enteligent™

800VDC-to-50VDC Power Delivery Architecture

Completing the DC-Native Power Stack for AI-Scale Data Centers

By Sean Burke, CEO | 2026

enteligent.com

Contents

Executive Summary

A previous white paper by Enteligent (800VDC-to-50VDC Power Delivery Architecture for Next-Generation AI & GPU Data Centers) established 800VDC distribution as the definitive power architecture for scaling next-generation AI and GPU data centers. This follow-on white paper extends that analysis deeper into the rack and server layer, demonstrating why the full value of 800VDC high-voltage DC distribution is uniquely realized when paired with next-generation **50VDC rack server power architectures.**

While 800VDC distribution solves the upstream limitations of traditional alternating current (AC) infrastructure, conventional in-rack server power supplies remain a significant inefficiency, cost, and reliability constraint. By replacing per-server AC power supplies with centralized, high-efficiency rack-level direct current (DC) conversion, operators and server manufacturers unlock substantial gains in electrical efficiency, cooling performance, physical density, and total cost of ownership.

This paper examines those gains in practical terms, including specific scenarios for **100 kW AI racks**, a **10 MW AI data center**, and higher loads. Implications for traditional enterprise servers, GPU compute platforms, and next-generation AI training clusters are also evaluated.

1. Completing the Transition from AC to DC

The 800VDC architecture postulated in the previous white paper reframed how power is delivered across the facility by eliminating unnecessary AC conversions, reducing current (the enemy of efficiency), and dramatically improving scalability. However, in most data centers, AC power still reappears at the worst possible location: inside every server chassis.

Traditional AC server power supplies were designed for an era of low-density computing. They are bulky, thermally stressed, and inherently inefficient, particularly under the highly dynamic loads characteristic of graphical processing unit (GPU) and artificial intelligence (AI) computing. Even at best-in-class efficiencies, each AC power supply introduces unnecessary losses, excess heat, and an unnecessary potential failure point directly inside the compute environment.

An 800VDC-to-50VDC rack architecture completes the DC-native model. Rather than at the rack or in the rack, AC-to-DC conversion is centralized at the building level by highly efficient utility-scale inverters. Power is distributed efficiently at high voltage, and a single rack-level conversion delivers native 50VDC directly to servers. Renewable on-premises PV solar generation and battery storage, which are natively DC themselves, are then most efficiently and better physically integrated into the power ecosystem. The result is a simpler, cleaner, and far more scalable power path aligned with how modern computing hardware actually operates.

While 800VDC distribution solves upstream power delivery constraints, conventional server power supplies remain a critical bottleneck. Individual AC power supplies within each server chassis or rack shelf are bulky, thermally stressed, and most operate at 86 to 92% efficiency.

2. Why 50VDC Has Become the Server Standard

As server power levels have increased, the limitations of legacy 12V internal power architectures have become untenable. Supplying tens of kilowatts at 12V requires extreme current levels which drives excessive copper losses, connector heating, and motherboard complexity. In response, the industry has converged on **48–54VDC server backplanes**, with 50VDC emerging as the practical midpoint.

At 50VDC, current is reduced by approximately four times compared to 12V delivery for the same power level. This reduction directly improves efficiency, simplifies mechanical design, and enables higher-performance voltage regulation modules (VRMs) which feed CPUs and GPUs. For AI accelerators that experience rapid step loads and large transient currents, a higher-voltage input rail significantly improves electrical stability and transient response.

This convergence is now reflected in Open Compute Project (OCP) standards and hyperscale server platforms, making 50VDC not an experimental choice, but the de facto foundation for next-generation server design.

3. Rack-Level Conversion with 800-50VDC Power Modules

The conversion bridge between an 800VDC facility backbone and a 50VDC server power bus is an in-rack 800VDC to 50VDC power conversion module. By consolidating power conversion at the rack level, a small number of redundant 800VDC power modules replace dozens of individual AC power supplies, reducing product cost. Fewer power supplies also reduce points of failure, which are typically mitigated by adding more redundant power supplies. Power is delivered to servers via busbar-based distribution, minimizing cabling losses and improving airflow within the rack.

