Amphenol
NETWORK SOLUTIONS™

Data Center Power Outlook for 2025-2030

White Paper

Written by: Brad Hawkins, *Data Center Systems Engineer*

November, 2025

Contents

Tables

Figures

1. Executive Summary

AI and Machine Learning have driven record data center spending and development as of mid-2025

This paper examines power trends for the next five years, assuming AI/ML remains the primary growth driver.

This paper evaluates 54VDC, AC, and HVDC power architectures, finding 400/800V HVDC set for rapid adoption especially in ultra-high-power AI applications. While all three architectures will grow, HVDC will see the most focus in standards development.

Key Market Segments

For the purposes of this paper, we consider the following market segments:

Type of Data Center	Location	Size	Use - Cases
Telco Edge	Cell Towers, COs	Micro - 1-20 Racks	Telco Services (BBU, RRH etc) Edge Computing, CDN Offload, ULL Services
Metro Edge	Near Cities	Small - 10-100 Racks	Medium Latency Storage, CDN Servers, Enterprise SaaS
Retail Colo	Inside/Near Cities	Medium - 100+ Racks	Enterprise SaaS, Small -Scale IT Services, Banks. Rented By-The-Rack
Wholesale Colo	Outside/In Between Cities	Large - 100's of Racks	Large - Scale Enterprise, Cloud Providers
Hyperscale Traditional	Rural - Cheap Land/ Power	Huge - 1000+ Racks	Cloud Services/Hosting, Storage
AI/ML Inference	Global, Close to Users	Small to Medium	AI Agents/Chat, Enterprise AI, Autonomous Technology, AR/VR, AI-as-a-Service (AlaaS), Generative AI
AI/ML Training	Rural - Cheap Land/ Power	Huge, Limited by Power	Internal Model Training, AlaaS, Some Generative AI Offloading

Figure 1 - Data Center Key Market Segments

Most figures and trends are separated into three categories – Edge, Colocation/Enterprise, and Hyperscale/AI. While these segments have their own intricacies in architecture and power-system design, the categories give a generalized sense of potential architecture changes in each group.

2. Data Center Power Market Outlook

Key Assumptions

Given the current state of the industry, we outline the core assumptions underlying this paper. The assumptions outlined below serve as the guiding force behind data center development in mid-2025, and many of the values in this paper are generalized based on several industry projections. All conclusions drawn in this paper should keep that in mind.

AI/ML Drives Spending

The first assumption is that AI/ML drives the majority of data center spend over the next five years. Given the rapid expansion of AI/ML data centers built over the last couple of years, this is not a surprise. The scale of projected spending is still immense.

In April 2025, McKinsey & Company published a paper quantifying its projected industry data center spending from 2025 to 2030. From 2025 to 2030, they project $6.6 trillion in required data center spend. Of this $6.6T, $5.2T is expected for AI/ML, while only $1.4T is projected for traditional data center spend (web hosting, emails, file storage, etc.).

As shown in Figure 2 , this $6.6T is divided into three categories: Builders (Construction, Infrastructure, etc.), Energizers (Power plants, transformers, PDUs, etc.), and Technology (GPUs, Switches, fiber optics, etc.). In all three categories, the AI/ML spend is projected to drastically outweigh traditional spend.

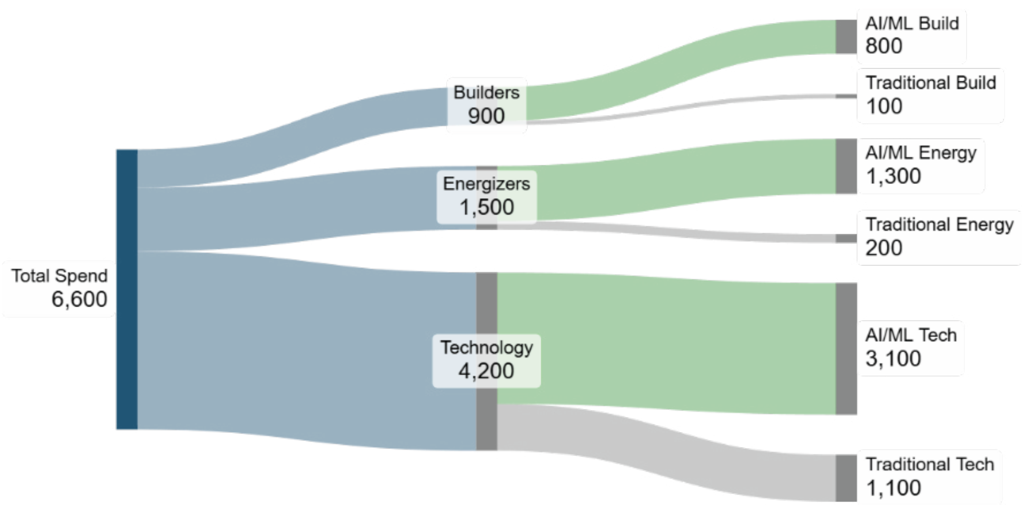

Figure 2 – McKinsey & Company Projected Data Center Spend

The majority of the projected spend will be allocated to technology, which isn't a surprise, given the high cost of AI chips.

It is estimated that a single rack built around NVIDIA's Blackwell CPU+GPU chip may cost in the range of $2-$3 million, although it remains to be seen if AMD's new AI chips will cause any downward price pressure.

Builders receive the smallest slice of the pie; however, it is still a substantial $900 billion. AI/ML will require both massive rural training data centers and smaller metro inference data centers. This number reflects the massive scale of the AI/ML buildout, but hyper-scalers are trying to minimize this non-revenue generating spend over the next few years.

