09 June 2026

Watts Up! Shift to 800V Rewiring AI DC Power Stack Driving Multi-Billion Dollar Opportunities; APH, FLEX, TEL, WOLF Best Positioned in Coverage

Legacy rack architectures are running out of runway as AI-driven compute, memory, and interconnect push power envelopes toward 600 kW and beyond, where excessive amperage, heat, and copper content render low-voltage designs physically and economically impractical. The industry's answer is a step-change to 800V, which cuts amperage by ~15x and tames the so-called "copper tax," albeit at the cost of introducing new technologies. While the destination is broadly agreed upon, the journey is likely to unfold in phases, with architectural approaches differing simplistically in where AC-to-DC conversion occurs, with some more palatable in the near-term given retrofit feasibility and supply chain maturity, while others are better suited in the long-term as power densities inflect higher. The directional shift is nonetheless clear, and is already shaking up the supplier landscape from silicon to substation in order to prepare for data center IT loads that are on track to triple to ~300 GW by 2030, and dollar content scaling even faster, with the transition unlocking multi-billion dollar opportunities across the power stack. Within our Hardware & Networking coverage, we see **APH, FLEX, TEL, and WOLF** as the best-positioned to capture the rewired power stack, each playing across distinct layers of the evolving architecture.

- **Multiple paths, one destination at 800V.** The transition to 800V centers simplistically on where rectification sits, with some near-term approaches disaggregating AC-to-DC conversion into units placed adjacent to the IT rack, leveraging cheaper, more mature components and lending themselves well to retrofits, while longer-term approaches centralize rectification at the data hall perimeter and distributes 800V at row scale directly into compute racks. Either way, the transition is restructuring the power stack, with WBG (wide bandgap) semis gaining ground over silicon, rack-scale power shelves and HVDC (high voltage DC) busways displacing conventional busbars and PSUs, and facilities overall evaluating fewer conversion stages, including via solid-state transformers (SSTs).

- **Power interconnect on track to double by 2028 to multi-billion dollar TAM.** In-rack power consumption is set to scale ~8x over the next 5 years, from 132 kW today to 1 MW, driving non-linear content growth in in-rack interconnects and inter-rack cabling. Importantly, power interconnect solutions are already estimated to represent ~$150 of content per XPU, and simplistically holding content constant against current XPU forecasts implies an AI power interconnect TAM expanding from ~$2 bn in 2025 to $4 bn+ in 2028 even before considering the tailwinds from incremental content opportunities with the evolving architecture, which compares separately to third-party forecasts for the broader power cable and busway markets growing in parallel from $10 bn to $20 bn.

- **Power management solutions trending toward robust double-digit billion dollar TAM.** Rack densification, dynamic AI loads and energy storage, and the migration to 800V are layering on top of an already robust DC build, driving outsized growth in overall power management content. Importantly, by applying estimates for ~$1 mn/MW of content to data center IT load forecasts, we estimate a TAM growing at nearly a +30% CAGR, from $20 bn in 2025 to $45 bn in 2028. By comparison, third-party forecasts separately size total power distribution solutions as growing at a 20%+ CAGR, from $50 bn to $90 bn over the same window.

- **WBG power devices represent a $5 bn+ opportunity.** The 800V pivot inside the rack, non-linear PSU scaling (2x power increase drives ~5x SiC content), and grid modernization through SSTs and battery storage are collectively redefining the addressable opportunity for SiC and GaN. As a result, the WBG market addressing AI data centers and grid infrastructure is forecasted to expand at a 30%+ CAGR, from ~$1 bn in 2025 to $5 bn+ by 2030, with long-term content reaching $25K $35K per MW, split roughly evenly between SiC and GaN.

