault interruption. In AC systems, natural current zero crossing helps simplify interruption. In DC systems, that advantage does not exist, making protection, isolation, and coordination significantly more demanding.
- As a result, high voltage DC protection is not a secondary consideration. It is one of the principal engineering challenges that must be solved for 800 VDC architectures to scale safely and reliably.

800VDC Distribution Bus

- From the DC switchgear, power is distributed through the 800 VDC bus to the rack environment.
- This bus serves as the high-power distribution backbone of the architecture. Whether implemented through busway, cabling, or other high-current interconnect structures, its purpose is to move large amounts of power efficiently across the facility.
- The benefit of this approach becomes more pronounced as power levels increase. For the same delivered power, a higher voltage means lower current. Lower current reduces resistive losses, eases conductor sizing pressures, and helps improve overall distribution efficiency. These advantages are increasingly relevant as the industry moves toward very high-density AI clusters and higher power per rack.

Rack Power Interface

- At the rack-level, the incoming 800 VDC supply is managed by a rack power interface, which may include protection, metering, monitoring, and control functions.
- This stage creates the handoff between facility-level DC distribution and IT equipment-level power delivery. It provides visibility into rack power behavior and supports controlled branching of power into the rack.
- As rack power grows, this interface is likely to become more intelligent and more integrated, serving not only as a protection point but also as a key node for system diagnostics, serviceability, and dynamic power management.

IT Equipment Power Entry and Protection

- Once power enters the IT equipment, it is received by an equipment-level power entry and protection stage.
- This is a more accurate description than the generic term “DC distribution board” when discussing servers, accelerators, and network equipment. In practice, the equipment-level entry stage may include connectors, protection elements, filtering, current sensing, and isolation features that prepare the incoming 800 VDC for conversion inside the chassis.
- This is an important architectural boundary. It is where facility-level high voltage DC becomes controlled equipment-level power.

Front End Isolated DC to DC Conversion

- Within the IT gear, the next step is front-end isolated DC to DC conversion.
- This stage converts the incoming 800 VDC into a lower intermediate bus voltage suitable for use inside the equipment. In many future architectures, this may become one of the most critical innovations in the chain. As higher voltage DC moves closer to the compute load, the performance, efficiency, density, and thermal behavior of this conversion stage will strongly influence the viability of the overall architecture.
- This is also the stage where fault isolation is expected to play a central role in protecting the IT equipment itself. The front-end isolated DC to DC converter is the most logical location for system-level fault isolation within the chassis, serving as the boundary that prevents upstream faults or abnormal conditions from propagating deeper into the server, accelerator, or switch. In practical terms, this stage is likely to become the primary protection barrier designed to prevent damage to downstream IT gear.
- It is also where safety, insulation, EMI control, and system packaging become highly important. Bringing 800 VDC into IT equipment can provide major efficiency benefits, but it also introduces stricter requirements for connector design, creepage and clearance control, fault containment and serviceability.

800VDC Switchgear and Protection

- Following rectification, power is routed through 800 VDC switchgear and protection equipment.
- This stage is especially important because DC fault interruption is more complex than AC fault interruption. In AC systems, natural current zero crossing helps simplify interruption. In DC systems, that advantage does not exist, making protection, isolation, and coordination significantly more demanding.
- As a result, high voltage DC protection is a primary engineering consideration that must be solved for 800 VDC architectures to safely and reliably scale.

Intermediate Bus and Board-Level Power Distribution

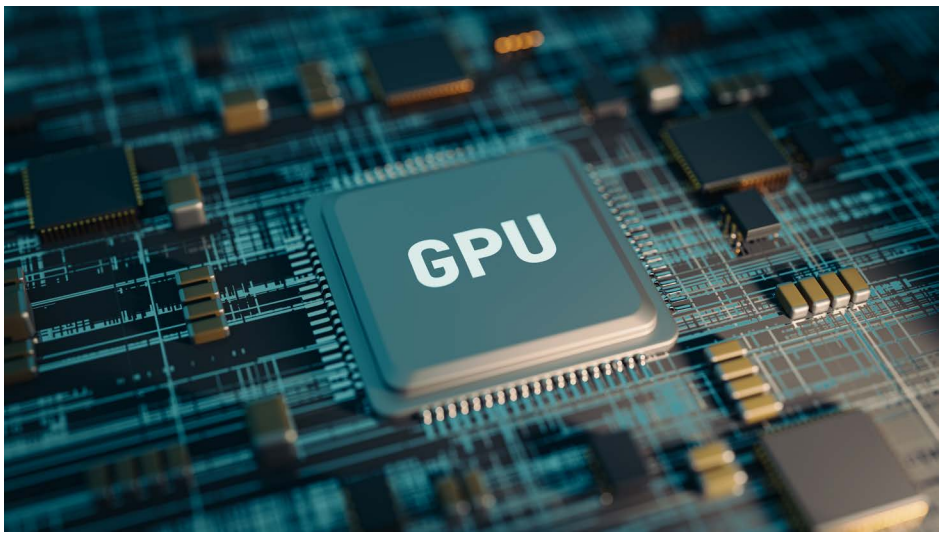
- After front end conversion, power is distributed inside the equipment through an intermediate bus and then across board-level power planes, copper structures, or localized bus bars.
- Inside servers and switches, power is often distributed through carefully engineered power planes, bus bars, or backplane structures designed to support high current, low impedance delivery.
- As compute density continues to increase, this layer of the architecture will likely receive greater attention. Mechanical integration, current sharing, transient performance, thermal design, and connector strategy will all become more important as designers work to move ever larger amounts of power closer to processors and accelerators.

