Global Semiconductors

Watts to Tokens: AI Power Semis and the Transition to 800V Data Centers

Industry Overview

25 May 2026

Equity
Global
Semiconductors

Vivek Arya
Research Analyst
BofAS
vivek.arya@bofa.com

Duksan Jang
Research Analyst
BofAS
duksan.jang@bofa.com

Michael Mani
Research Analyst
BofAS
michael.mani@bofa.com

Didier Scemama >>
Research Analyst
MLI (UK)
didier.scemama@bofa.com

Mikio Hirakawa >>
Research Analyst
BofAS Japan
mikio.hirakawa@bofa.com

AI Power: transformative opportunity for analog semis

Power is becoming a critical constraint in AI scaling. We estimate rising compute density drives rack power capacity up 100x from 10-15kW in traditional cloud servers to 1.5MW for Nvidia's Feynman platform (CY29/30). Existing infrastructure cannot meet demand, requiring wholesale revamp of power delivery from grid to rack to GPU/XPU core. This creates a transformative opportunity for analog semi vendors to shift mix away from cyclical auto/industrial demand toward secular, durable AI markets, where diverse architectures, new components, novel materials (wide-bandgap semis), and products create an unparalleled chance to differentiate and take share. In this report, we build a bottoms-up AI analog semi industry model translating accelerator/rack demand into content pools across components (VRM, IBC, SST, etc.) and device types (Si, SiC, GaN, etc.) into supplier revenue share across low power (<200kW) and high power (600kW+) racks. We estimate ~233GWs deployed through CY30 expands today's $7.9bn AI analog market to $27bn by CY30 (28% CAGR) inside a $1.7tn AI data center systems LT TAM.

Data Centers: $25bn TAM led by 100x jump in rack power

We estimate rack content grows ~25x from $36K today to nearly $300K per 600kW rack and approaching $1mn in the MW-class era. The TAM from "rack-to-core" expands to $25bn CY30 from $7.6bn today (27% CAGR). Value shifts towards components closest to the accelerator including multi-phase voltage regulator modules (Infineon, TXN, ADI, Renesas) and intermediate bus converters (IBC; ON, Infineon,) but also optics (STMicro, ADI, TXN). As high-performance power density becomes critical in 800 VDC architecture, silicon carbide (SiC) and gallium nitride (GaN) are greenfield new markets.

Power Infra: additional $2bn TAM as microgrids inflect

Select strategic opportunities in power infrastructure grow to $38.9K per MW during the 800 VDC evolution from $12.4K per MW today. The "grid-to-data hall" TAM grows to $1.8bn by CY30 from ~$245mn today (49% CAGR). Legacy equipment makes way for emerging tech like solid-state transformers (SSTs) and solid-state circuit breakers (SSCB) as facilities transition to microgrids. SiC is a big materials winner but we see ample analog IC, MCU, and sensing content in a market where there was little before.

Revenue opp led by analog ICs; discretes gain most share

Best positioned vendors are those with: (1) broadest portfolio spanning the power tree across multiple device types; (2) products that meet high voltage and elite reliability requirements (similar to auto/industrial semi products); (3) provide system-level design expertise and optimization from grid-to-core. TXN's leading power semi franchise gives it the highest share in the market. Infineon boasts the broadest AI portfolio across Si, SiC, and GaN and may gain the most share CY25-30 from grid-to-core. ADI enjoys the third biggest revenue opportunity and also gains share enhanced by the Empower acquisition. ON has high leverage to novel SiC and GaN tech boosting share of wallet.

>> Employed by a non-US affiliate of BofAS and is not registered/qualified as a research analyst under the FINRA rules.
Refer to "Other Important Disclosures" for information on certain BofA Securities entities that take responsibility for the information herein in particular jurisdictions.
BofA Securities does and seeks to do business with issuers covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision.
Refer to important disclosures on page 49 to 51. **12977687**