Table Of Contents

Great Scott! Rack Power Shifting to 600 kW and Beyond

Rising power and copper tax driving a shift to higher voltage architectures

The shift from low- to high-voltage data center architectures is being driven by the limitations of legacy solutions as power envelopes rise into the 600 kW+ range per rack, led by power-hungry next-gen compute, memory, and interconnect. As illustrated in Figure 1, as power requirements climb higher, legacy low-voltage architectures (*e.g., 48V/54V racks*) draw exorbitant amperage, resulting in excessive heat and requiring a significant increase in copper volume to maintain acceptable current, thermals, and voltage, which impacts both cost and physical feasibility. For example, with a 54V architecture, Nvidia highlighted one could see power shelves consume 64U of rack space at 1 MW and require up to 200 kg of copper busbar (*see blog here*). By increasing voltage, one can moderate this "copper tax" by reducing the amperage by ~15x, albeit at the expense of added safety considerations and additional or more complex components, such as those related to conversion, protection, etc.

Figure 1: Rack Comparison - *Higher Volt, Lower Amp, Less Copper*

	Blackwell		Rubin		Future …	
GPU	85 kW		500 kW			
CPU	10 kW		65 kW			
Switch Tray	10 kW		15 kW			
NICs, DPUs, etc.	5 kW		10 kW			
Other	10 kW		10 kW			
Power	120 kW		600 kW		1 MW	
Volt	54V	800V	54V	800V	54V	800V
Amp	2,222 A	150 A	11,111 A	750 A	18,519 A	1,250 A

Source: J.P. Morgan estimates.

Multiple paths to be taken as data centers transition to 800V

While the overall shift to 800V for data centers resonates against the backdrop of rising power demands, there remain different paths going forward, with various architectures being pursued both near- and long-term. *For example:*

- Nvidia provided a specification for an 800V two-wire system that distributes 800V at row scale. In this design, centralized rectifiers (*AC-to-DC conversion*) at the data hall perimeter feed busways that supply multiple compute racks directly, with conversion to the GPU's low-voltage rail occurring inside the rack.
- By comparison, an OCP specification is a ±400V three-wire bipolar configuration (*nominally 800V line-to-line*) in which rectification occurs in disaggregated "sidecar" power racks placed adjacent to the IT rack rather than at the perimeter.

Importantly, the sidecar specification is viewed by some as more of a transitional architecture, given its lower voltage, leverage of cheaper and mature components, and ability to be retrofitted into existing facilities. By comparison, the Nvidia specification may become more prevalent in new, greenfield data centers, particularly toward the end of the decade, as power requirements inflect higher.