Evidence of this minimization can be seen in recent papers around Meta's rapid ramp-up of their data center builds. The company is building new data centers in ultra-light tent-like structures, with minimal backup redundancy such as diesel generators. For AI and ML builds, legacy standards often do not apply, particularly in large training data centers.

The third category of spend is the most relevant to this paper – $1.5 trillion in expected spend on energizer products over the next five years. AI/ML is a series of chips solving trillions of math problems every second, which requires immense power. This paper will delve into the growth of the data center power industry, as well as the causes and effects of projected power architecture changes.

Enterprise/Consumer Embraces AI/ML

The McKinsey projection of $6.6 trillion in data center spend relies on sustained momentum in AI/ML. While past hype cycles have come and gone, current rapid adoption suggests continued growth in both consumer and enterprise markets is likely.

This paper reflects the belief in sustained growth of AI/ML; however, it should be noted that there is some risk that companies are overspending relative to public adoption. Similarly, there is a chance that the industry will continue to see breakthroughs, and adoption will ramp up. All figures and projections in this paper should be considered with that caveat in mind.

No Major Macro Economic Disruptions

The final key assumption in this paper is that the industry does not see any major macro-economic disruption. Over the past five years, we have witnessed a pandemic, a large-scale war in Eastern Europe, trade tensions, and significant supply chain disruptions. Despite that, the data center market has continued to grow and innovate. Given this, the threshold for major disruption is high. This paper assumes no severe recessions, wars, or trade breakdowns.

General Power Density Trends

According to the McKinsey paper, under continued momentum, it is projected that global yearly power demand for data centers will grow 3.5 times by 2030. The paper indicates that by 2030, the market will see approximately 31GW of additional AI capacity added per year. To put this number into perspective, the average nuclear power plant in the United States has just over 1 GW of capacity. Grand Coulee Dam, the largest power station in the US, has a max production capacity of 6.8GW.

Power density is the total power usage of all equipment in a rack, or more generally, or the amount of heat that must be dissipated. For example, a 10KW rack would require removing 10KW of heat, similar to several space heaters at 1.5KW each.

An estimated projection of rack power densities for greenfield builds is shown in Figure 3, separated by the data center market. Sources show a variety of current and projected rack densities and the figures represent an average of these numbers.

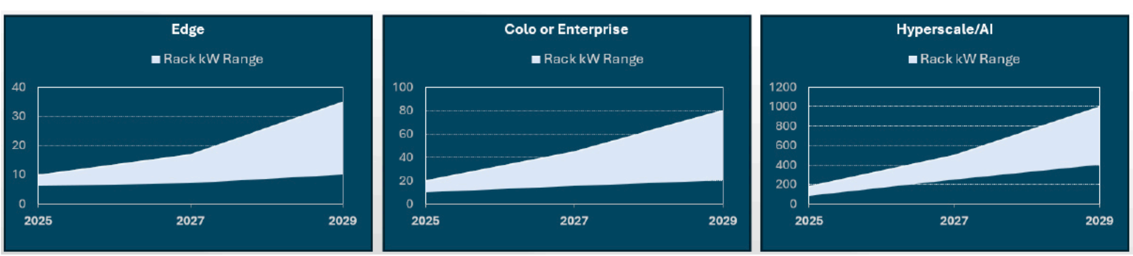

Figure 3 - Projected Greenfield Rack Power Density by Market, 2025-2029

In the edge data center market, most greenfield builds will remain low-density; however, AI may drive certain sites to higher densities. While an ideal edge data center would have vast HVAC and space capacity, the need for low latency can drive these sites into buildings/spaces not designed for moving large amounts of heat. Given that, some greenfield builds will keep a lower power density at or below 10kW per rack. At the high end, the AI-Impact will be felt over the next five years. It is certain that some AI applications (like Autonomous machines) will require ultra-low latency inference.

Given the large power draw of inference chips, edge data centers will see occasional deployment of much higher-power racks, up to 30-50KW, and perhaps beyond as liquid-cooling becomes widespread (liquid cooling is covered more in-depth later in this paper). If consumer adoption of AI grows faster than expected, these densities may grow quicker as well.

For Colocation and Enterprise, power growth will be like the edge market. However, given their typically purpose-built design (with centralized HVAC), their baseline rack density for greenfield is higher. At the low end, given that this market is based on customer/enterprise demand, rack densities will remain at or below 20kW. Like the edge market, inference deployments may take rack densities to upwards of 100kW, if not more. The market's power densities may grow faster or slower depending on the enterprise's adoption of AI.

On the hyperscale side, the growth in power density is on track to be massive. While current air-cooled rack densities max out around 50kW, liquid cooling and new power architectures will enable hyperscale companies to climb well beyond that limit. This subject is covered in depth through this paper, but in general it's expected that greenfield hyperscale rack power densities will climb into the hundreds of kW over the next five years, if not beyond 1MW.

Cost Pressures Across Data Market Segments

As rack power densities rise, the three market segments covered above will face various challenges as they attempt to minimize extraneous power draw not dedicated to generating revenue. A simple method of quantifying this unwanted power draw is the well-known Power Usage Effectiveness (PUE) equation, specifically:

$$PUE = \frac{\text{Total Facility Power}}{\text{IT Load}}$$

A "perfect" data center would have a PUE of 1.0, but this theoretical value can't truly be reached, although hyperscale companies get close. As of 2025, PUEs from 1.5-2 are common for edge data centers, while Colocation/Enterprise sites tend to be under 1.5 and hyperscale's under 1.2. This reflects that hyperscale companies tend to prioritize total system efficiency, while edge sites typically have other cost pressures that take precedence over efficiency (although it is still important to both). In Figure 4 below, the general differing cost pressures of the three data center markets are summarized.