This shared-conversion model mirrors the evolution seen in hyperscale networking and storage, where centralized, modular power shelves have long replaced per-device supplies to improve efficiency and serviceability.

4. Efficiency in Practice: A 100 kW AI Rack Scenario

The efficiency advantages of the 800VDC-to-50VDC architecture become most apparent at high rack power densities. Consider a modern AI training rack operating at **100 kW,** populated with GPU servers drawing highly dynamic loads.

In a traditional AC architecture, power flows through multiple conversion stages — a utility transformer, uninterruptible power supplies (UPSs), power distribution units (PDUs), and finally the server AC power supply — before reaching the on-board voltage regulation modules (VRMs). The compounded efficiency of this chain typically falls between 78 and 85%. At 100 kW of IT load, this means 18 to 28 kW is lost as heat within the electrical path, much of it inside the rack itself.

With an 800VDC-to-50VDC rack architecture, that same 100 kW rack operates at approximately 94 to 95% end-to-end efficiency. Waste power drops to roughly 5 to 6 kW. The net difference — on the order of **15 to 20 kW per rack** — saves up to 20% in power need for the same computing performance, with corresponding reductions in thermal load, cooling requirements, electrical stress, and data-center power needs.

For AI racks that already challenge cooling limits, eliminating this heat is often the difference between feasible and impractical rack densities.

5. Limitations of Next Generation 1MW Sidecar Concepts

The limitations of the AC approach are magnified as AI rack densities exceed 100 kW and move toward power supply concepts that range from 400 kW to 1 MW, where the industry is increasingly exploring "sidecar" power architectures. These systems, offered by vendors like Delta Power, Megmeet, and others, typically place a rack-adjacent cabinet near high-density IT rows to perform AC-to-DC conversion (often to ~800VDC) and, in some cases, incorporate local energy storage for ride-through and transient smoothing. This energy storage is typically a combination of batteries and capacitors. Often, the vendors position it as battery backup, but the reality is the additional energy storage is required to respond to the highly dynamic load conditions.

Sidecar solutions validate a critical industry conclusion: traditional AC distribution cannot scale to megawatt-class racks. While sidecars represent a transitional step toward high-voltage DC, Enteligent's position is that the lowest total cost of ownership and highest operational efficiency are achieved when high-voltage direct current (HVDC) is implemented as a facility-native architecture paired with rack-level 800VDC-to-50VDC conversion, rather than as thousands of distributed AC/DC islands.

Sidecar architectures reduce some upstream conversion losses but multiply the number of conversion assets deployed across a campus. Each sidecar introduces additional fans, capacitors, control electronics, monitoring points, and service requirements. At hyperscale, this results in higher operational complexity, increased maintenance overhead, and greater exposure to component failure. By contrast, an 800VDC backbone concentrates AC-to-DC conversion into fewer, higher-utilization systems while relying on compact, efficient conductor infrastructure and highly efficient rack-level converters to serve the IT load.

Cabling strategy further differentiates the two approaches. Sidecar designs often require substantial copper conductors, either on the AC side feeding the sidecar or on the low-voltage side distributing power from the sidecar to servers. Enteligent's architecture applies voltage where it is most effective: high voltage (800VDC) for distance and low voltage (50VDC) only where short, high-current paths are unavoidable — inside the rack and server chassis. This minimizes copper mass, thermal loss, and installation complexity.

From a mechanical and space perspective, sidecars shift power equipment out of the IT rack but reintroduce it onto the white-space floor as additional cabinets that must be cooled, routed, and maintained. While this architecture acknowledges the physical realities driving interest in megawatt-class racks, it adds the additional cabinet footprint of a second rack space which is required for the sidecar power itself, resulting in the same 500 kW average per rack space anyway. Enteligent's rack-level approach preserves IT rack density while avoiding the proliferation of large rack-adjacent enclosures. Power conversion is integrated cleanly into the rack using compact 2U DC shelves, maintaining consistent airflow and simplifying facility layouts.