Figure 2: Power Infrastructure Evolution

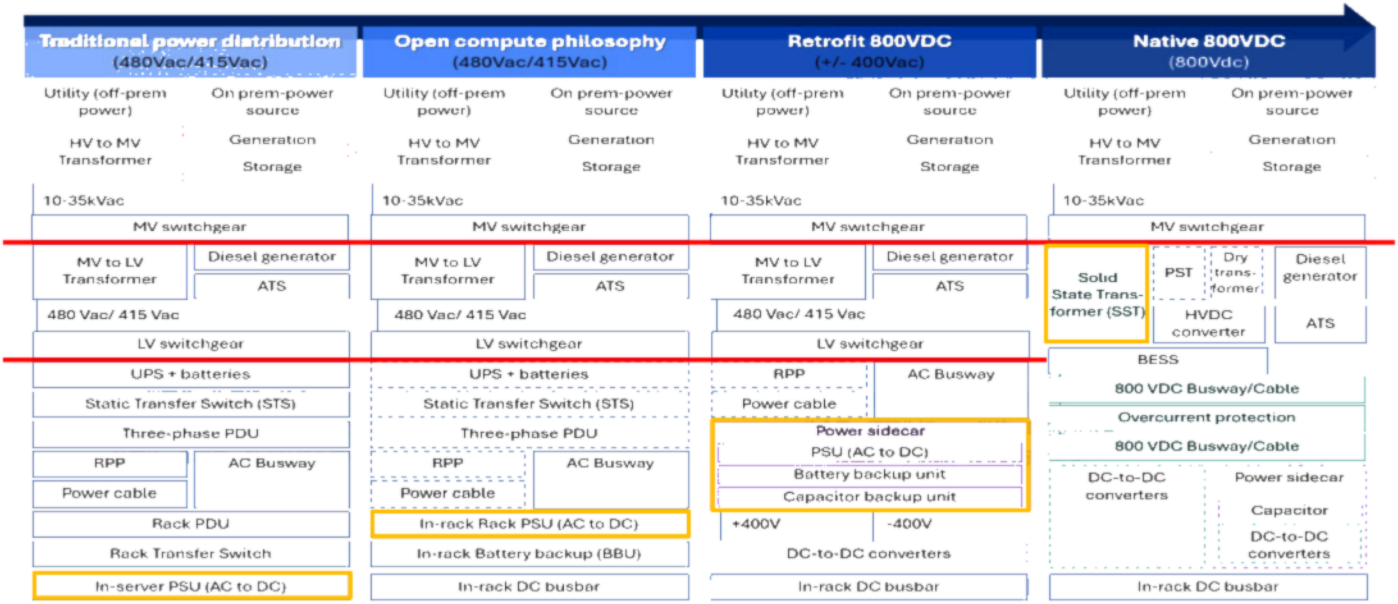

Note: Front-end is largely identical across architectures; Primary change is the shift of the AC-to-DC conversion boundary more upstream.
Source: Omdia.

New Voltage, New Victors

Shift to higher voltage driving a rewire of the stack and with it, beneficiaries

Irrespective of which path eventually gets taken, the shift toward higher voltage architectures is restructuring the AI data center power stack from the semiconductor layer through the rack and ultimately the facility itself, each with its own evolving supplier ecosystem and architectural implications. *For example*:

- At the component level, wide-bandgap semiconductors are increasingly being considered as replacements for traditional silicon, with SiC devices for example gaining traction in high-voltage rectification and power conversion stages while GaN is finding growing adoption in high-frequency downstream DC-DC conversion.
- At the rack level, conventional 48V/54V busbar architectures and multi-kW server power supply units (PSUs) are likely to give way to much higher-capacity rack-scale power shelves and high-voltage DC busways capable of supporting substantially higher rack densities while improving space efficiency.
- Distributed energy-buffering architectures are also being explored, incorporating combinations of rack-level battery backup units (BBUs), capacitive buffering, and centralized battery systems to manage loads.
- At the same time, the move to high-voltage DC requires a new generation of protection, fault management, and serviceability technologies given the significantly more demanding arc flash (*e.g., violent release of energy*) and fault interruption requirements.
- At the facility level, the industry is increasingly evaluating architectures that reduce the number of conversion stages between the utility feed and the compute rack, including medium-voltage AC converted directly into 800V DC at the data hall perimeter using rectification systems, and potentially solid-state transformer (SST)

architectures.