Timestamp: 25 May 2026 10:00AM EDT

Contents

Power is the ultimate AI scaling constraint	3
Compute roadmap: 1MW per rack by CY30	4
GW model: AI needs 233GW through CY30	7
How power is delivered today	8
800 VDC: power delivery for the GW era	9
Analog semis: key beneficiaries in high power	14
Data center: $25bn TAM from rack to core	18
Power Supply Units (PSU)	19
Intermediate Bus Conversion (IBC)	20
Server Board (GPU/XPU power, CPU, VRM, multi-phase)	23
Other components (protection, sensors, optics, etc.)	26
Power Infra: $2bn TAM from grid to hall	29
Energy Storage System (ESS)/Uninterruptible Power Supply (UPS)	30
Solid-State Transformer (SST)	31
Solid-State Circuit Breakers (SSCB)	32
AI Power Analog Semi Model	34
Glossary	44

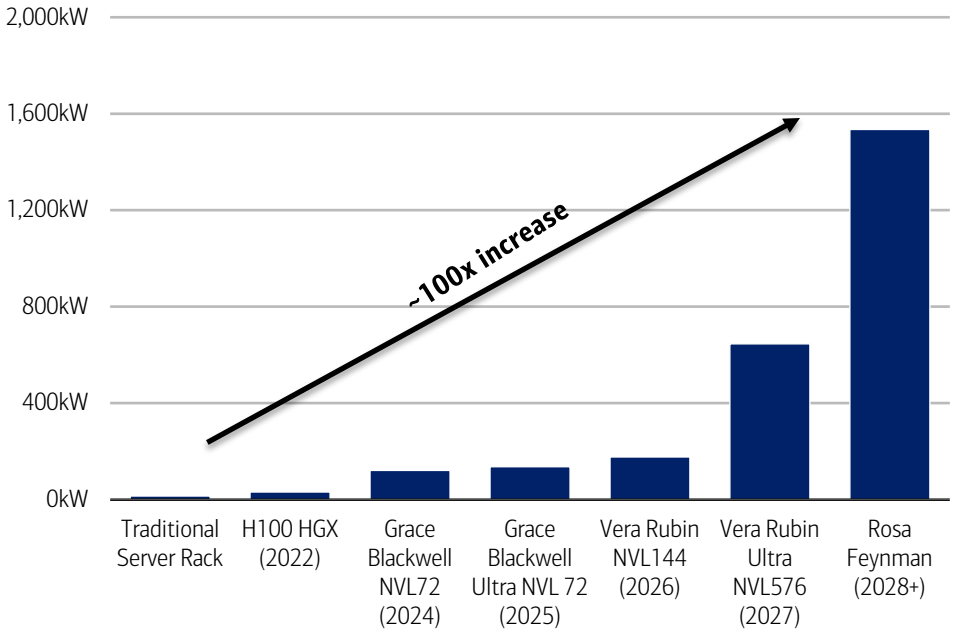
Power is the ultimate AI scaling constraint

Multiple bottlenecks are emerging in the multi-year AI infrastructure buildout – memory, optics, leading-edge logic wafers, advanced substrates, etc. – but among the greatest constraints to scaling is arguably power.

Higher compute density is translating to higher power demands

Traditional cloud data center operators mainly worried about compute space as power and cooling infrastructure occupied a significantly smaller physical footprint. When the AI investment cycle began, server CPUs transitioned to power hungry GPUs, ordinary internet traffic evolved into intense training and inference workloads, and heavy computational needs turned into heavy power needs. As demand surges and clusters become larger, the need for higher compute density per server rack (e.g. Hopper had only 8 GPUs in a node while Blackwell's scale-up domain is 72) to optimize performance means power per rack also goes up. Whereas a traditional cloud rack may consume 10-15 kilowatts (kW), Blackwell generation racks require 100-120kW, and future generations like Feynman may drive a staggering over 1 megawatt per rack (enough to power up to 1000 US homes) as up to 576 GPUs are packed into a single node (exhibit 1).

Exhibit 1: By the end of the decade, we think Nvidia's roadmap could lead to >1.5 megawatt racks for the Feynman era, nearly 100x higher than a standard server CPU rack at 10-15 kilowatts
Power capacity per rack (kilowatts) across various generations of Nvidia platforms vs. a traditional server rack



Source: BofA Global Research estimates, Nvidia, company reports

BofA GLOBAL RESEARCH

Performance-per-watt is a critical metric for operators as it determines the amount of tokens you can generate per unit of power, and thus, the amount of revenue. Data Centers operate inside a fixed power envelope which is heavily dependent on location, so maximizing access to power and converting it to tokens in an efficient manner is crucial.