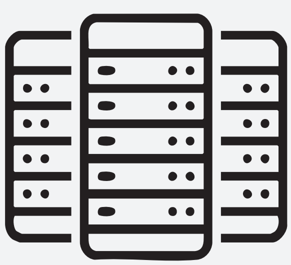

Example Case at PUE of 1.2

Simplified Example

Wholesale PNW Power Cost: $0.0349/kWh

100MW Data Center – Hourly Power Cost: $3,490

Hourly Power Cost – Non IT-Load: $582

Yearly Total Power Op-Ex: $30.5M

Yearly Non-IT Power Op-Ex: $5.1M

Data Center Type	Cost of Grid Power	Price Of Copper	Cost of Power Distribution Components	Total System Efficiency	Cost of Installation	Heat Generation
Edge PUE: 1.5-2+	HIGH	LOW	MEDIUM	MEDIUM	HIGH	LOW
Colo/Enterprise PUE: 1.2-1.5	HIGH	LOW	MEDIUM	MEDIUM	MEDIUM	HIGH
Hyperscale/AI PUE: <1.2	LOW	MEDIUM	HIGH	HIGH	LOW	HIGH

Figure 4 - Cost Pressures Driving Architecture Changes in Data Centers

Edge Data Center Cost Pressures

For Edge data centers, certain cost pressures are unavoidable. As the United States faces growing demand for grid power, the costs of this power will likely rise. This will be an unavoidable fact for edge data centers over the next couple of years. Additionally, due to their smaller size, installation costs can't scale as well for larger data centers, so products and processes that minimize labor are important.

It's also important to note that, although a cost-pressure like heat generation is rated low for the edge market, it remains significant for those who design edge power systems. An edge data center might not have hot/cold containment as well-designed as a large data center, but it would still benefit from reducing HVAC costs using something like in-rack Coolant Distribution Units (CDUs) for liquid cooling.

Colocation and Enterprise Data Center Cost Pressures

Colocation and Enterprise data centers face similar challenges due to commonalities in their design. Like edge sites, grid power costs remain a large impact, although depending on location these sites may receive favorable power costs. For colocation facilities, heat generation is large cost-pressure. While a colocation data center typically charges customers for power/space, the cost of HVAC is already factored into this, which means it's a competitive advantage to have a lower-cost HVAC system.

Total system efficiency is somewhat important to colocation and enterprise facilities. However, given the applications, widespread power options are expected to be available regardless of system efficiency. These sites typically involve a one-time purchase of power equipment and installation, resulting in lower cost pressures in those categories compared to Edge.

Enterprise/AI Data Center Cost Pressures

With power demand surging in hyperscale and AI, nearly every aspect of power-system design is under heavy cost pressure, except perhaps installation, which accounts for a smaller share of hyperscale budgets than in other markets. As previously mentioned, these sites typically focus on PUE and are pursuing efficiency improvements more aggressively than colocation or edge.

Unlike the other markets, the cost of grid power is typically not a major concern – if they can obtain power, they can usually negotiate favorable rates; however, power availability is the biggest challenge. As the US grid struggles to keep up with power growth, many hyperscale companies need to generate their own power. It's not uncommon for 100MW+ data centers to have their own power-generating turbines, making the cost of distribution components a significantly larger share of the initial budget.

Similarly, given these rising power demands, hyperscale and AI data centers are facing a significant cost in managing the heat generated by the racks. As mentioned earlier, liquid cooling is essential at higher rack power densities. For AI data centers, liquid cooling is becoming increasingly widespread and maturing rapidly as a product line, as it is an effective way to manage ongoing heat costs.

Similarly, given these rising power demands, hyperscale and AI data centers are facing a huge cost managing heat generated by the racks. As mentioned before, liquid cooling is a must at higher rack power densities. For AI data centers, liquid cooling is becoming widespread and maturing quickly as a product line, as this is an effective way to manage ongoing heat costs.

3. Architecture Overview and Design Implications

Standard Power Architectures in Data Centers

Before diving into the three major power architectures (AC/±54VDC/HVDC) in data centers, it is worth noting that the following architectures are simple examples – in reality, any deployed power system will have its own distinct setup, and many will mix AC and DC as needed by their site equipment.

That said, this architecture overview will hopefully provide the reader with an understanding of the general layouts and advantages and disadvantages of each technology.

Alternating Current (AC) Power Architecture Overview

In the United States, AC remains the predominant power architecture used in data centers nationwide (although this may be changing soon, as discussed below). In Figure 5, a standard AC architecture is illustrated, extending from the grid to the device. The input into the data center is typically high-voltage three-phase AC power, typically something like 13.8kV.

The high-voltage AC is stepped down to a lower voltage for use inside the data center using a transformer located either inside or just outside the data center campus. It will typically then go through a transfer switch (which is also connected to the backup generator), before flowing into the uninterruptible power supply (UPS).

The UPS generally will rectify the AC voltage to DC, for the purposes of charging the backup battery system (which bridges the gap from power outage to generator start-up), before inverting back to AC for distribution throughout the data center (potentially flowing through another stepdown transformer, depending on the distribution voltage). Inside the equipment itself will be a power supply unit (PSU) that performs the final rectification to 12V DC (as computer chips require DC power for operation).