Reliability at scale is another key distinction. Some sidecar systems integrate battery modules or large capacitor banks to manage GPU transients, but these additions increase component count and thermal stress at the edge of the infrastructure. Enteligent's architecture achieves resilience through a simpler power chain, elimination of server AC power supplies, and the use of modular, hot-swappable rack-level DC shelves. This enables graceful degradation and maintenance without abrupt server outages.

Critically, Enteligent's approach aligns with the industry's long-term roadmap toward DC-native data center campuses. Sidecars risk becoming stranded assets as operators transition toward centralized rectification, HVDC distribution, and direct integration with renewable generation and battery storage. By making HVDC the facility backbone from the outset, the 800VDC-to-50VDC solution enables a smoother transition to the future data center power requirements without redundant conversion stages or architectural dead ends.

In practical terms, sidecar solutions confirm the applied physics: HVDC is required for megawatt-class racks. Enteligent's 800VDC-to-50VDC rack power architecture builds on that reality to deliver a more scalable, serviceable, and economically durable end state — one that supports both near-term AI deployments and the next decade of compute growth.

6. Quantified Impact Summary: Rack and Facility Scenarios

To ground the architectural discussion in practical outcomes, Table 1 summarizes the electrical, thermal, and operational deltas between traditional AC architectures and an 800VDC-to-50VDC rack power model.

Table 1 — Impact Summary for AI-Scale Deployments

Metric	Traditional AC Architecture	800VDC → 50VDC Architecture	Delta / Benefit
100 kW AI Rack – End-to-End Efficiency	78–85%	94–95%	+10–17 percentage points
100 kW AI Rack – Conversion Losses	18–28 kW	5–6 kW	15–20 kW less waste heat
Rack Cooling Load	High localized PSU heat	Centralized, lower-loss heat	10–15% cooling reduction
Max Practical Rack Density	~40–60 kW typical	100–150+ kW scalable	2–3× density headroom
10 MW Facility – Conversion Losses	~2.2 MW	~0.55 MW	~1.6 MW avoided losses
10 MW Facility – Annual Energy Waste	~19 GWh/year	~4.8 GWh/year	~14 GWh/year saved
Cooling Infrastructure Sizing	Larger chillers/CDUs	Smaller mechanical systems	Lower CapEx + OpEx
Power Components per Rack	20–40 AC PSUs	3–6 DC power shelves	Fewer failures, simpler ops

This comparison highlights that the transition to rack-level DC conversion does not deliver marginal gains—it fundamentally reshapes the electrical and thermal profile of overall data center infrastructure.

7. Efficiency Waterfall: Where the Power Is Won and Lost

One of the clearest ways to visualize the benefit of the 800VDC-to-50VDC architecture is through an efficiency waterfall that tracks power losses from the utility connection to the processor voltage regulators.

In a traditional AC data center, incoming utility power is first stepped down by transformers, then processed through a double-conversion UPS. Double-conversion UPS systems typically operate at 93 to 97% efficiency, while eco-mode operation can approach 98 to 99%, but might not meet all reliability requirements. In parallel, rack power supplies are commonly operated at partial load due to redundancy and capacity headroom, reducing their efficiency by 1 to 3 percentage points relative to nameplate values. From there, it flows through low-voltage AC distribution and rack PDUs before finally being rectified again inside each server by an AC power supply. Each of these stages operates at relatively high efficiency in isolation, yet their cumulative effect is severe. By the time power reaches the VRMs feeding CPUs and GPUs, as much as 15 to 22% of the original energy has already been converted into heat.

In contrast, the 800VDC-to-50VDC architecture collapses this cascade of losses. A single, high-efficiency central rectifier converts AC to DC once. Power remains in DC form throughout the facility, and a rack-level DC converter — operating at greater than 97% efficiency — delivers native 50VDC directly to servers. The result is a dramatically flatter waterfall, where only 5% to 6% of incoming energy is lost before reaching the compute hardware.

For a 100 kW AI rack, this difference is stark. Under an AC model, roughly 20 kW is dissipated across transformers, UPS systems, PDUs, and server power supply units (PSUs). Under the DC model, losses are reduced to less than 6 kW, and that heat is concentrated in fewer, better-cooled locations. When multiplied across hundreds of racks, the waterfall explains why DC-native architectures unlock both higher density and lower operating cost simultaneously.