Point-of-Load Voltage Regulation

- The next stage is point-of-load regulation, typically performed by voltage regulator modules located close to the processor, accelerator, or memory device.
- These regulators convert the intermediate bus voltage to tightly controlled low-voltage rails required by advanced semiconductor devices. This stage is essential because CPUs and GPUs operate at very low voltages but very high currents, with fast transient demands and little tolerance for poor regulation.
- Point-of-load regulators are also expected to provide the highest speed fault isolation in the system because they sit immediately adjacent to the load. Their proximity to the processor or accelerator places them in the best position to respond rapidly to fault events before damage can occur. In future high-power AI systems, these regulators will need to isolate quickly while still supporting aggressive load transients.
- As GPUs and other advanced processors switch states rapidly, the power system must distinguish between a true fault condition and a legitimate high-speed load transient caused by the device turning on, turning off, or moving between operating modes. This creates and opportunity for future innovation to improve this discrimination, enabling regulators to respond fast enough to protect the load without creating false trips or unnecessary interruptions during normal operation.
- For AI systems in particular, point-of-load performance is increasingly tied not only to efficiency and transient response, but also to the intelligence and speed of fault handling at the load edge.

GPU or Processor

- The final stage is the GPU or processor itself, where the delivered energy is converted into computational work.
- At this point, power has traveled from utility scale transmission voltage to highly regulated low voltage rails at the silicon level. The objective of the full architecture is to make that journey as efficient, scalable, safe, and controllable as possible.



800VDC INDUSTRY PREDICTIONS

Looking ahead, one of the most important questions is not whether 800VDC can be used at the facility level, but how far into the rack and into the IT equipment it can practically go.

800VDC Will Move Closer to the Compute Load

In the near term, the most likely path is that 800VDC will move directly to the rack-level and then into the equipment chassis, where isolated DC to DC conversion will take place at the front end of the IT gear. Over time, more of the traditional rack-level power shelf functionality may migrate into the server or accelerator enclosure itself. This could produce more integrated architectures in which the boundary between facility power delivery and IT power conversion becomes increasingly blurred.

Sidecar Architectures Will Enable the Transition

At the same time, sidecar architectures are likely to remain important in the transition because they provide a practical boundary for AC to DC conversion, protection, energy storage integration, and serviceability while the ecosystem matures. Until solid state transformers are more established in the ecosystem and within the broader power distribution network, the more likely near term model is one in which busway remains AC and sidecars or row-level power systems continue to perform the AC to DC conversion before feeding power into the IT gear.

Over time, however, as 800VDC distribution becomes more prevalent, the architecture is likely to shift toward direct DC feed from facility busway into the IT rack, reducing the need for a separate sidecar power rack. In that future model, solid state transformers become especially important because they offer a path to convert AC directly into 800VDC, creating a more direct route from the upstream power system into the rack-level DC bus.

Power Distribution Will Become More Intelligent

Rack power interfaces will become more intelligent, with tighter integration of monitoring, protection, telemetry, and control. As power density rises, the rack will increasingly function as an active electrical node rather than a passive distribution point.

Interconnect and Power Conversion Will Continue to Evolve

Connectors, busbars, and equipment power entry technologies will become more important differentiators. As 800VDC moves closer to IT equipment, interconnect design will need to address not only current carrying capability and thermal performance, but also insulation coordination, fault containment, serviceability, and safety. At the same time, the front end isolated DC to DC stage inside the equipment is likely to become a major area of innovation because it sits at the critical boundary between facilit-level high voltage DC and the protected internal power architecture of the IT gear.

Protection Will Become a Critical Differentiator

Point-of-load regulators will need to evolve beyond efficient voltage conversion alone. They will increasingly be expected to deliver ultra fast local protection while distinguishing true electrical faults from legitimate, very fast load transients created by GPUs and other advanced processors. That capability may become one of the most important enablers of stable high power compute platforms as processor current slew rates continue to increase.