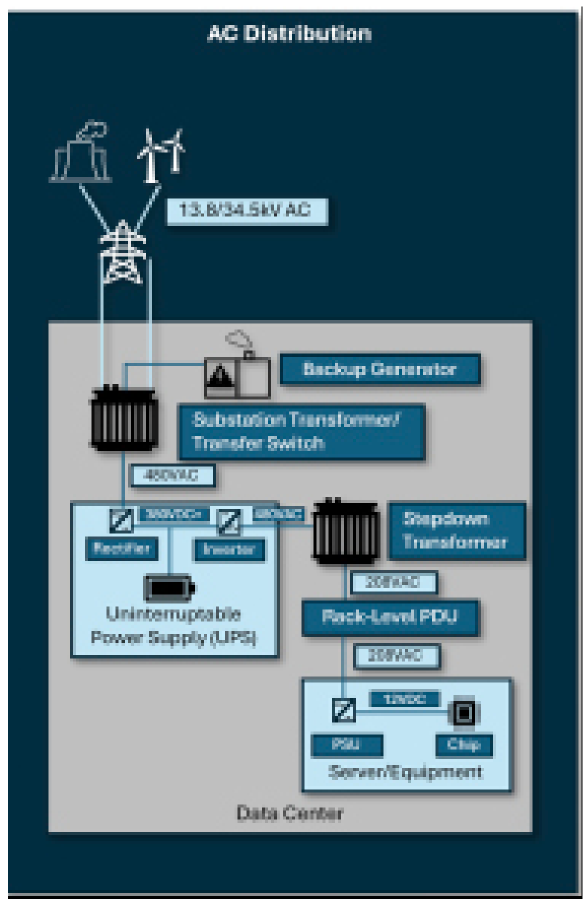

Figure 5 - AC Power Distribution Layout

±54VDC Direct Current Power Architecture Overview

±54VDC has been used for many years in the United States, but it's typically only at the smaller-scale telco level, not used very often at larger data centers. That said, over the last few years, AI/ML equipment has driven more 54V use-cases in data centers.A typical ±54V system looks like a battery bay. In the bay, like the AC system, the AC voltage will be rectified to DC for the purpose of charging the battery backup system.

Unlike a typical AC system, once rectified to ±54VDC, the power will stay at that voltage throughout the facility. This allows the battery backup system to work natively with the same voltage required by the equipment. Once the ±54VDC reaches the equipment, an internal DC-DC converter will drop the voltage to 12V or below, which the equipment's chips require.

As covered in-depth in the next section on architectural efficiencies, ±54VDC does provide some efficiency advantages compared to AC power; however, this is typically outweighed by the higher current (meaning more copper is needed), as well as a smaller installer base than AC. It should be noted that nearly the entire ±54VDC architecture can be baked into a single in-rack device (battery, rectification, and fusing).

High-Voltage Direct Current (HVDC) Power Architecture Overview

HVDC is actively used in many different industries; however, for US-based data centers, the rollout is still in its infancy. In many ways, HVDC can be the ideal solution for powering data centers; however, safety issues and low adoption in the past have meant the technology has never reached widespread use.

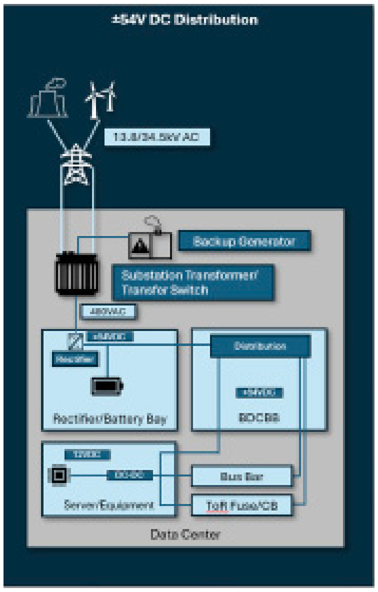

Figure 6 - 54VDC Power Distribution Architecture

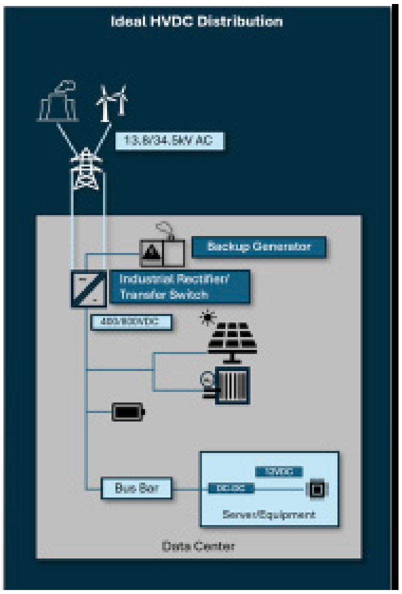

Figure 7 - HVDC Power Distribution Architecture

HVDC requires a completely different architecture than ±54VDC or AC power. When high-voltage AC enters the facility, it will typically flow into the transfer switch (usually also connected to a backup generator), followed by a massive industrial rectifier system. The rectifier outputs 400 VDC or 800 VDC, depending on the HVDC scheme in use. A key benefit is compatibility with other HVDC sources, such as solar or fuel cells, which can feed the HVDC bus directly. Unlike AC, you only need to match voltage, not phase.

Inside the data center, new HVDC architectures are being designed around simple busway-based distribution systems. A large rail system will typically deliver HVDC to each rack, and inside that rack, a rear-mounted busway makes it easy to add or remove equipment, without technicians needing to get anywhere close to the dangerous high-voltages.