Bringing gigawatt-scale power online takes time

The gating factor in power stems having to build the entire physical infrastructure flow for a new data center from grid interconnects, substations, transformers, switchgear, cooling and more. It takes time to bring this capacity online. Key equipment is also constrained, including gas turbines (5-year delivery timeline) and transformers (2–3-year lead times) Sites that have secured grid power may be able to reach 1GW capacity in 1-4 years from construction but a greenfield project could take as long as 5 years, according to the International Energy Agency, before project delays.

Compute roadmap: 1MW per rack by CY30

Hopper H100 HGX was Nvidia's first flagship anchor accelerator platform for the AI-era launched in 2022 and utilized Hopper GPUs where thermal design power (TDP) fetched only 0.7kW per package while other functions like switches, CPUs, and NICs, represented a small part of the power budget all in an architecture that was air-cooled (exhibit 2).

As we moved into Blackwell (GB200 NVL72), total GPU power budget jumped +400% from 17kW to 86kW as GPU TDP increased and GPU die count multiplied to 144. Networking and CPUs became more relevant driving up power consumption and an entirely new liquid-cooled rack, Oberon, was required to handle the heat generation. Total power increased from 25kW per rack to 100-120kw per rack.

As the launch of Rubin gets underway, power climbs again albeit to a lesser extent, but Rubin Ultra in CY27 represents another inflection to over 600kW per rack (+240% platform-over-platform), with its significantly larger scale-up domain of up to 576 GPUs, higher TDPs in the 3kW range, 144 Vera CPUs equipped for agentic AI, and powerful NVSwitch 6 inside the next-generation Kyber rack form factor. This is when we would expect to see initial deployments of 800 VDC due to the power demands.

While we can only speculate on specifications as official details are sparse, these same trends should persist as we approach the Feynman era (likely launches CY29/30) where architectures begin to migrate to hybrid microgrids (more on this later) and compute TDP likely takes another step up again. The roadmap then enters the zone of 1MW class racks where the sheer amount of power density packed into the physical space of filing cabinets necessitates completely new infrastructure.

Exhibit 2: The main driver of higher rack power advancing towards 1MW per rack by CY30 (Feynman) is GPU where TDP (thermal design power) is approaching 5kW per package. Rising switching, NIC, CPU, and DPU power budgets are also factors.
Detailed power budget estimates across various generations of Nvidia platforms

Nvidia						
Compute	**Hopper (H100)**	**GB200 (Blackwell B200)**	**GB300 (Blackwell Ultra/B300)**	**Rubin (VR200)**	**Rubin Ultra (VR300 Ultra)**	**Feynman (RF200)**
Platform	**H100 HGX Rack**	**Grace Blackwell NVL72**	**Grace Blackwell Ultra NVL 72**	**Vera Rubin NVL72**	**Vera Rubin Ultra NVL144**	**Rosa Feynman NVL576**
Rack Architecture	Air-Cooled	Oberon; Liquid-cooled	Oberon; Liquid-cooled	Oberon; Liquid-cooled	Kyber; Liquid-cooled	Hybrid microgrid
GPU	**Hopper**	**Blackwell**	**Blackwell Ultra**	**Rubin**	**Rubin Ultra**	**Feynman**
GPU Package Count	32	72	72	72	144	576
GPU Die Count per Package	1	2	2	2	4	2
Total Dies per Rack	32	144	144	144	576	1152
GPU Power per Package Consumption (kW)	0.7	1.2	1.4	1.8	3.6	1.8
GPU Power Budget (kW)	**22.4**	**86.4**	**100.8**	**129.6**	**518.4**	**1036.8**
CPU	**Xeon/EPYC**	**Grace**	**Grace**	**Vera**	**Vera**	**Rosa**
CPU Count	8	36	36	36	72	288
CPU Power Consumption (kW)	0.30	0.30	0.30	0.35	0.35	0.35
CPU Power Budget (kW)	**2.4**	**10.8**	**10.8**	**12.6**	**25.2**	**100.8**
NIC	**CX-7**	**CX-7**	**CX-8**	**CX-9**	**CX-9**	**Next Gen CX**
NIC Die Count	32	72	72	72	144	576
NIC Power Consumption (kw)	0.03	0.03	0.03	0.03	0.03	0.05
NIC Power Budget (kW)	**0.8**	**1.8**	**2.2**	**2.2**	**4.3**	**28.8**
NVLink Switch	**NVSwitch3**	**NVSwitch5**	**NVSwitch5**	**NVSwitch6**	**NVSwitch7**	**Next Gen NVSwitch**
Switch Die Count	12	18	18	18	36	144
Switch Power Consumption (kW)	0.3	0.8	0.8	1.1	2	2
Switch Power Budget (kW)	**3.0**	**14.4**	**14.4**	**19.8**	**54.45**	**288**
DPU	**BlueField-3**	**BlueField-3**	**BlueField-3**	**BlueField-4**	**BlueField-4**	**Next Gen BlueField**
DPU Die Count	3	18	18	18	36	144
DPU Power Consumption (kW)	0.1	0.1	0.1	0.2	0.2	0.2
DPU Power Budget (kW)	**0.2**	**1.4**	**1.4**	**2.7**	**5.4**	**28.8**
PSU Conversion Losses (kW)	0.7	3.5	3.9	5.7	22.8	31.4
Cooling (kW)	1.5	1.5	1.5	2	12	15
Other (kW)	0.5	1.0	1.0	2.0	3.0	5.0
Overhead Power Budget (kW)	**2.7**	**6**	**6.4**	**9.7**	**37.8**	**51.4**
Total Rack Power (kW)	**32**	**121**	**136**	**177**	**646**	**1535**