Figure 3: Next Generation 800V Architectures

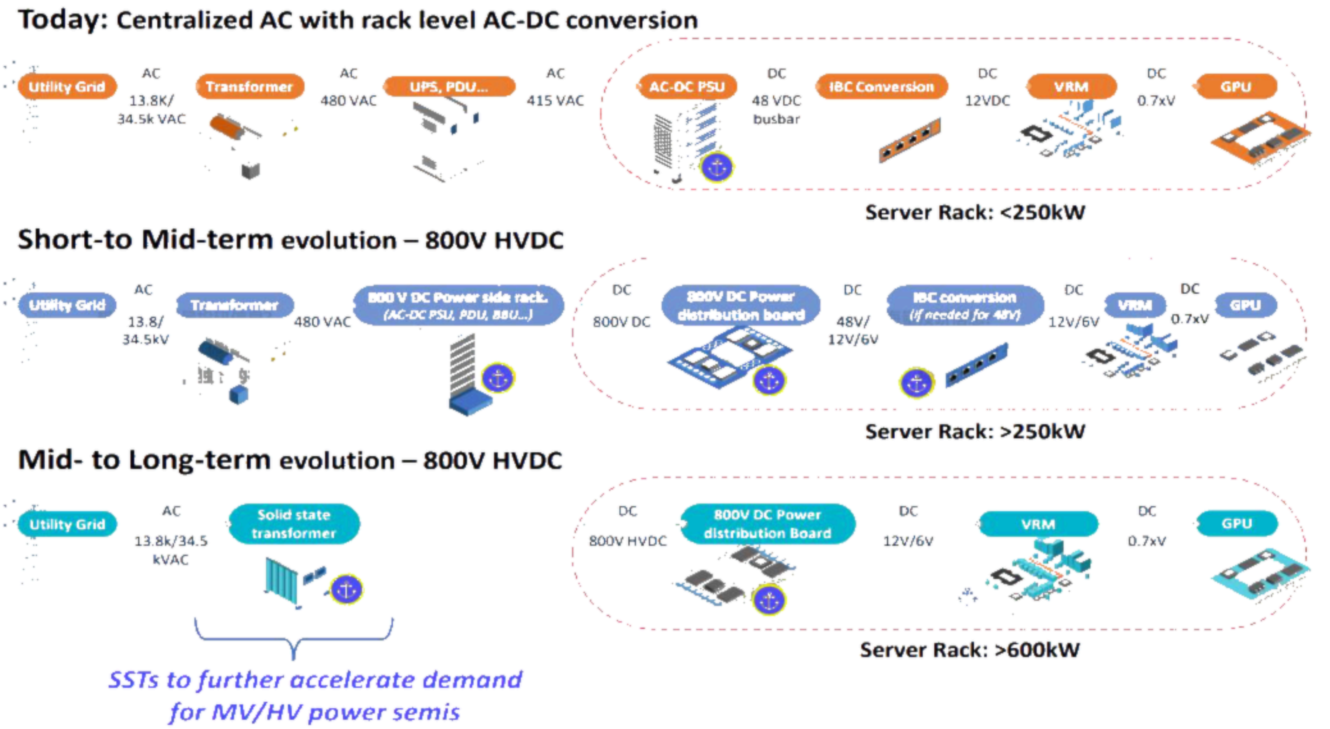

Note: "Today" boasts 80%-85% energy efficiency vs. "Short-to Mid-term" of up to 90% and "Mid-to Long-term" of over 90%.
Source: Navitas.

APH, FLEX, TEL, WOLF Best Positioned in Hardware & Networking Coverage Across the Hardware & Networking coverage, Amphenol (APH), Flex (FLEX), TE Connectivity (TEL), and Wolfspeed (WOLF) are the best positioned in our view to benefit from the data center transition to 800V architectures, with each company participating across the aforementioned "rewired stack" to varying degrees. Key highlights across each company include:

- **Amphenol (APH)** participates in the 800V transition through its rack-level interconnect portfolio, which encompasses busbar connectors and assemblies, board-level power connectors, blind-mate power couplings, and cable assemblies. Amphenol's competitive advantage lies in its breadth of qualified SKUs spanning across signal, power, and RF, combined with a strong execution track record across hyperscalers and XPU players and meaningful organic and inorganic optionality. Key peers include BizLink, Molex, and TE Connectivity.

- **Flex (FLEX)** participates in the 800V transition through its rack-level power and system integration portfolio, which encompasses power shelves, power supply units (PSUs), capacitor bank units (CBU), battery backup units (BBUs), liquid-cooled busbars, switchgear, packaged substation, and fully integrated power racks. Beyond already being named as a key Nvidia partner in power, Flex benefits from portfolio breadth across power, cooling, and compute IP on top of rack integration, in addition to a global manufacturing footprint spanning the U.S., Mexico, Malaysia, Hungary,

6

and Brazil, among others. Key peers include Bel Fuse, Delta, Eaton, LITEON, Megmeet, Schneider Electric, and Vertiv, as well as rack integrators such as Foxconn and Jabil. *We estimate the power business is 32% of total revenues for the CPI segment or the SpinCo and on track to exceed 40% by 2027.*