Comparing Power Architectures and Their Efficiencies

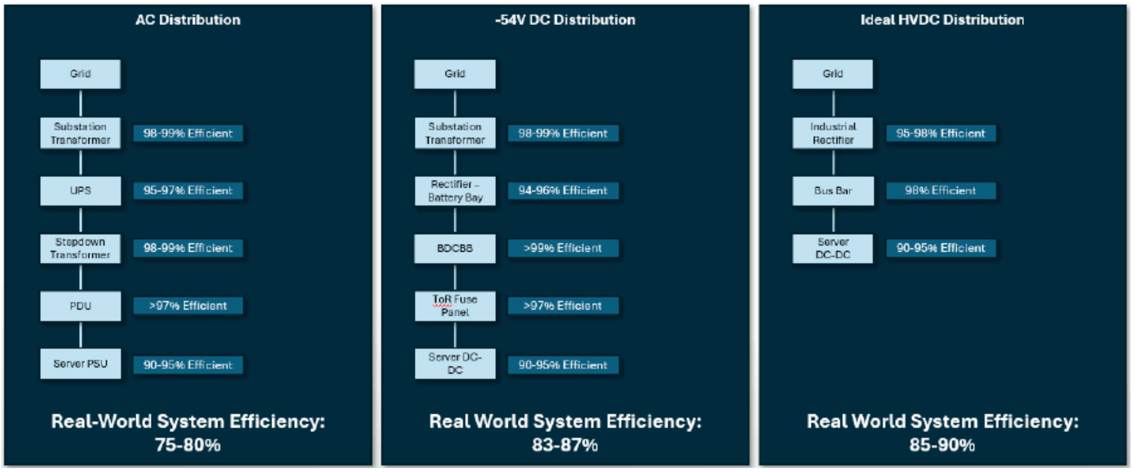

Figure 8 - Power Architecture Efficiencies - Grid-to-Chip

Comparing the efficiencies of the three architectures is useful for general trends, but in realistic settings, the actual efficiencies may vary wildly based on several factors.

In any given power architecture, multiple sources of inefficiencies can be found. Whether it's a transformer, a rectifier, or the cabling itself, some percentage of power flowing through the equipment will be lost as heat. Generally, as shown in Figure 8, the more conversion stages a system has (spots where voltage may change from AC to DC or be stepped down), the worse the system's efficiency. Additionally, a typical system using a lower voltage (like ±54VDC) will have higher currents, which leads to higher heat loss in cabling and lower total efficiency.

For the purposes of comparison, AC has the lowest total system efficiency. Although data center AC power systems have seen significant efficiency improvements over the past two decades, they inherently face limitations due to conversion losses, since chips ultimately operate on DC voltage.

±54VDC has slightly better real-world efficiency, depending on its setup. In this architecture, you typically have only one transformer step, and rectifiers can achieve very high efficiencies, depending on the load and conditions. Given that, it might be expected that ±54VDC would have displaced AC in new builds. However, as mentioned previously, a lower voltage means a higher current, which effectively drives the need for larger cables. For a small telco site (where ±54V is prevalent), this extra cost of copper isn't a showstopper; however, for larger data centers, the cost of cables may not be worth the extra percentage point or two of efficiency.

High-Voltage DC (400/800V DC) has the potential to achieve the highest overall system efficiency among the three architectures, due to its higher voltage and a typically simpler architecture with fewer conversion points. Given that, it begs the question of why the data center market is primarily standardized on AC vs DC power (with some notable exceptions).

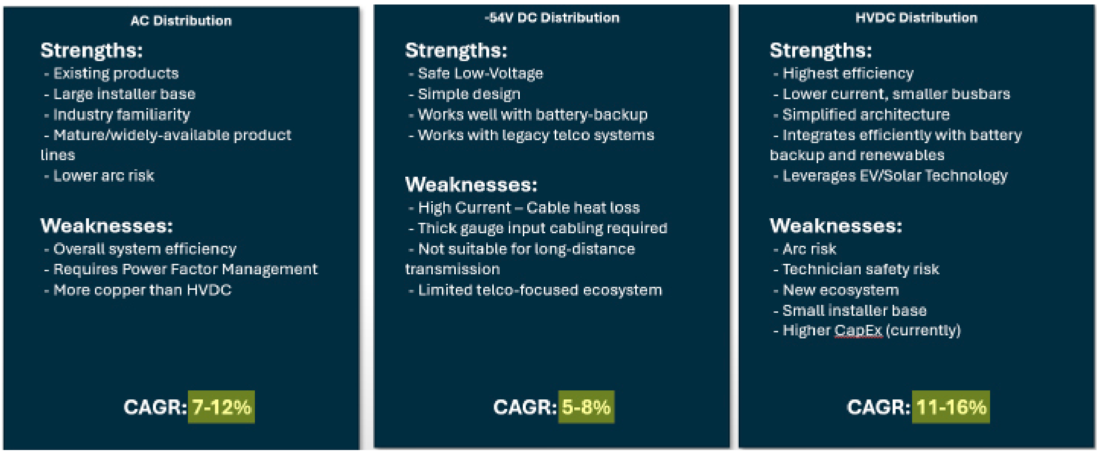

Figure 9 - Power Architecture Pros and Cons

In Figure 9, a simplified list of pros and cons of each technology can be found, outside of the major efficiency differences covered previously. This gives some reasoning behind the dominance of AC up to now – there is a very large installed base for AC power, with a cohesive and mature product base. A properly designed AC system will come within a few percentage points of a DC system in terms of efficiency, so there hasn't been a major motivation to change up to now.

HVDC has the highest potential efficiency among the three technologies; however, several constraints have hindered its adoption to date. The first key weakness of HVDC is the challenge of suppressing electrical arcs. In a high-voltage AC system, any potential electrical arcing can be extinguished more easily than in an HVDC system, due to the AC system's voltage crossing zero 120 times a second. This means that any potential HVDC system must include robust protection, including smart power solutions like solid-state breakers and digital fault detection.

