Solving Data Center Power Bottlenecks: a Techno-Economic analysis of DC Superconductors

Introduction

The explosion of AI and cloud services is driving unprecedented power demand in hyperscale data centers. Facilities that once drew only tens of megawatts are now entering the gigawatt era [1]. Within the IT space itself, a single rack may require 500kW-1MW of power by 2030 [2] meaning a single data hall could consume 100-400 MW of power. At a system wide scale, conventional AC power on copper cables must now carry enormous currents and are approaching physical limits. For a single MW rack, up to 200 kg of copper busbar would be required and an associated 1 GW center could require up to half a million tons of copper [10]. The status quo is unsustainable: traditional copper cabling/architecture wastes energy as heat (up to 10% loss in efficiency [3]) and forces operators to devote quadratically more floor space and capital to power delivery instead of computing. In effect, the power distribution has become a bottleneck for data center growth.

In this paper, we will further lay out the existing architecture found in hyperscale data centers and the primary pain-points and bottlenecks. Next, we will describe the proposed system-level change (DC on superconductors) and the key deltas associated with the switch. Finally, an economic analysis focused specifically on the superconducting aspect will be presented, with the key assumptions (levers) laid out for transparency.

Boundary definition: Unless explicitly stated otherwise, all CAPEX and OPEX figures refer to an installed feeder system including conductor, cryostat, terminations, cryoplant, instrumentation, protection, commissioning, and EPC. Civil expansion costs are included only when the baseline requires additional routes and the superconducting alternative avoids them. Reported currents and resistive losses are computed at the stated DC bus voltage and feeder configuration (bipolar vs. monopolar; number of parallel feeders; number of terminations). Cryogenic electrical input is derived from an explicit heat budget divided by stated COP.

Existing architecture and painpoints

Conventional data centers distribute power from the grid as three-phase AC, which is converted at the substation to medium voltage AC and is distributed to the pods on campus; it is on this distribution ring that on-site generators are also incorporated. Within a pod, the voltage is stepped down again to ~480 VAC for use in the building. On this line, the UPS system must do a double conversion from AC/DC and DC/AC to charge reserve batteries. Another distribution point sends ~415 VAC power to the racks and PSUs, where the voltage is again stepped down to the operational voltage ~54 VDC. From this architecture comes four main painpoints:

- **Power Distribution Inefficiency:** AC systems require multiple transformer phases, up to 5 times [4]. With transformers operating at ~98% efficiency, this can lead to ~10% losses from transformers alone. Dual-conversion UPS architectures incur inherent efficiency penalties due to AC–DC–AC power conversion, with cumulative losses typically on the order of ~3–4 [5]. Finally, Joule heating in busbars has a dual effect: it directly dissipates electrical power and imposes an additional thermal load on the building HVAC system, which can reduce system efficiency by up to ~10%[6].
- **Density and Footprint:** The combination of sharply increasing per-rack power requirements, reliance on three-phase AC, and low-voltage distribution necessitates dozens of large-gauge copper cables/busways. While this approach was feasible in prior generations of data centers, the associated capital expenditures for physical space, routing complexity, and cabling infrastructure are becoming a major limiting factor. As a result, new builds must allocate an increasing share of capital to supporting power infrastructure rather than to compute itself.
- **Architecture Complexity:** With multiple conversion stages, extensive distribution lines, and redundancy requirements, system complexity scales non-linearly with total power draw. Each additional node introduces a new potential point of failure, reducing overall system reliability. The reliance on massive parallel feeders further complicates routing and physical layout. At projected rack power levels, this translates into hundreds of kilograms of copper per rack over distances of hundreds of meters. This quadratic growth in cabling also consumes an increasing share of floor space dedicated solely to power distribution."
- **Integration of Renewables:** DC-native renewable sources such as solar must undergo power conversion to integrate with AC-based microgrids, introducing an additional layer of complexity. More broadly, the requirement for continuous power availability necessitates pairing renewables with battery storage systems, which themselves require multiple conversion stages for integration.