Efficiency Waterfall: 100 kW Rack

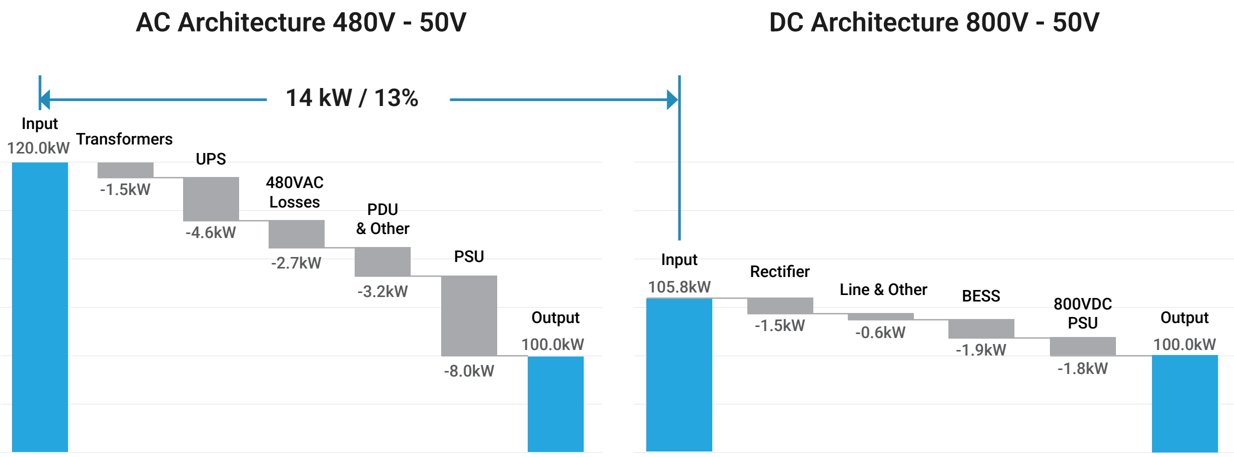

8. Thermal and Cooling Implications

Every watt lost in power conversion must ultimately be removed by the cooling system. Traditional AC server power supplies concentrate heat sources inside each chassis, introducing local hotspots, obstructing cooling airflow, and increasing fan power and failure rates.

By removing AC PSUs entirely, the 800VDC-to-50VDC architecture shifts heat generation out of the server and into highly efficient, centrally-cooled rack modules. In practice, operators gain meaningful reductions in rack inlet temperatures, smoother thermal profiles, and lower overall cooling demand.

In air-cooled environments, this translates into higher allowable rack densities within traditional temperature limits. In liquid-cooled AI racks, lower electrical losses reduce the size and power requirements of pumps, cold plates, and coolant distribution units.

9. Space, Density, and Mechanical Advantages

AC power supplies consume valuable chassis volume and impose rigid mechanical constraints on server design. Removing them enables more compact server layouts, higher compute density per rack unit, and improved airflow paths.

At the rack level, busbar-based DC distribution replaces bulky cable whips, simplifying installation and freeing space for compute hardware. These gains compound rapidly in GPU and AI racks, where physical space is already constrained by accelerators, memory, and cooling hardware.

The result is not merely higher efficiency, but more usable compute per square foot of data center floor space.

10. Facility-Scale Impact: A 10 MW AI Data Center

When scaled across an entire facility, the benefits of rack-level DC conversion become economically decisive. To illustrate the facility-scale impact, consider a 10 MW AI data center built around 100 kW racks."

Under a traditional AC architecture operating at 82% efficiency, approximately 2.2 MW of power is lost to conversion and distribution losses. Under an 800VDC-to-50VDC architecture operating at 94.5% efficiency, losses drop to roughly 0.55 MW. The delta — about 1.6 MW — represents power that no longer needs to be generated, distributed, or cooled.