Beyond 800VDC

Looking beyond 800VDC, current industry work is looking broadly at the options and may credibly land at low voltage DC range up to 1,500VDC. In practical terms, this suggests a likely sequence for the market: first, 800VDC becomes established for high density AI deployments; next, the industry gains operating experience, matures protection and interconnect standards, and then evaluates whether future hyperscale and AI factory architectures justify moving higher within the LVDC envelope.

TE ROLE IN THE 800VDC INDUSTRY

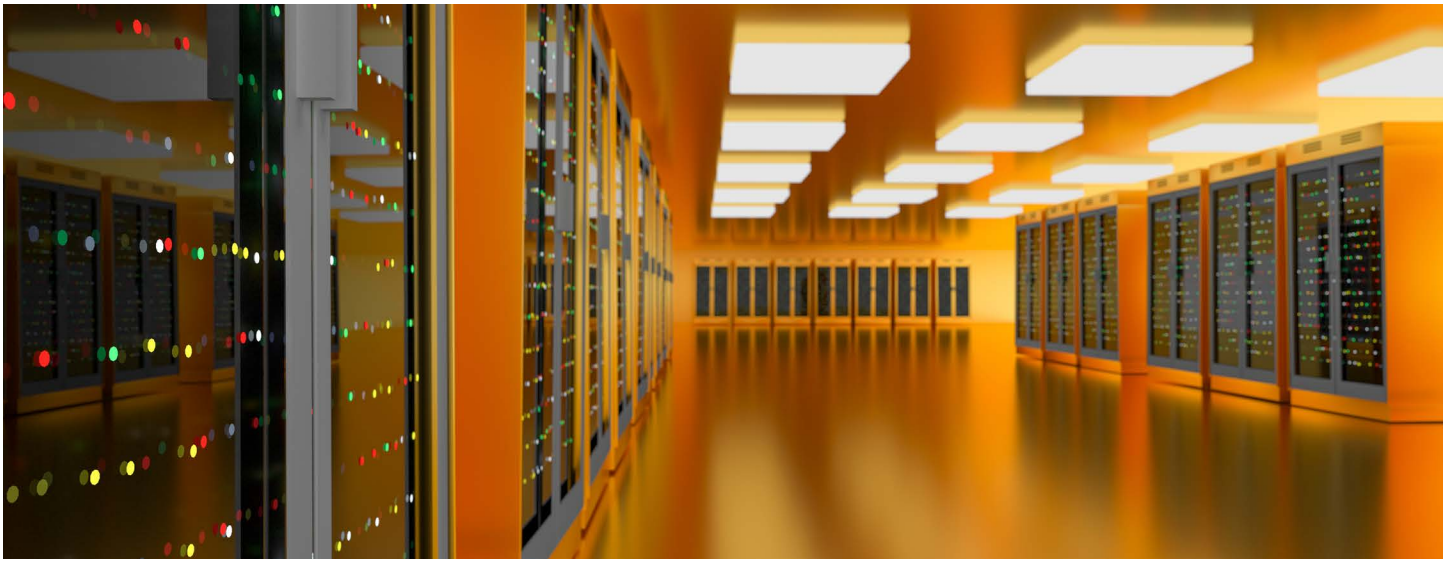
As the industry evaluates 800VDC as a future data center power architecture, the conversation often centers on power conversion, system efficiency, and protection. Equally important, however, is the interconnect and power distribution layer that makes the architecture physically deployable. This is where TE plays a key role.

High voltage DC architectures place new demands on connectors, busbars, cable assemblies, sensing interfaces, and equipment-level power entry. As voltage and power density increase, the interconnect system must do more than carry current. It must also support insulation coordination, creepage and clearance management, fault containment, serviceability, thermal performance, and reliable integration into increasingly compact equipment. These challenges become even more important as high voltage DC moves closer to the rack and potentially into the IT chassis.

TE is well positioned to help enable this transition through technologies that support both the facility-level and equipment-level power path. At the infrastructure layer, TE capabilities in high power interconnects, cable assemblies, and busbar-related power distribution can help address the demands of moving large amounts of power efficiently and safely across the data center. At the rack and equipment boundary, TE can support architectures that require robust power entry, high voltage connector systems, sensing, monitoring, and mechanically integrated distribution solutions.

Inside the equipment, the role of interconnect becomes even more strategic. As front-end isolated DC to DC conversion, intermediate bus distribution, and board-level power delivery evolve, designers will need solutions that reduce impedance, manage heat, support compact packaging, and enable reliable service access. Designers will need to focus on connector solutions and also system-related integration considerations, spanning electrical, mechanical, thermal, and reliability considerations.