These safety concerns pose a significant barrier to the implementation of HVDC, compounded by a limited supplier base. However, over the last 15 years, the United States has developed a substantial amount of Solar and EV-charging infrastructure. Both technologies utilize HVDC, and as a result, the number of manufacturers in the industry has grown dramatically, along with the number of installers familiar with HVDC. Several data center operators and manufacturers have also announced their intention to transition to HVDC for specific applications.

Projected Greenfield Architecture Trends Across Data Center Segments

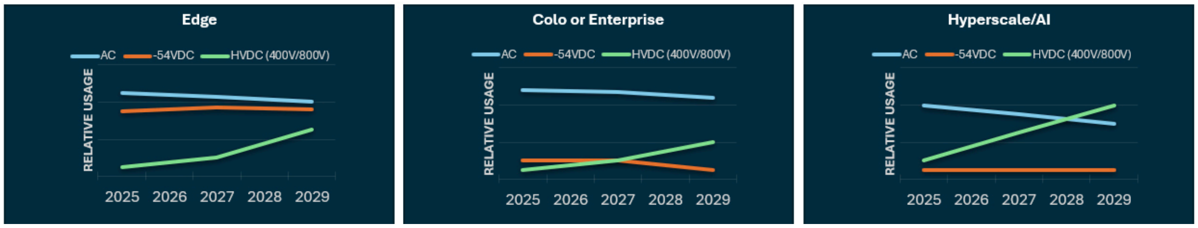

Figure 10 - Projected Trends - Relative Use of Each Power Architecture in Greenfield Builds, By Segment

In Figure 10, the estimated annual growth rates are listed for all three architectures. The numbers are generalized based on various sources [7]/[8]/[9]/[10]. In general, all three power architectures are expected to continue experiencing robust growth, driven primarily by the rapid expansion of AI. AC is expected to see the fastest growth over the next couple of years. However, as hyperscaler companies transition to HVDC, it's expected that AC growth will slow, and by the end of the decade, HVDC could reach widespread adoption.

In Figure 10, the projected relative usage of these technologies is shown. Note that this is a comparative analysis between the three, so AC will continue to see growth despite its downward-trending relative usage. As mentioned in our key assumptions, these trends assume that AI/ML will continue to drive most of the data center's development over the next five years.

The edge data center market will continue to be primarily driven by local, low latency use cases moving forward. It's expected that the edge will host more AI Inference racks, which will have several knock-on effects. In the short term, ±54VDC is expected to experience slight relative growth at the edge due to +54V Inference GPU racks; however, this is anticipated to be displaced as NVIDIA and other manufacturers standardize around HVDC. By the end of the decade, HVDC should be trending upward for edge, and the space will more than likely see a lot of mixed architectures (serving both HVDC and lower voltage DC, as well as AC)

For the Colocation/Enterprise, AC could remain very strong through the end of the decade. These facilities typically drive their decision-making on what their customers need. Depending on how popular rentable AI equipment becomes, HVDC may see some growth at the end of the decade, especially as these sites build a cohesive power offering to cover all applications. ±54VDC will likely remain low and possibly decrease further in this space as AC and HVDC dominate the market.

Hyperscale is the segment where the HVDC architecture appears poised to capture a significant portion of the market share for greenfield AI builds, particularly upon the release of NVIDIA's Kyber Rack-Scale systems in 2027. It remains to be seen whether non-AI equipment also embraces HVDC, but the scale of the AI buildout should rapidly decrease the price of HVDC rectification, power delivery, and chips, which could lead to adoption across the industry.

4. Implications for AC and ±54VDC

In short, it's predicted that the market for AC and lower-voltage DC will remain strong and continue to grow, at least for the next couple of years. In the AC market, the major manufacturers will continue to pursue small efficiency improvements; however, many of these same manufacturers are shifting their focus to HVDC systems.

The low-voltage DC market will also remain strong, particularly in sites that are difficult to upgrade, such as micro-edge sites. While HVDC may displace this in some sites by the end of the decade, hybrid systems may be required instead. If the needs for AI necessitate deploying GPUs in small sites, there may be limited use cases for DC-DC Boost Conversion (from ±54 to HVDC) or other DC conversion panels.

Notes on Other Technologies

Related HVDC Technologies

The data center space doesn't operate in a vacuum, and it makes sense to reflect on other industry spaces that may be using similar technologies. Companies like Google have mentioned that they are moving to HVDC due to the supply chains recently established around the electric vehicle.[11] Many newer electric vehicle models use 400VDC or 800VDC power systems, the same as those future data centers are planning to use.

Like EVs, many existing solar systems are using HVDC (although typically in the 1000-1500V range). While not the same as the HVDC voltages for data centers, it's relatively simple to convert this DC voltage from solar into a usable voltage for data centers.

These industries have already faced many of the same challenges that data centers will soon face, and the data center industry is wise to learn from their experiences. There will likely be significant overlap between the three industries moving forward, specifically in areas such as overcurrent protection, liquid-cooled distribution, and connectorization.

Liquid Cooling Adoption

Although not directly related to HVDC, liquid cooling is important to cover in the context of future data center architectures. As rack power densities climb, air as a cooling medium cannot move enough heat to effectively cool newer racks. While the adoption of liquid cooling has been discussed for years, new high-power racks will require it to operate effectively. Over the next few years, data centers will see both technologies develop in parallel, and it is more likely than not that HVDC will need to incorporate liquid cooling into its products.

As previously mentioned, other industries like EV-Charging have already explored technologies like liquid-cooled cabling for HVDC. Data centers will likely adopt similar products for busbars and other distribution, especially as they try to drive PUE down towards 1 for new AI/ML builds.