- **TE Connectivity (TEL)** participates in the 800V transition through its rack-level interconnect portfolio, which encompasses busbar connectors, board-level power connectors, power couplings, and cable assemblies. TE Connectivity is advantaged in that its historical exposure to automotive 800V solutions for EVs is transferable to data center use cases spanning across "grid-to-chip", including AC input, power shelf, busbar, and IT gear interfaces. Key peers include Amphenol, BizLink, and Molex. *We estimate the power business is 25% of total data center related revenues.*
- **Wolfspeed (WOLF)** participates in the 800V transition through its SiC power semiconductor portfolio, which encompasses bare die MOSFETs, packaged discrete MOSFETs, and power modules that address both SST for grid-to-facility applications and rack-level power shelves. Wolfspeed is uniquely positioned as a U.S.-based, vertically integrated supplier across SiC substrates through to devices with an extensive patent estate, inclusive of historical exposure to automotive 800V solutions, with manufacturing concentrated in New York and North Carolina. Key peers include Coherent, Infineon, Mitsubishi, Navitas, Onsemi, Power Integrations, ROHM, and STMicro. *We estimate the power business is ~10% of total revenue.*

Sampling of Key Suppliers Highlighted by Nvidia

Illustrated below is a sampling of the key suppliers Nvidia highlighted as collaborators across the data center electrical ecosystem (see blog here).

- **Silicon Providers:** AOS, Analog Devices, Efficient Power Conversion, Infineon Technologies, Innoscience, MPS, Navitas, onsemi, Power Integrations, Renesas, Richtek, ROHM, STMicroelectronics, Texas Instruments.
- **Power System Components:** Bizlink, Delta, Flex, Lead Wealth, LITEON, Megmeet.
- **Data Center Power Systems:** ABB, Eaton, GE Vernova, Heron Power, Hitachi Energy, Mitsubishi Electric, Schneider Electric, Siemens, Vertiv.

800V Transition Driving Multi-Billion Dollar Opportunities Across Power Stack

DC IT load to expand at a +25% CAGR, anchoring strong TAM expectaions

Underlying market assumptions across key power components and systems are largely anchored to data center power capacity forecasts, which appear directionally similar across industry stakeholders with total data center power expected to approach 300 GW by the end of the decade. For example, Omdia forecasts data center IT load expanding at nearly a +25% CAGR, growing from 100 GW in 2025 to 300 GW in 2030, implying more than 35 GW of new capacity coming online annually.

Figure 4: Data Center IT Load Forecast

GW of Power

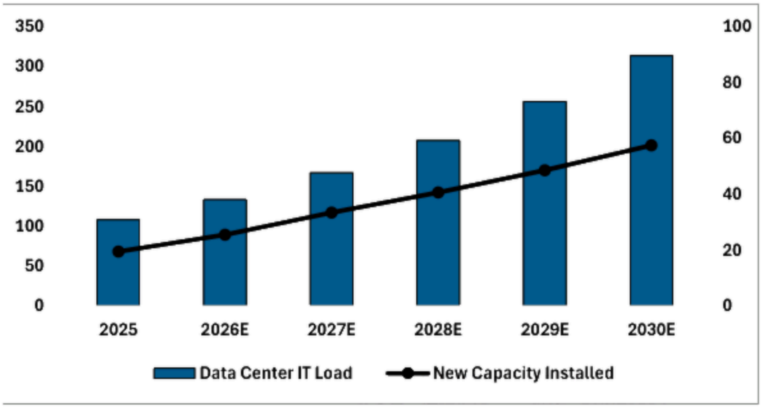

Source: Omdia.

Power interconnect market growing at double-digit CAGR to multi-billion TAM

The power interconnect opportunity within AI data centers has been framed around the non-linear scaling of in-rack power consumption, which is expected to rise ~8x, from ~132 kW today to as much as 1 MW over the next five years. This trajectory underpins a multi-vector content expansion driven by the proliferation of in-rack power interconnects, liquid bus bars, and inter-rack cable connectivity required to deliver and thermally manage higher power densities, which is separate but correlated to the build-out of upstream power infrastructure.