The Case for DC Power Architecture

Switching to DC distribution addresses many of these inefficiencies by supplying the IT space directly and eliminating redundant AC/DC conversion stages. The benefits of this approach have already been demonstrated; for example, a Swiss data center implementing a 380 VDC architecture reported approximately 10% higher overall efficiency than an equivalent AC design, along with up to 20% grid-to-chip energy savings driven by combined electrical and cooling reductions. [7]. Capital costs for the DC facility were approximately 15% lower than those of a traditional AC design, and the system required roughly 25% less physical space due to a simplified power chain [7]. This reduction in complexity results from eliminating intermediate transformers and redundant rectifiers, as DC architectures require fewer power conversion and conditioning stages. In short, DC power delivery is inherently more efficient and reliable for IT loads because it streamlines the path from the grid to the server.

DC distribution reduces integration requirements for renewable energy systems. Because batteries and solar photovoltaics inherently produce DC, a DC data center can incorporate these sources with minimal power conversion. For example, Google transitioned to 48 VDC rack-level power with lithium-ion batteries at each rack, eliminating intermediate AC/DC conversion stages [8]. More broadly, DC power architectures are naturally compatible with DC power sources, as they avoid unnecessary inverters and rectifiers, resulting in a system that is more conducive to renewable integration.

The industry is increasingly converging around DC power standards. The Open Compute Project and major hyperscalers are driving adoption of high-voltage DC architectures to meet emerging power density requirements. Google, Meta, NVIDIA, and others are collaborating around 400–800 VDC distribution, drawing on advances from the electric-vehicle sector, which already operates at these voltage levels. NVIDIA has announced plans to support 800 VDC architectures by 2027 to enable next-generation 500 kW–1 MW racks, replacing legacy low-voltage systems [9]. By converting 13.8 kVAC directly to 800 VDC at the facility level, such architectures eliminate multiple conversion stages and can improve end-to-end efficiency by up to ~5% [10]. In addition, DC distribution directly addresses the physical limits of AC systems; maintaining a 1 MW rack at 54 V would require nearly 200 kg of copper for a single rack [10]. While public designs and vendor roadmaps remain heterogeneous and early-stage, the implications of this shift are difficult to ignore.

Superconducting Cables: DC's Complimentary Tech

For the purpose of this paper, we define a superconducting cabling system as the following: A superconducting power system includes the primary conductor, the necessary cryostat to keep the conductor at operating temperatures, the current feedthroughs that interface room-temperature copper conductors with cryogenic superconductors, sensing hardware/software to prevent and mitigate off-nominal events such as quenches or coolant loss, high current safety features such as solid-state breakers/shunts/resistors, and finally the primary cryoplant. This system is intended to be placed within the high-current sections of a DC data-center **Fig 1.**