At typical electricity and cooling costs, this efficiency improvement alone can easily translate into hundreds of thousands of dollars per year in operating expense (OpEx) savings, while also reducing upstream electrical and mechanical infrastructure sizing. Over a 10-year lifecycle, the avoided energy and cooling costs compound into multi-million-dollar savings.

11. Financial Impact of Density and Efficiency

Reducing losses can also enable higher power density. By lifting deliverable IT load from 7.8MW to 9.45MW within the same 10MW facility envelope, the solution improves PUE from ~1.28 to ~1.06. In practical terms, the site can support more kW per rack / per hall and more compute per square foot without any incremental grid, cooling, or "overhead" penalty, unlocking 1,650 kW of additional usable IT capacity (+21%).

This creates clear pathways to value for both real estate owners and operators: for landlords and wholesale/colo providers, the incremental IT capacity can be monetised as additional billable IT kW (and/or a pricing premium for higher-density capability), increasing stabilized revenue and NOI; for compute operators, the same capacity uplift converts into more saleable compute output (or higher utilization headroom) within the existing power envelope. The value uplift is then quantified either by capitalising incremental NOI under contracted kW-based models, or by DCF-ing the incremental gross profit under compute-hour models, often translating to tens of millions in value depending on achievable pricing, occupancy/utilization, incremental OpEx, and cap/discount rates. Secondary benefits include reduced energy/overhead for a fixed IT load, where savings are retained and avoided, or deferred capital expense (CapEx) expansion and schedule risk, to deliver the same compute output.

In a traditional data center, the main constraint is often **space**. In a 10MW AI data center, the constraint shifts to **power density** — how much usable power (and cooling) you can concentrate into each rack and hall, because new generation AI chips consume materially more power, and deliver significantly more compute. That compute allows an AI rack to generate 10x to 15x more revenue than a traditional rack in the same square footage. Therefore, **power density** translates directly into **revenue density**.

Table two summarizes the differences between traditional data center computing and AI-Optimized.

Comparison of Revenue and Performance Metrics (definitions and context at Appendix A)

Metric	Traditional Data Center	AI-Optimized Data Center
Primary Revenue Unit	Rack space and bandwidth	Compute hours (GPU/TPU time)
Typical Power Density	5–10 kW per rack	50–100+ kW per rack
Revenue per Rack	~$1,500 – $3,000 / month	~$20,000 – $50,000+ / month
Customer Lock-in	Low (easy to migrate data)	High (ecosystem/model dependency)
Margin Structure	Stable, utility-like (20-30%)	High risk/High reward (40-60%)

The DC-led approach brings together a unique blend of efficiency, density and economics which ultimately drives revenue density, enabling structural cost reductions and scalable growth in value.

12. Cost and Reliability Implications

From a capital management perspective, centralized rack power conversion within a DC infrastructure reduces the total number of power components deployed across the facility. Server manufacturers benefit from lower bill-of-materials costs, simplified validation, and elimination of region-specific AC power variants. Data center operators benefit from fewer field-replaceable units, reduced spare inventories, and simplified maintenance workflows.

Reliability improves as well. AC power supplies are among the components with the highest failure rates in servers, primarily due to thermal stress and their reliance on electrolytic capacitors. Rack-level DC power shelves, by contrast, are designed for redundancy, hot-swapping, and operation at lower thermal stress, enabling graceful degradation instead of abrupt server outages. The centralization of AC/DC conversion upstream in a large, efficient and robust inverter significantly reduces the likelihood of failure and removes thermal loads from the compute floor.

13. Applicability Across Compute Classes

Modern data centers rarely operate a single workload type. Most facilities support a heterogeneous mix of:

- General-purpose cloud and enterprise servers (CPU-centric, moderate density)
- GPU compute clusters for high-performance computing (HPC), analytics, and rendering
- AI training and inference systems with very high and dynamic power demand
- High-capacity storage systems with steady, long-duration loads

Traditionally, this diversity has been served by a combination of 240 VAC, 208 VAC, and 415/480 VAC distribution, often within the same facility. While functional, this multi-voltage AC approach increases electrical complexity, reduces efficiency, and limits operational flexibility.