Source: BofA Global Research estimates, Nvidia

BofA GLOBAL RESEARCH

Compute and networking contribute the most to rising power

As observed, when CPUs and GPUs improve, TDP typically increases 20% generation over generation thus leading to higher server power over time. As GPUs are networked together (via NVLink or another protocol for non-NVDA platforms) to upgrade performance by functioning as a synchronized node of compute, the use of copper interconnects for scale-up means that GPUs cannot be too far from each other (limited reach). Thus, limitation of tightly integrating GPUs on the same copper domain at limited distances creates a direct relationship between maximum power density and the ability to drive maximum performance as you need to pack as many GPUs possible into a smaller physical space in order to scale (Nvidia calls this the performance-density trap). Moving from Hopper to Blackwell led to a 75% increase in TDP but a 3.4x increase in rack power density which in turn translated to a 50x increase in performance. Every subsequent increase in the scale-up networking domain size is now translating into up to 2-4x increases in total power (exhibit 3).

Exhibit 3: Large inflections in power tend to coincide with architecture shifts like when Nvidia migrated to Oberon racks for Blackwell and eventually Kyber for Rubin Ultra. GPUs generally consume most power (65-75%) followed by Switches (10-13%)
Nvidia power budget analysis spanning Hopper to Feynman

Power Budget Analysis (kW)						
Compute Platform	Hopper (H200) H100 HGX	GB200 (Blackwell B200) Grace Blackwell NVL72	GB300 (Blackwell Ultra/B300) Grace Blackwell Ultra NVL 72	Rubin (VR200) Vera Rubin NVL144	Rubin Ultra (VR300 Ultra) Vera Rubin Ultra NVL576	Feynman (RF200) Rosa Feynman
GPU	22.4	86.4	100.8	129.6	518.4	1036.8
CPU	2.4	10.8	10.8	12.6	25.2	100.8
NIC	0.8	1.8	2.2	2.2	4.3	28.8
Switch	3.0	14.4	14.4	19.8	54.5	288.0
DPU	0.2	1.4	1.4	2.7	5.4	28.8
Overhead	2.7	6.0	6.4	9.7	37.8	51.4
Total Rack Power	**32**	**121**	**136**	**177**	**646**	**1535**
Power Budget Analysis (% budget mix)						
GPU	71.1%	71.6%	74.2%	73.4%	80.3%	67.6%
CPU	7.6%	8.9%	7.9%	7.1%	3.9%	6.6%
NIC	2.5%	1.5%	1.6%	1.2%	0.7%	1.9%
Switch	9.5%	11.9%	10.6%	11.2%	8.4%	18.8%
DPU	0.7%	1.1%	1.0%	1.5%	0.8%	1.9%
Overhead	8.6%	5.0%	4.7%	5.5%	5.9%	3.3%
Total Rack Power	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**
Power Budget Analysis (platform % change)						
GPU	-	286%	17%	29%	300%	100%
CPU	-	350%	0%	17%	100%	300%
NIC	-	125%	20%	0%	100%	567%
Switch	-	380%	0%	38%	175%	429%
DPU	-	500%	0%	100%	100%	433%
Overhead	-	122%	7%	52%	290%	36%
Total Rack Power	**-**	**283%**	**13%**	**30%**	**266%**	**138%**