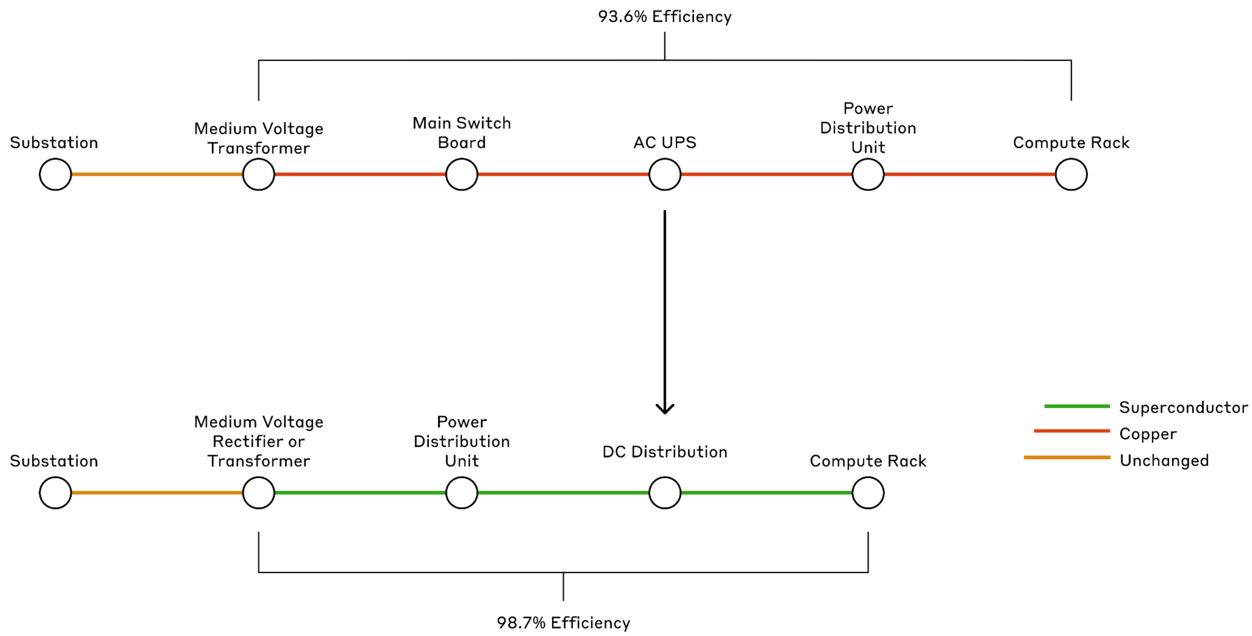

Fig 1: A comparison of traditional AC and future DC infrastructures, with the scope of superconductors included.

While DC distribution addresses conversion losses, high-temperature superconducting (HTS) cables tackle the complementary challenge: resistive losses and the physical limitations of copper conductors. An HTS cable can carry current with effectively zero electrical resistance under operating conditions, enabling near-lossless power transmission over facility-scale distances [6]. As a result, a superconducting link can deliver power with negligible voltage drop and without Joule heating. This represents a step-change for data-center power infrastructure, where copper distribution can dissipate up to ~10% of delivered power as heat [6]. By eliminating I^2R losses, superconductors both improve electrical efficiency and materially reduce downstream cooling requirements.

Another major advantage lies in the current density achievable with superconducting materials. Second-generation HTS tapes can carry hundreds of amperes within a single tape measuring approximately 10 mm × 0.1 mm [11]. At the cable level, HTS systems can conduct orders of magnitude more current than copper for the same cross-sectional area. For example, delivering 300 MW at 33 kV would typically require dozens of large copper conductors, whereas a single superconducting cable with a cross-section on the order of ~250 mm could accommodate the same load. Even at lower voltages, properly engineered HTS cables can support currents in the tens of kiloamperes. This dramatic increase in current density substantially reduces cable footprint, allowing a single superconducting link to replace dozens of parallel copper runs. Industry estimates indicate that superconducting solutions can reduce substation land use and cable-tray requirements by approximately 30%, while eliminating extensive civil works otherwise required to route hundreds of copper cables. In indoor busway applications, fewer conductors translate to cleaner layouts and simplified installation.

Superconducting power cables operate at cryogenic temperatures and are housed within vacuum-insulated cryostats that minimize heat ingress. When implemented in concentric, coaxial, or bipolar geometries, these systems can also substantially suppress external magnetic fields, with the degree of cancellation dependent on geometry and installation. Unlike conventional busbars, HTS cables do not function as large resistive heaters or electromagnets and can therefore be routed closer to IT equipment without imposing additional HVAC or EMI constraints. The primary thermal load arises from the refrigeration required to maintain cryogenic operation. Cryostat heat leak is typically on the order of sub-watt per meter to a few watts per meter, depending on cable diameter, operating temperature, and installation conditions. At very high current levels (25–100 kA class), termination and current-lead losses can become the dominant contributors to the thermal budget. Accordingly, any techno-economic assessment must explicitly account for (i) cryostat heat leak, (ii) termination heat loads that scale with current, and (iii) the assumed coefficient of performance (COP) of the refrigeration system at the operating temperature.

Protection and fault behavior remain non-trivial for 400 to 800 VDC distribution operating at tens of kiloamperes. Although superconductors can transition to a resistive state during a quench, this behavior is not a substitute for a purpose-designed fault current limiter and can be destructive without engineered shunts, fast quench detection, coordinated DC protection, and well-defined fault-clearing strategies. Accordingly, this paper treats protection architecture, including DC breakers, grounding schemes, and quench response, as a go or no go constraint that must be fully specified for any deployment.