A single 800VDC distribution architecture, stepped down locally to 50VDC (or similar rack-level bus), provides a unified electrical platform that simplifies infrastructure while improving performance across all workload types.

For traditional enterprise racks, the 800VDC-to-50VDC architecture delivers immediate gains in efficiency and reliability. For GPU compute racks, it removes copper and thermal bottlenecks that increasingly limit scalability. For AI training clusters exceeding 80 to 100 kW per rack, it is often the only architecture capable of delivering power safely, efficiently, and economically.

As new data centers are built, or older data centers are retrofitted with new power infrastructure, the 800VDC power modules can power all classes of servers – traditional compute, storage, telecom, and GPU/AI servers with a universal power infrastructure.

At the same time, high-voltage DC distribution minimizes current levels for high-density racks while remaining highly efficient for lower-density cloud and storage deployments. Because all equipment ultimately operates internally on low-voltage DC, providing a common high-voltage DC input simplifies server power design and enables standardized rack-level power shelves. The result is a power delivery system that performs efficiently and reliably across the full spectrum of data center workloads, from steady-state storage to highly dynamic AI compute.

14. Conclusion

As AI workloads push data centers beyond the practical limits of traditional AC-based power delivery, incremental improvements are no longer sufficient. The combination of **800VDC facility distribution and rack-level 800VDC-to-50VDC conversion** represents a complete, DC-native power architecture aligned with the realities of modern computing requirements.

For server manufacturers, it simplifies platforms and reduces cost. For data center operators, it unlocks higher density, lower operating expense, and long-term scalability. For utility-scale electricity suppliers, it eases the power they need to supply to the site by over 20%. For AI infrastructure as a whole, it removes a fundamental bottleneck to continued growth.

About Enteligent

This architecture is not a future concept - it is the foundation required to power the next decade of high-density, AI-driven data centers. Enteligent is at the forefront of this shift with its new Rack Core family of DC power modules, designed for this critical role of in-rack conversion of 800VDC supply to 50VDC output, enabling modern high-performance servers. Enteligent *RackCore* DC power modules are air-cooled, rack-ready in a 2U form factor, and available in 20 kW, 40 kW, and 100 kW power capacities.

enteligent™

Building intelligent infrastructure for an electrified world

Enteligent is a Silicon Valley–based developer of Direct Current (DC) power infrastructure for the world's fastest-growing, highest-density loads including data centers and electric vehicle charging. Built on a proprietary DC-DC platform, Enteligent's solutions eliminate unnecessary AC conversions to deliver higher efficiency, lower total cost of ownership, and more scalable power architectures. From DC-powered EV chargers to world-first 800V DC power supplies for AI data centers, Enteligent is powering the shift to DC-native systems.

To learn more, visit enteligent.com.

Contact:

Sean Burke, CEO, Enteligent Inc.
sean.burke@enteligent.com

15. References

High-Voltage DC Power for Data Centers

1. **IEC 62776** – *High Voltage Direct Current (HVDC) Power Systems for Data Centers.* International Electrotechnical Commission.
2. **Electric Power Research Institute (EPRI)** – *DC Power for Data Centers: Assessing Efficiency Gains and Infrastructure Requirements.*
3. **Lawrence Berkeley National Laboratory (LBNL)** – *Energy Efficiency in Data Centers: Analysis of Power Conversion Losses.*
4. **Uptime Institute** – *Next-Generation Power Architectures for High-Density and Hyperscale Facilities.*
5. **IEEE Transactions on Power Electronics** – Multiple peer-reviewed papers on HVDC rectification and conversion efficiency.