Source: BofA Global Research estimates, Nvidia

BofA GLOBAL RESEARCH

We think this relationship between packing more GPUs in a single domain, tighter integration, and larger networking topologies will be the common themes commanding higher power capacity in future Nvidia platforms for years to come.

Alternative platforms to Nvidia also see rising power due to similar trends

Similarly, we think these trends apply to Nvidia alternatives also in the market (Exhibit 4). As of now, competing platforms offered by AMD or Intel are less power hungry as compute lags a generation behind Nvidia and there is less integrated networking. Custom ASICs are generally lower power consuming as key functions are optimized for specific tasks making them more efficient. Yet we still believe future platforms will converge towards higher power capacities as compute and networking intensity rise as this will be necessary for them to become competitive options against Nvidia's leading performance-per-watt offering. This evolution is evident in the latest platforms like AMD Helios where power budget exceeds 100kW.

Exhibit 4: Other merchant GPUs and custom ASIC platforms are seeing rising power over time due to higher compute and networking density
Merchant GPU and custom ASIC power budget analysis

Other Merchant and Custom					
Company	AMD	AMD	Intel	AWS	Google
Platform	MI350X Server Rack	Helios MI400X	Gaudi Server Rack	Trainium 3 UltraServer	v7 Ironwood Cube
XPU	MI355X	MI400X/MI455X	Gaudi 3	Trainium 3	TPU v7 Ironwood
XPU Package Count	64	72	40	64	64
XPU Die Count per Package	1	2	1	2	2
Total Dies per Rack	64	144	40	128	128
XPU Power per Package Consumption (kW)	1.4	2	1	1	1
GPU Power Budget (kW)	89.6	144.0	40	64	64
CPU	EPYC Turin	EPYC Venice	Xeon Granite Rapids	AWS Graviton	Axion
CPU Count	16	36	10	16	16
CPU Power Consumption (kW)	0.40	0.45	0.35	0.25	0.25
CPU Power Budget (kW)	6.4	16.2	3.5	4.0	4.0
NIC/Scale Out	Pensando/ConnectX	Pensando Vulcano	ConnectX-7	AWS Nitro v6	OSFP Optical
NIC Die Count	64	72	40	144	256
NIC Power Consumption (kw)	0.05	0.03	0.03	0.03	0.01
Scale-out Top of Rack (ToR) Switch	InfiniBand/Ethernet	Tomahawk 6	Ethernet	Spine-leaf	OCS chassis and ToR
Scale-out Top of Rack (kW)	2.0	Included above	2.0	3.0	3.0
NIC/Scale-Out Power Budget (kW)	5.2	1.8	3.0	7.3	5.6
Scale Up Fabric	InfinityFabric	UALink	RoCE	NeuronSwitch-v1	ICI 3D Torus and OCS
Switch Die Count	Mesh on chip	12	On-chip	24	Only external OCS
Switch Power Consumption (kW)	0	1.5	0	0.6	0
Switch Power Budget (kW)	0	0	0	14.4	0
DPU	Pensando NIC	Pensando NIC	BF-3	Nitro DPU	Axion
DPU Die Count	0	0	0.5	2	0
DPU Power Consumption (kW)	0.1	0.1	0.1	0.2	0.2
DPU Power Budget (kW)	0.0	0.0	0.0	0.3	0.0
PSU Conversion Losses (kW)	3	5	1	4	1
Cooling (kW)	2	4	2	4	3
Other (kW)	1	2	1	2	1
Overhead Power Budget (kW)	6	11	4	10	5
Total Rack Power (kW)	107	173	51	100	79