Solving the Key Pain Points with DC Superconductors

By combining DC architectures with superconducting cables, next-generation power delivery solutions directly address the major pain points of AC copper systems:

- **Eliminating Resistive Losses:** Superconducting DC cables virtually eliminate Joule heating losses in power distribution. In large data centers, conventional copper feeders dissipate several percent of delivered energy as heat, whereas HTS lines operate with near-zero resistive loss under steady-state conditions. The result is higher overall system efficiency and lower operating costs. For facility operators, this translates into reduced wasted electricity and a corresponding decrease in cooling demand, since power distribution no longer injects heat into the building.
- **Power Density and Compact Design:** A superconducting DC link can deliver megawatts of power through a single cable, offering on the order of 10x the current capacity and roughly $\frac{1}{20}$ of the physical footprint of an equivalent copper solution, which would otherwise require dozens of parallel conductors. This directly addresses the space and weight limitations of copper distribution, where cabling can exceed 200 kilograms per rack. As a result, power trunks and busways can be made significantly slimmer, reclaiming valuable floor space for IT equipment and reducing overall build-out costs.

- **No Voltage Drop Constraints:** In copper-based systems, low-voltage DC distribution is distance-limited, as voltage drop forces power rooms and battery banks to be placed close to the load. Superconducting DC cables remove this constraint by enabling stable DC delivery over long distances with negligible voltage drop. As a result, data center designers gain significant flexibility in equipment placement. Conversion units, battery UPS rooms, and rectifiers can be centralized or sited wherever convenient, including off-site locations, without concern that distant racks will experience voltage sag [13]. In principle, a single centralized HVDC rectifier, with appropriate redundancy, could uniformly feed all rows through HTS distribution, a topology that is impractical with copper conductors. This ability to deliver stable DC over distance also simplifies scaling, as new racks or pods can be added to an existing distribution backbone without requiring localized power rewiring.
- **Reduced Conversion and Cooling Equipment:** A DC superconducting architecture reduces the total number of components in the power chain. Under conventional AC distribution, each rack may require multiple power supply units or dedicated power distribution units, and at extreme power densities, power shelves can occupy up to 64U of rack space [10]. Google's deployment of distributed rack-level batteries demonstrated simpler power distribution and improved reliability. When extended to the facility scale, this approach enables a single centralized conversion stage that feeds both the distribution network and energy storage, with only a final DC to DC conversion occurring at the rack level. In addition, because superconducting cables do not introduce resistive heating into the space, cooling systems can be more tightly aligned with actual IT heat loads. This reduction in parasitic cooling demand contributes to improved power usage effectiveness. Industry estimates suggest that combining high-voltage DC distribution with efficient power delivery can reduce total cost of ownership by approximately 30% through combined power and cooling savings [10].
- **Renewable Integration:** A DC superconducting architecture enables more direct integration of renewable and backup energy sources within data centers. As discussed, DC power systems natively accommodate solar photovoltaics, fuel cells, and battery storage without the conversion losses associated with AC inversion. Superconducting DC links can connect large on-site battery installations or solar arrays directly to the facility's internal DC bus with negligible transmission loss. This capability allows facilities to ride through grid disturbances more effectively, reduce peak demand, and optimize energy dispatch. More broadly, by reducing distribution losses and parasitic cooling loads, a superconducting DC system supports sustainability objectives by lowering total energy consumption and associated emissions, while giving operators greater control over when and how power is sourced and utilized.

In summary, transitioning from conventional AC copper architectures to DC distribution on superconductors directly addresses the principal bottlenecks facing modern high-power data centers. This approach improves efficiency, increases achievable power density, reduces space and weight requirements, and enhances overall system reliability. The factors that historically constrained adoption of DC and superconducting technologies, including cost, complexity, and immature standards, are rapidly diminishing. Industry momentum continues to build, with hyperscalers, equipment vendors, and power providers converging around DC-based architectures.

In parallel, superconducting cable technology has matured substantially, with costs declining significantly over the past decade. While prior literature often frames adoption thresholds in terms of conductor-level cost metrics such as dollars per kiloampere-meter, turnkey deployed systems cannot be accurately represented by tape cost alone. Accordingly, this paper evaluates superconducting economics within an explicit bill-of-materials boundary and treats $/kA-m as a sensitivity parameter rather than a point forecast. Early pilot deployments further demonstrate viability. For example, a recent demonstration delivered 3 MW through a single 800 VDC superconducting cable in a physically representative data center environment, achieving near-zero transmission loss and safely integrating within standard facility constraints [15]. The system demonstrated the ability to carry 10x more power over 5x the distance, while occupying a 20x smaller physical footprint required by copper-based solutions.