4. Major HVDC Stakeholders

Operators

Hyperscale and NVIDIA HVDC Adoption

This paper, up to this section, has mentioned projected trends and ideal/hypothetical scenarios. There are real-world standards/deployments that are driving these projections.

It is no secret that NVIDIA and Google are two of the most important stakeholders in the AI/ML market moving forward. Up to now, NVIDIA has been the primary supplier of chips for nearly the entire AI industry. Although more chip-makers are entering the market, it's still appropriate to assume that the industry will follow NVIDIA's lead when it comes to architecture changes. Google tends to forge its own path when it comes to certain data center design aspects, but it does work alongside other hyperscale companies in the Open Compute Project, highlighting the need for rapid adoption of these technologies.

NVIDIA has recently announced that, starting with their Kyber rack-scale systems in 2027, they will adopt 800V HVDC as their standard power architecture, primarily due to the reduction in space/copper-usage needed for power delivery compared to 54VDC.[12] This design relocates rectification of HVAC to HVDC at the data center edge, distributing HVDC throughout the entire facility.

Given the limited previous adoption of HVDC, NVIDIA is aware that the market needs to mature rapidly. As a result, they are partnering with companies such as STMicro and Texas Instruments on the HVDC chip portion, and with Delta, Eaton, Schneider, Vertiv, and others on the power system portion. The HVDC market is expected to experience rapid development in the data center space over the next few years, with hyperscale companies also working on the formal adoption of HVDC.

Mt Diablo – Microsoft, Meta, and Google's HVDC Power Sidecar

In the fall of 2024, Microsoft and Meta formally announced their Mt Diablo project, which, like NVIDIA's announcement, is intended to separate the compute and power rectification portions of the network into separate areas.[13] In spring 2025, Google also joined this project, and together they released the first Mt Diablo HVDC spec in August 2025.[14]

The specification is standardized on a +/-400VDC three-conductor system, which can also be operated in a two-conductor version, depending on the equipment's needs. For this system, a rack-level rectification cabinet is laid out and will ideally support compute loads as rack power densities climb towards the 1MW level.

Removing the rectification step from the equipment racks will simplify the design of data centers moving forward. Liquid cooling systems will be easier to design around the compute equipment, and power/compute can be upgraded incrementally, separate from each other, without being tied together due to shared rack space.

HVDC Manufacturers

Given the upcoming adoption of HVDC by hyperscale companies, including NVIDIA, manufacturers have expanded their HVDC development. Below are a few notes on selected major manufacturers in the HVDC space, and what they have publicly announced:

Silicon Providers

NVIDIA has announced partnerships with Analog Devices, Infineon, Innoscience, MPS, Navitas, ON Semiconductor, Renesas, ROHM, STMicroelectronics, and Texas Instruments for HVDC chip silicon products. Many of these companies specialize in Gallium Nitride (GaN) or Silicon Carbide (SiC) chips, which are tolerant to higher voltages and are ideal for next-generation HVDC AI systems.

Power Systems and Components

For HVDC power systems and their components, NVIDIA has announced partnerships with Delta, Flex Power, Lead Wealth, LiteOn, Megmeet, Eaton, Schneider Electric, and Vertiv. These companies will provide rectifiers, PSUs, DC-DC converters, and many other larger-scale power system products for HVDC. Some of these companies have publicly announced new, large-scale systems for HVDC, such as Delta's Infrasuite, or Vertiv's Compute Function Unit.