Power Budget Analysis (kW)					
Company	AMD	AMD	Intel	AWS	Google
Platform	MI350X Server Rack	Helios MI400X	Gaudi Server Rack	Trainium 3 UltraServer	v7 Ironwood Cube
GPU	89.6	144.0	40.0	64.0	64.0
CPU	6.4	16.2	3.5	4.0	4.0
NIC/Scale-out	5.2	1.8	3.0	7.3	5.6
Switch	0.0	0.0	0.0	14.4	0.0
DPU	0.0	0.0	0.0	0.3	0.0
Overhead	6.0	11.0	4.0	10.0	5.0
Total Rack Power	107	173	51	100	79
Power Budget Analysis (% budget mix)					
GPU	83.6%	83.2%	79.1%	64.0%	81.5%
CPU	6.0%	9.4%	6.9%	4.0%	5.1%
NIC/Scale-out	4.9%	1.0%	5.9%	7.3%	7.1%
Switch	0.0%	0.0%	0.0%	14.4%	0.0%
DPU	0.0%	0.0%	0.1%	0.3%	0.0%
Overhead	5.6%	6.4%	7.9%	10.0%	6.4%
Total Rack Power	100.0%	100.0%	100.0%	100.0%	100.0%

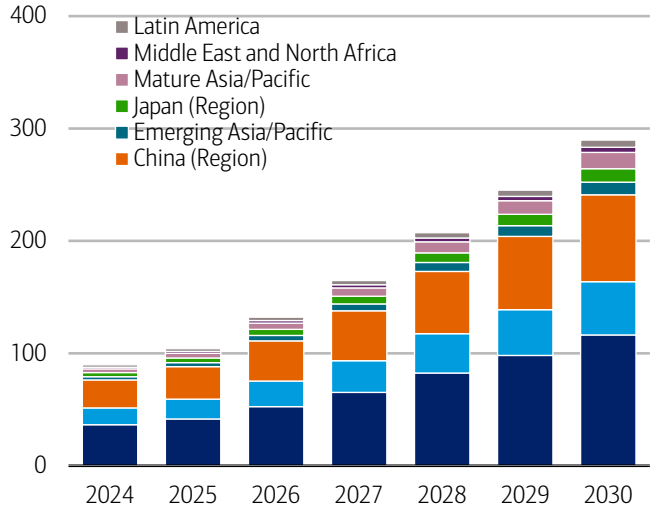
Source: BofA Global Research estimates, AMD, Intel, AWS, Google

BofA GLOBAL RESEARCH

GW model: AI needs 233GW through CY30

Per the International Energy Agency (IEA), electricity consumption from data centers was almost 500TWh in CY25, or 1.5% of global electricity demand. Based on planned global project pipelines for data centers in early CY26, if all projects are realized, installed capacity could double in the coming years from just around 100GW to almost 200GW in the coming years. However, we think AI/compute estimates for current and upcoming GPU/XPU platforms suggest demand is well above that.

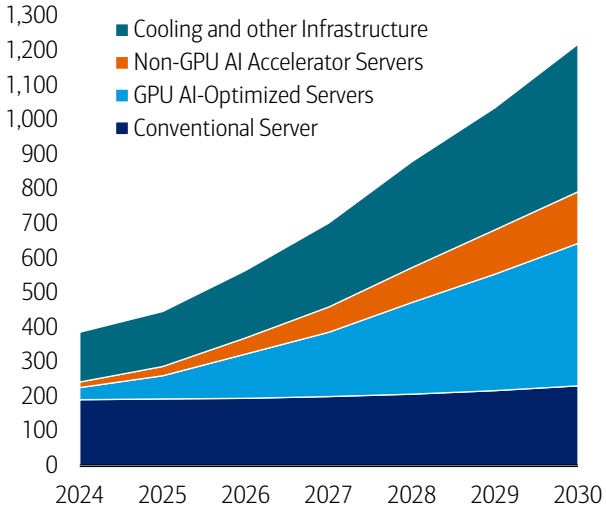
Exhibit 5: Cumulative data center power demand globally is expected to increase from ~100GW today to almost 300GW by CY30
Data center power demand by region (GW)



Source: BofA Global Research, Gartner

BofA GLOBAL RESEARCH

Exhibit 6: Layering on top of conventional server power demand, AI is catalyzing electricity consumption to new heights
Data center electricity consumption (TWh)