Figure 5: AI Power Interconnect Market Forecast

$ in Billions, Units in Millions

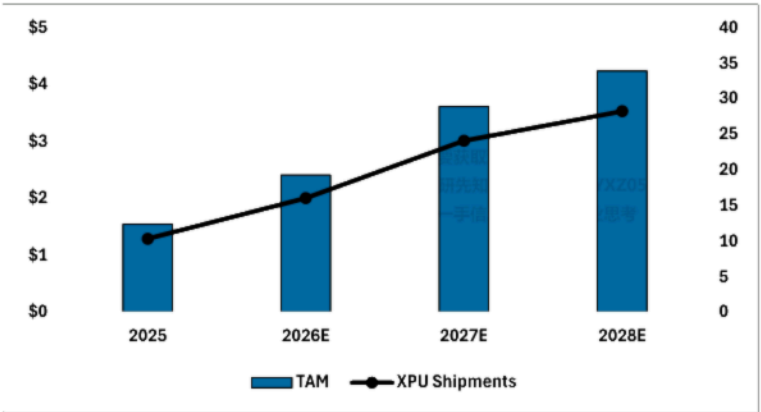

By comparison, Omdia forecasts the total power distribution cable and busway market growing at a 20%+ CAGR from $10 bn in 2025 to $20 bn in 2028, also underscoring a robust growth trajectory for power interconnect solutions overall.

Figure 6: Total Power Distribution Cable and Busway Market Forecast

$ in Billions

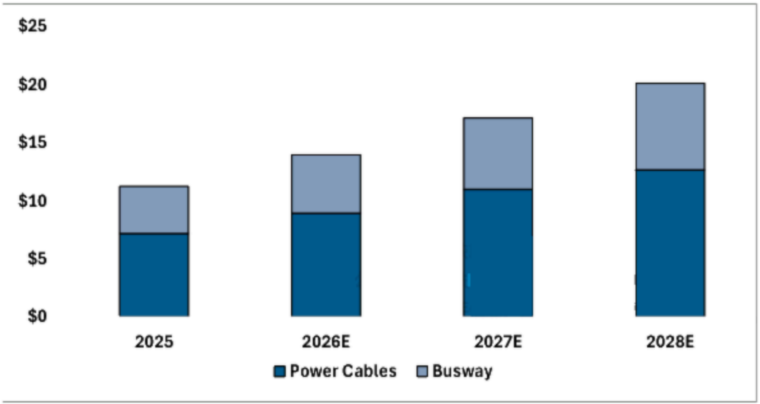

Source: Omdia.

The power interconnect opportunity is most relevant for companies in our coverage such as Amphenol (APH) and TE Connectivity (TEL).

Power distribution solutions expanding at robust rate to $50 bn+ TAM

Vertiv forecasts a data center content opportunity in the $3.0 $3.5 mn range per MW, of which roughly ~30% is accounted for by power management solutions with the rest related to thermal management, IT systems, infrastructure solutions, etc. Underpinning the content opportunity particularly as it relates to power are various trends, including: 1) rack densification, with AI racks moving from ~140 kW today toward MWs in the future, pushing traditional busbar, copper, and connector capacity to physical limits; 2) dynamic, synchronous nature of AI loads driving demand for differentiated power smoothing solutions (e.g., energy storage); 3) proliferation of energy storage across various domains; 4) shift to medium-voltage power architectures; 5) the transition to 800V architectures, with multiple deployment paths (e.g., sidecar, centralized MVDC UPS, SST solutions); and 6) on-site power generation.

Importantly, as illustrated in Figure 7, when leveraging Vertiv's expectation for power management content (*~$1 mn per MW or 30% of the total content opportunity*), we estimate the power management TAM expanding at almost a +30% CAGR from $20 bn in 2025 to $45 bn in 2028 or roughly in line with Omdia's forecasts for new annual data center capacity coming online.

Figure 7: Power Management Solutions Market

$ in Billions

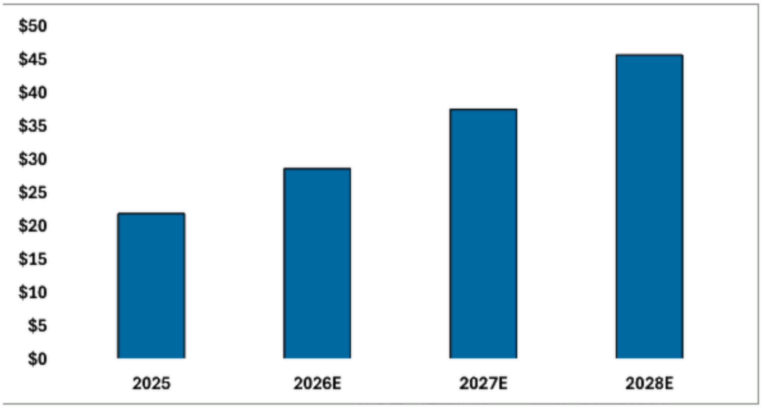

Source: Vertiv, Omdia, and J.P. Morgan estimates.