TE's value in 800VDC architecture is therefore not limited to any single product family. It lies in enabling the full power distribution chain with interconnect and integration solutions designed for higher voltage, higher current, and more demanding AI infrastructure environments. As data center architecture continues to evolve, TE can help customers bridge the gap between power system theory and deployable system design.



CONCLUSION

800VDC is more than an incremental change in distribution voltage. It represents a broader rethinking of how power can be delivered from the utility entrance to the compute load in future AI infrastructure.

Its promise lies in a more direct and efficient architecture: earlier conversion to DC, lower current for a given power level, fewer conversion stages, and a power path better aligned with high density racks, battery backed systems, and increasingly dynamic compute environments. At the same time, the shift to 800 VDC brings new technical demands, particularly in protection, fault isolation, interconnect design, and equipment-level integration.

The most important challenge is not simply moving higher voltage DC closer to the load. It is doing so in a way that is safe, reliable, and intelligent. Front end DC to DC conversion stages will likely become the primary protection boundary for IT equipment, while point-of-load regulators will need to respond at high speed near the processor itself. As GPU power profiles become more dynamic, the industry will need new approaches that can distinguish true faults from legitimate, high speed load transients without compromising protection or uptime.

For that reason, the future of 800 VDC will be shaped not only by advances in power conversion efficiency, but also by innovation in fault detection, isolation speed, system intelligence, and high-performance interconnects. Success will depend on how well these technologies evolve together.

As the industry pushes toward ever higher rack power and greater infrastructure efficiency, 800 VDC stands out as one of the most credible pathways forward. It offers a strong foundation for the next generation of scalable digital infrastructure, while also pointing toward a future in which power architectures become more integrated, more distributed, and more tightly coupled to the needs of the compute load itself.

Call to Action: Explore 800VDC solutions with TE

The move to 800VDC is not simply a voltage change. It is a system architecture decision that affects conversion strategy, protection philosophy, interconnect design, and long term scalability. As data center operators and equipment designers evaluate how to support higher rack power and more demanding AI workloads, the interconnect and power distribution layer will be a critical part of making these architectures practical.

TE is helping customers address these challenges across the power path, from facility-level distribution to rack power interfaces and equipment-level integration.

Connect with our experts to discuss your 800VDC requirements.

Connect with us to get support on Data Center & AI Solutions

Additional Resources

Empowering AI & Data Center Connectivity
For Next-Gen Speed, Efficiency and Reliability

VISIT THE WEBSITE

Connectivity Solutions Empowering the AI Evolution
Learn how TE is helping to empower the AI evolution with our robust portfolio of interconnect solutions

WATCH THE VIDEO

The Evolution of Data Centers and the Future of Cooling and Connectivity
Explores how data centers are evolving to meet increasing power and cooling demands, with a focuson new design approaches and connectivity solutions.

READ THE TREND PAPER

A More Sustainable Data Center
With the rise of generative AI, data centers have become a growing consumer of electricity.

READ THE TREND INSIGHT

Thermal Performance
Design challenges for increasing computing performance in higher-power data racks while managing thermal performance.

READ THE TREND INSIGHT

Connect With Us

We make it easy to connect with our experts and are ready to provide all the support you need.
Visit **te.com/support** to chat with a Product Information Specialist.

te.com/AI

TE, TE Connectivity, TE connectivity (logo), and EVERY CONNECTION COUNTS are trademarks owned or licensed by the TE Connectivity plc family of companies.
Other product names, logos, and company names mentioned herein may be trademarks of their respective owners.

While TE has made every reasonable effort to ensure the accuracy of the information in this document, TE does not guarantee that it is error-free, nor does TE make any other representation, warranty or guarantee that the information is accurate, complete, correct, reliable or current. TE reserves the right to make any adjustments to the information contained herein at any time without notice. TE EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES REGARDING THE INFORMATION CONTAINED HEREIN, INCLUDING, BUT NOT LIMITED TO, ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. In no event will TE be liable for any direct, indirect, incidental, special or consequential damages arising from or related to recipient's use of the information. It is the sole responsibility of recipient of this information to verify the results of this information using their engineering and product environment. Recipient assumes any and all risks associated with the use of the information. Antenna performance may vary. TE is a component manufacturer, and customer and/or end-user is responsible for all end-use compliance and regulatory requirements. These marketing materials do not constitute an offer or agreement of sale, and TE Connectivity shall be bound by and subject to only those terms and conditions set forth in a written agreement referencing specific products and signed by an authorized representative of TE.

©2026 TE Connectivity. All Rights Reserved.

09-26