Direct Comparison with Existing Technology

To concretize the benefits of a DC superconducting system, key metrics are compared against a proposed DC architecture of DC on copper systems. Specifically, the pre-fab, ~75 MW pods that will make up future hyperscale data centers are compared.

	Copper	Superconductor	
DC Voltage Level	800		Volts
Conductor assembly current density	0.208	50	kA/cm^2
Nominal total power	75		MW
Conductor assembly cross section (total)	903	9.84	cm^2
Conductor assembly mass per meter	809	15	kg/m
One way length	200		m
Total electrical resistance	0.171		mΩ
Temperature rise above ambient	20		K
Cable temperature at nominal current	323	65	K
Voltage drop at nominal current	16		V
Electrical power for cooling		243.9	kW
CAPEX	9.2	8.5	$MM
Operating energy (losses + machinery)	1,500	243.9	kW
Electricity price	40		$/MWh
OPEX annualized	525,600	85,462	$

Table 1: Comparison between a Copper and Superconducting implementation for a 200m distribution line within a 75 MW compute block. Cooling power is computed from the heat budget and COP assumptions (Appendix I)

As shown in the table above, superconducting feeders primarily address space, routing, and voltage-drop constraints at very high current. Whether they improve OPEX depends strongly on the full cryogenic heat budget and the refrigeration COP, and must be evaluated on a case-specific basis. Pair these improvements with the potentially lower CAPEX and relatively high TIR from recent pilots [15], and traditional arguments against the technology begin to significantly diminish.

Techno-Economic Feasibility and Associated Levers

To determine the economic feasibility of DC superconductors within this environment, we must first quantify the existence and magnitude of the power-distribution bottleneck in AI-scale data centers, specifically the economic envelope in which 800 VDC superconducting distribution becomes a clear winner versus copper under explicit cost and thermal constraints.

Key messages

- Data center electricity consumption is projected to more than double from ~415 TWh (2024) to ~945 TWh (2030), driven materially by accelerated (AI) servers and associated infrastructure growth [16]
- Industry momentum is converging on higher-voltage DC inside the data hall: 400 VDC for up to 1 MW IT racks [8] and 800 VDC ecosystem for AI factories [10]
- A DC architecture reduces conversion stages and is supported by field results such as ABB/Green's 380 VDC facility reporting higher efficiency and lower capex/space versus comparable AC designs.
- At 800 VDC, copper distribution is constrained primarily by current density, voltage-drop design limits, routing congestion, and the knock-on cost of heat removal. Superconducting feeders remove Joule heating loss and voltage-drop constraints while delivering very high ampacity in a compact, EMI-contained cryostat.

Product concept and explicit economic constraints

This feasibility envelope is anchored to two input constraints provided for the proposed superconducing solution:

- **Superconducting cable variable price:** $453.33 per KA-m (i.e., $566.67 per MW-m at 800 VDC)
- **Cryogenic plant electrical draw:** 243.9 kW
- **Cryostat thermal leak capability:** 0.3 W/m (cold side)
- **Termination thermal leak capability:** 45 W/kA

- **Bus voltage:** 800 VDC distribution

Techno-economic breakpoints: when superconducting distribution wins

Superconducting distribution becomes a clear winner where copper is forced into extreme conductor areas to meet voltage-drop/temperature limits, or into excessive parallel runs that create routing and installation complexity. A useful engineering proxy is voltage-drop design: if a feeder must be designed for ~1% drop at 800 VDC, the required copper cross-section scales with power and distance; total copper mass then scales approximately with the square of the route length. Superconducting cable cost, by contrast, scales linearly with length.

Feeder (MW)	Length (m)	Cu mass pair at 1% drop (t)	SC cable cost
1	200	3.9	$113,334
1	500	24.1	$283,335
2	200	7.7	$226,668
2	500	48.3	$566,670
5	200	19.3	$566,670
5	500	120.7	$1,416,675
10	200	38.6	$1,133,340
10	500	241.4	$2,833,350

Table 2: A comparison between copper weight and superconductor cost at different lengths and power throughputs, not including cryogenerator

At multi-megawatt ratings and 500 m class routes, copper mass (and therefore material, support, and installation burden) becomes extreme, while superconducting cable cost remains linear in length. This is the feasibility wedge: constrained distribution feeders (multi-MW) and/or long internal routes, especially where routing congestion, busway weight, and voltage-drop placements constraints drive expensive copper solutions.