Server, GPU, and AI Power Architecture

6. **Open Compute Project (OCP)** – 48V Rack Power Architecture and Design Specifications.
7. **NVIDIA Corporation** – *H100, H200, and DGX Platform Power and Thermal Design Guides.*
8. **AMD** – *Instinct MI300X Accelerator Power and Thermal Specifications.*
9. **Intel Corporation** – *VR13 / VR14 Power Delivery Design Guides.*
10. **Google / IEEE** – *Eliminating Traditional UPS Systems Using DC Power Architectures in Hyperscale Data Centers.*

DC-DC Conversion, VRMs, and Server Power Efficiency

11. **Delta Electronics** – *48V and 50V VRM Efficiency Topologies for AI Servers.*
12. **Megmeet** – *2025 OCP Global Summit, October 22, 2025*
13. **Vicor Corporation** – *Factorized Power Architecture for High-Current GPU Loads.*
14. **Texas Instruments** – *High-Voltage to Low-Voltage DC Conversion for Modern Compute Systems.*
15. **Infineon Technologies** – *High-Efficiency HV-to-LV DC-DC Conversion for Data Center Applications.*

Thermal Management, Cooling, and PUE

16. **ASHRAE TC 9.9** – *Thermal Guidelines for Data Processing Environments.*
17. **The Green Grid** – *Power Usage Effectiveness (PUE): A Comprehensive Examination.*
18. **Microsoft Azure Energy & Sustainability Team** – *Thermal and Electrical Impacts of AI Racks Exceeding 100 kW.*

Electrical Conductors, Copper Usage, and Ampacity

18. **National Electrical Code (NEC)** – Tables 310.16 and 310.17 (Conductor Ampacity).
19. **IEEE Std. 835** – *Standard Power Cable Ampacity Tables.*
20. **Southwire Company** – *THHN/THWN-2 Copper Wire Ampacity and Cost Data.*
21. **Cerrowire Industrial Division** – *Copper Gauge and Ampacity Reference Charts.*

Energy Storage, Renewables, and DC Integration

22. **IEEE 1547** – *Interconnection of Distributed Energy Resources with Electric Power Systems.*
23. **National Renewable Energy Laboratory (NREL)** – *PV String Voltages and Battery Energy Storage DC Bus Architectures.*
24. **Tesla Energy** – *Megapack Technical Data Sheet (HVDC Bus Characteristics).*

Enteligent Technical Documentation

25. **Enteligent Inc.** – *800VDC Power Delivery Architecture and Advantages for Next-Generation AI & GPU Data Centers.*
26. **Enteligent Inc.** – *Rack Core 800VDC-to-50VDC Power Module Product Specification.*
27. **Enteligent Internal Analysis** – Modeled efficiency, thermal, and power-loss scenarios for 100 kW AI racks and 10 MW data center deployments.

16. Appendix A: Financial Metrics — Definitions and Context

The financial metrics presented in this comparison are intended to contextualize the operational and engineering distinctions between traditional and AI-optimized data center architectures. They are indicative rather than prescriptive, reflecting broad industry benchmarks rather than any specific deployment.

Revenue per Rack refers to the monthly recurring revenue generated per rack unit, encompassing power, cooling, connectivity and compute provisioning. In AI-optimized environments, this figure rises substantially due to the premium commanded by high-density GPU compute and the managed services that surround it.

Customer Lock-in describes the relative ease with which a customer can migrate to a competing provider. Traditional colocation customers face low switching costs; AI workload customers are typically embedded within software ecosystems, trained model repositories and proprietary toolchains that make migration operationally complex and commercially unattractive.

Margin Structure reflects earnings before infrastructure debt service and depreciation as a proportion of revenue. The wider range in AI-optimized facilities reflects the capital intensity of the build and the variability of utilization during ramp-up, offset by significantly higher revenue per unit once stabilized.

Power Usage Effectiveness (PUE) is the standard industry metric for measuring data center energy efficiency. It is expressed as the ratio of total facility energy consumption to the energy consumed by IT equipment alone. A PUE of 1.0 would represent perfect efficiency — every watt drawn from the grid reaching the compute hardware directly. In practice, traditional data centers operate at PUE values of 1.5 or above, with best-in-class facilities approaching 1.2. The 800VDC-to-50VDC architecture's reduction in conversion losses directly improves PUE by shrinking the gap between total facility consumption and useful IT load, with modelled outcomes approaching 1.06 in optimized deployments.

All figures reflect market conditions as of 2024–2025 and should be treated as directional indicators for comparative purposes.

enteligent™