By comparison, Omdia forecasts power distribution solutions, including many of which that Delta, Vertiv and/or Flex provides, growing at a 20%+ CAGR from $50 bn in 2025 to $90 bn in 2028.

Figure 8: Total Power Distribution Solutions TAM

$ in Billions

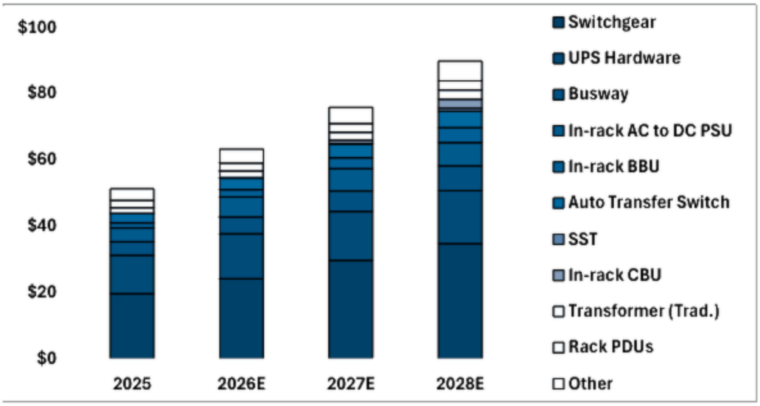

Source: Omdia.

The power distribution solutions opportunity is most relevant for companies in our coverage such as Flex (FLEX).

Wide-bandgap power devices growing at 30%+ CAGR to $5 bn+ opportunity

As illustrated in Figure 9, the wide-bandgap power device market addressing AI data centers and energy & grid infrastructure is forecasted by Navitas to expand at a 30%+ CAGR from ~$1 bn in 2025 to $5 bn+ in 2030, led by trends incl... architectural pivot to 800V power delivery within AI racks; 2) th... PSU power levels from 5-10kW today to 18.5-30kW, where a 2x ... drives a ~5x increase in SiC content; and 3) the modernization of ... through deployment of SSTs, battery energy storage systems, etc. ... the 219 298 GW of projected AI datacenter demand by 2030. Th... opportunity of $25 K $35 K per MW in the long-term, of which ... associated with SiC inside and outside the data center with the re... by GaN, primarily inside the data center.

10

The wide-bandgap power device opportunity would be most relevant for companies in our coverage such as Wolfspeed (WOLF).

Figure 9: Wide-Bandgap Power Device Market

$ in Billions

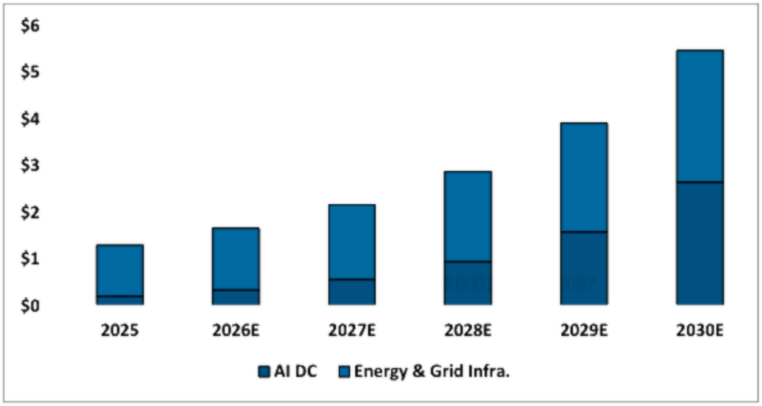

Source: Navitas.