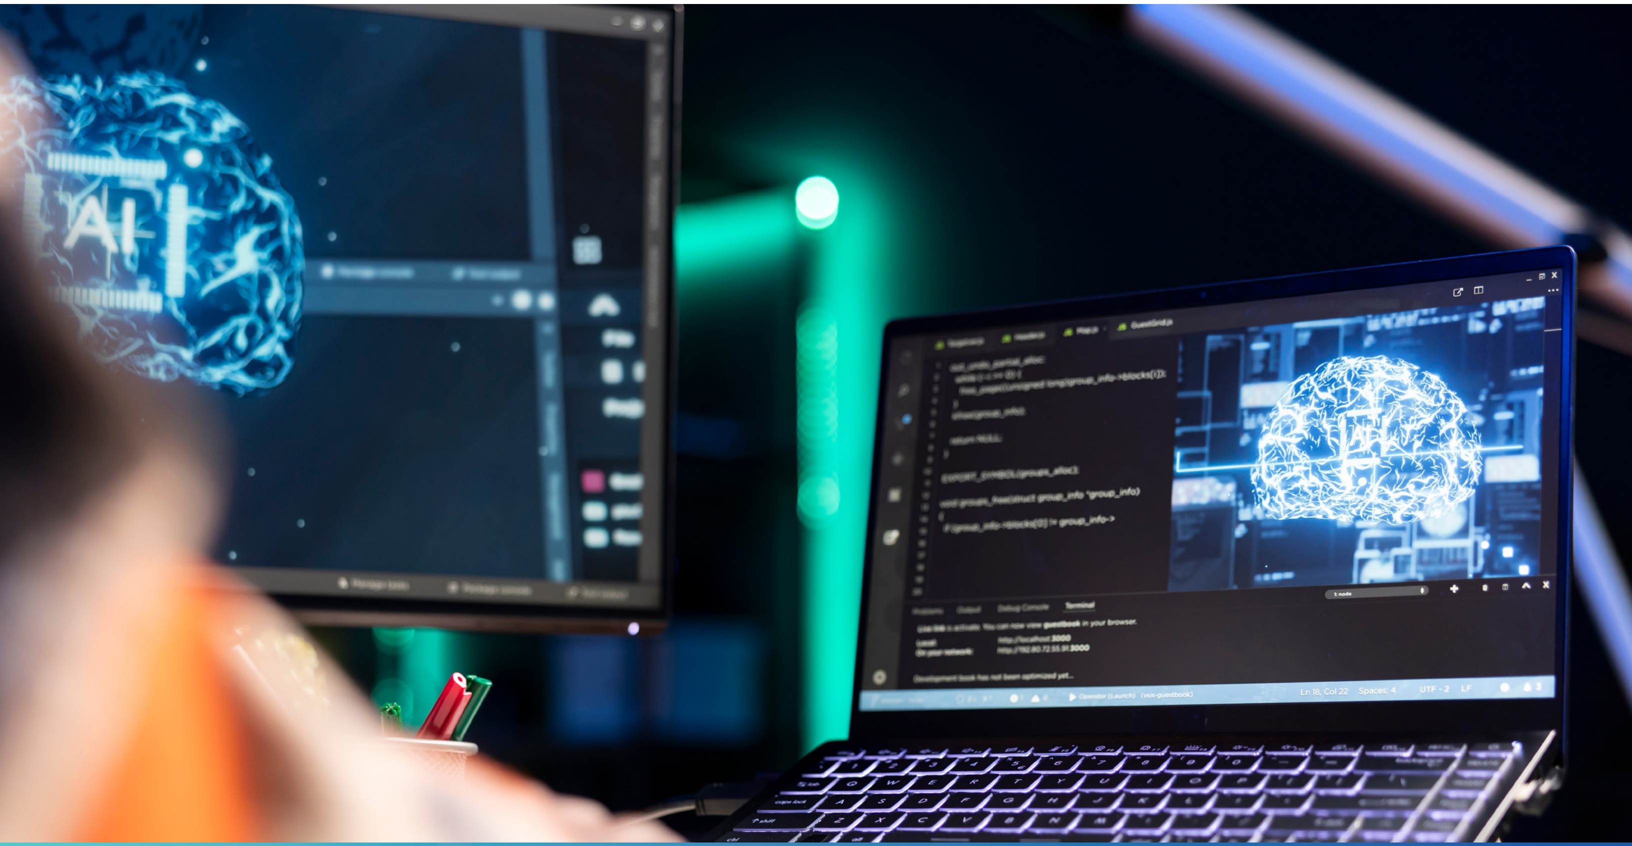

5. Summary

Prepared for informational use reflecting public information as of September 2025, this paper indicates that the data center market is undergoing a rapid transformation, both on the compute side for AI/ML development, as well as on the power side with HVDC. If AI/ML continues to drive development over the next five years, HVDC is expected to be a predominant power architecture, and potentially the primary choice for greenfield data centers. Both operators/builders including NVIDIA, Google, Microsoft, and Meta, as well as power manufacturers like Delta and Vertiv, are dedicating resources to developing standards and equipment for the technology. The next five years are likely to be a period of rapid change for the power industry.

6. References

[1] https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers

[2] https://www.tomshardware.com/pc-components/gpus/nvidias-next-gen-blackwell-ai-gpus-to-cost-up-to-dollar70000-fully-equipped-servers-range-up-to-dollar3000000-paper

[3] https://semianalysis.com/2025/07/11/meta-superintelligence-leadership-compute-talent-and-data/

[4] https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers

[5] https://world-nuclear.org/information-library/country-profiles/countries-t-z/usa-nuclear-power

[6] https://sixfoldgroup.com/2024/09/23/power-usage-effectiveness-pue

[7] https://www.gminsights.com/industry-analysis/high-voltage-direct-current-capacitor-market

[8] https://www.futuremarketinsights.com/papers/dc-power-systems-market

[9] https://straitsresearch.com/paper/dc-powered-servers-market

[10] https://www.fortunebusinessinsights.com/data-center-power-market-109792

[11] https://cloud.google.com/blog/topics/systems/enabling-1-mw-it-racks-and-liquid-cooling-at-ocp-emea-summit

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

[13] https://techcommunity.microsoft.com/blog/azureinfrastructureblog/mt-diablo---disaggregated-power-fueling-the-next-wave-of-ai-platforms/4268799

[14] https://www.opencompute.org/documents/ocp-specification-diablo-400-v0p5p2-2025-05-30-pdf

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

[16] https://www.deltapowersolutions.com/en/mcis/data-center.php

[17] https://www.ainvest.com/news/nvidia-vertiv-forge-future-ai-factories-cutting-edge-800-vdc-power-solutions-2505/

About the Author

Brad Hawkins
Data Center Systems Engineer

Brad Hawkins is a dedicated Data Center Systems Engineer, having joined Amphenol Network Solutions in 2014 as an intern and earned a degree in Engineering Physics from Whitworth University. With a strong background in Application Engineering and Product Management, Brad's expertise was instrumental in creating and expanding the Custom Configured Solutions division within the company. Beyond his professional accomplishments, Brad finds enjoyment in outdoor adventures such as camping, as well as indulging in video games, golfing, and tinkering with old technology. He embraces a forward-thinking mindset, constantly considering how the industry is evolving and strategically positioning Amphenol Network Solutions to be at the forefront of future growth.

About Amphenol Network Solutions

At Amphenol Network Solutions, we are driven by a passion for innovation and a relentless commitment to creating customized solutions that seamlessly integrate with your unique requirements. With our deep understanding of fiber optic technology, we specialize in creating tailored solutions that anticipate and adapt to the rapidly evolving demands of your network. Through our responsive support, unwavering commitment, and ongoing collaboration, we ensure that our solutions are ready to deliver superior performance and reliability.

509.926.6000
getinfo@amphenol-ns.com
www. amphenol-ns.com