Deployment wedge and blockers that a superconducting 800 VDC system can remove

- Wedge applications: 200-500+ MW hyperscale and an AI-factory sites; campus-scale builds; long-route distribution corridors; and upgrades where copper routing or busway weight is a schedule blocker
- Blocked by copper: voltage drop forces power rooms and battery/rectifier placements close to load; superconducting DC removes distance limits and enables centralized conversion layouts.
- Blocked by copper mass: 1 MW class racks at low voltage imply prohibitive copper mass; higher-voltage DC reduces current, and superconducting feeders further compress distribution footprint.
- Blocked by reliability/maintenance: fewer conversion stages and fewer parallel conductors reduces failure points and field maintenance complexity; superconductors also offer inherent fault-limiting behaviour via quench dynamics.

TAM, SAM, and SOM

We size the opportunity around incremental data center power demand growth to 2030 and the share likely to be served by prescale operators. The IEA projects an increase from ~415 TWh (2024) to ~945 TWh (2030). This implies ~60.5 GW of additional average load. Assuming a 70% utilization-to-nameplate factor yields ~86 GW of new nameplate capacity; applying a 60% large-site/hyperscale share yields ~52 GW addressable for high-power HVDC backbone distribution.

We express addressable demand in MW-m of constrained high-current corridors. For simplicity, we use an effective corridor length per delivered MW (L_{eff}), intended to represent the subset of feeders where conduit/civil constraints bind (not total campus wiring).

$$TAM\ (\$) = (Incremental\ qualified\ MW) \times (L_{eff} \times m/MW) \times (\$/MWm)$$

Using this equation, we can estimate the TAM for two lengths:

- $L_{eff} = 200m$: **$5.9B**
- $L_{eff} = 500m$: **$14.7B**

Due to undetermined specifics regarding effective lengths encompassing both brownfield and greenfield, a range of 6 - 15B will be used rather than a single figure.

For the SAM, apply a conservative near-term adoption filter: only a fraction of new builds and major expansions will adopt 800 VDC internal distribution immediately. If 10-30% of the large-site segment adopts HVDC-superconducting backbones in the 2026-2030 window, the serviceable market is approximately **$0.59B-$4.4B.**

Lastly, the SOM can be estimated utilizing a project-based approach. Using a pricing scheme of $113,000/MW, a single 200 MW data center project could equate to ~**22.6M**. This means that a conservative plan like 1 pilot deployment then 1-2 full sites by 2030 can yield **~22.6M-90.4M** cumulative revenue. When considering the trend for >500 MW data centers, however, this figure could potentially double.

Conclusion

As computing continues to scale, the old approaches to power distribution are hitting a wall. AC on copper systems waste too much energy and space to support next-generation racks. Superconducing DC feeders, while not a general-purpose replacement for all copper instances within a data center today, have the opportunity to mitigate high conduit counts, unnecessary civil expansion, thermal derating, and schedule risks given appropriate application spaces.

Crucially, this isn't just a lab experiment or wishful thinking. Major operators and vendors are already moving toward DC distribution. In parallel, superconducting cable solutions are reaching commercial readiness, with companies like Nexans, VEIR, and VESC showcasing successful pilot programs in their respective fields. Early adopter sites will enjoy higher efficiency and density, gaining a competitive edge in performance per watt. Over time, as volumes rise and costs further decline, DC superconducting power distribution is poised to become a new norm for hyper-scale facilities.

Appendix

I.

Cable heat leak	0.3	W/m
Termination efficiency	45	W/kA
Thermal load (200m)	5.69	kW
COP at operating temperature	0.07	W/W
Resulting electrical input	81.3	kW
Electrical input for 75 MW	243.9	kW

The dominant cold-load term is termination heat, not distributed cryostat heat leak; our electrical input figures assume four terminations (bipolar, both ends).

References

(1) OpenAI. (2025, July 22). *Stargate advances with 4.5 GW partnership with Oracle*. https://openai.com/index/stargate-advances-with-partnership-with-oracle/

(2) Carlini, S. (2025, October 16). *The 1 MW AI IT rack is coming, and it needs 800 VDC power*. Schneider Electric Blog. https://blog.se.com/datacenter/2025/10/16/the-1-mw-ai-it-rack-is-coming-and-it-needs-800-vdc-power/

(3) ENERGY STAR. https://www.energystar.gov/products/data_center_equipment/16-more-ways-cut-energy-waste-data-center/reduce-energy-losses-power-distribution-units-pdus/

(4) Digi-Key Electronics. (2011, July 7). *Evaluating DC vs. AC power in data centers*. Digi-Key. https://www.digikey.com/en/articles/evaluating-dc-vs-ac-power-in-data-centers

(5) DataCenterDynamics. (2021, May 14). *New, efficient UPS modes: Too "risky" for many.* Data Center Dynamics. https://www.datacenterdynamics.com/en/opinions/new-efficient-ups-modes-too-risky-for-many/

(6) Nexans. (2025). *Data centers of the future powered by superconducting cables for scalable and sustainable growth*. https://www.nexans.com/perspective/data-centers-of-the-future-powered-by-superconducting-cables-for-scalable-and-sustainable-growth/

(7) ABB. (2012, May 30). *World's most powerful DC data center online*. https://new.abb.com/news/detail/12816/worlds-most-powerful-dc-data-center-online

(8) Peabody, C., & Sizikov, G. (2025, February 25). *How we got to 100 million cells in our global Li-ion rack battery fleet*. Google Cloud Blog. https://cloud.google.com/blog/topics/systems/100-million-li-ion-cells-in-google-data-centers/

(9) Peabody, C., & Huffman, A. (2025, April 29). *Enabling 1 MW IT racks and liquid cooling at OCP EMEA Summit*. Google Cloud Blog. https://cloud.google.com/blog/topics/systems/enabling-1-mw-it-racks-and-liquid-cooling-at-ocp-emea-summit/

(10) Blake, M., Hsu, M., Goldwasser, I., Petty, H., & Huntington, J. (2025, May 20). *NVIDIA 800 VDC architecture will power the next generation of AI factories*. NVIDIA Developer Blog. https://developer.nvidia.com/blog/nvidia-800-v-hvdc-architecture-will-power-the-next-generation-of-ai-factories/

(11) Lue, J.W. & Gouge, Michael & Duckworth, Robert. (2005). Over-current testing of HTS tapes. Applied Superconductivity, IEEE Transactions on. 15. 1835 - 1838. 10.1109/TASC.2005.849306.

(12) Hajdasz, S., Kempski, A., Solak, K., Marc, M., Rusinski, J., & Szczesniak, P. (2023). *Critical current degradation in HTS tapes for superconducting fault current limiter under repeated overcurrent*. **Applied Sciences, 13**(7), 4323. https://doi.org/10.3390/app13074323

(13) American Superconductor Corporation. (2012). *Superconductor LVDC data center power cables* (LVDC_BRO_0412) [Brochure]. https://www.amsc.com/wp-content/uploads/LVDC_BRO_0412_web.pdf

(14) Wang, B. (2023, September 3). *Markets and future costs of superconducting magnets and superconducting wire*. *Next Big Future*. https://www.nextbigfuture.com/2023/09/markets-and-future-costs-of-superconducting-magnets-and-superconducting-wire.html

(15)VEIR. (2025, November 12). *VEIR successfully demonstrates 3-megawatt superconducting power delivery for AI data centers, eliminating critical growth bottleneck*. https://veir.com/veir-successfully-demonstrates-3-megawatt-superconducting-power-delivery-for-ai-data-centers-eleminating-critical-growth-bottleneck/

(16) Reuters. (2024, December 20). *U.S. data-center power use could nearly triple by 2028, DOE-backed report says*. Reuters. https://www.reuters.com/business/energy/us-data-center-power-use-could-nearly-triple-by-2028-doe-backed-report-says-2024-12-20/

Prepared by: Aether Transmission Systems
Contributors: Nathan Groth, Ishaan Panda