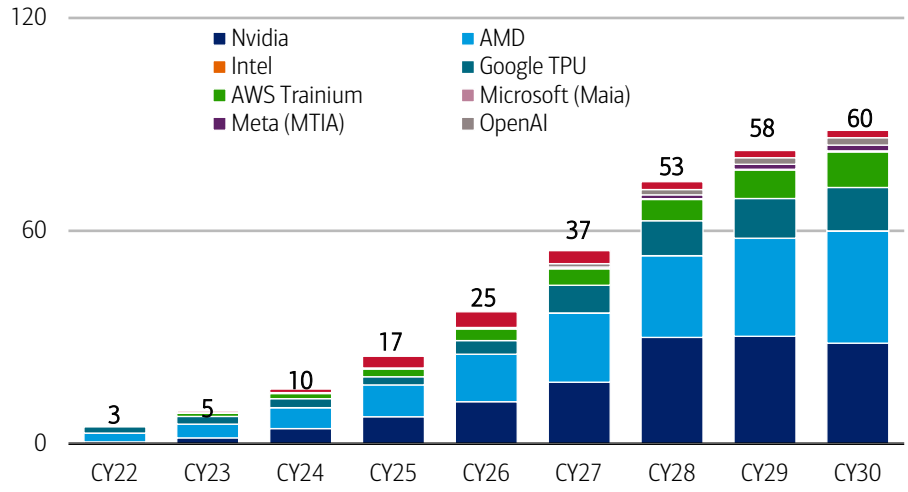


Source: BofA Global Research, Gartner

BofA GLOBAL RESEARCH

Our bottoms-up analysis of accelerator and rack demand for both merchant GPUs and custom ASICs suggests that a cumulative 233GWs needs to be added between CY25-30. GW adds could escalate from ~17GW in CY25 to ~60GW by CY30, underscoring the urgency to move to new power delivery architecture.

Exhibit 7: We estimate AI accelerator demand implies that GWs installed per year for compute could 4x from 15GW in CY25 to 60GW in CY30, a cumulative 233GWs deployed during the period
Implied GWs deployed across major accelerator platforms CY24-30



Source: BofA Global Research estimates, Nvidia, AMD, Intel, Amazon, Google

BofA GLOBAL RESEARCH

How power is delivered today

We have established why power density is rising with each subsequent rack generation and how that is driving a massive inflection in implied GW demand for compute. Now, we move on to evaluating the current state of infrastructure and why it is ill-equipped to address the power demands of tomorrow.

Grid-to-chip today: multiple conversions leading to higher inefficiency

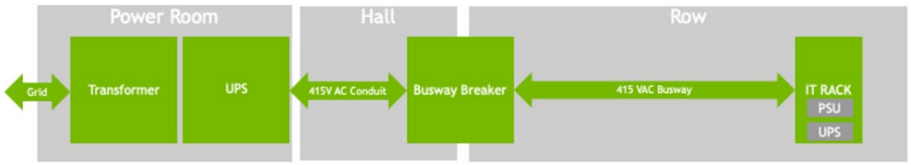
Distribution begins from a generation source, i.e. a power plant (gas turbine, nuclear, solar, etc.) which is then stepped up to ultra-high voltage (230-765kV) via a transformer for long distance transmission. It is better to keep electricity in high voltage form for as long as possible due to greater efficiency. This is because power (P) equals voltage (V) times current (I) and resistance leads to losses which scale with current squared (see below). For the same delivered power, doubling voltage halves current and quarters the losses. Therefore, high voltage for as long as possible is ideal for distribution.

$$P\ (Power) = V(Voltage) \times I(Current)$$

$$P_L\ (Power\ Loss) = I(Current)^2 R(Resistance)$$

Today, utility power is largely centered around alternating current (AC) due to legacy infrastructure, cheaper and more reliable transformers when compared to direct current (DC), and the need for costly and more complex power electronics when dealing with DC. However, DC is increasingly being adopted in the grid for long-range transmission due to its superior efficiency (lower losses) and advances in technology which make handling DC more economical (e.g. solid-state transformers) as we will discuss later. Once the high-voltage AC arrives at a data center campus, it gets stepped down at a substation to medium voltage (10-35kV), then again at on-site transformers to 480V or 415 three-phase AC and then distributed to the data hall. Inside the data hall, each rack receives 480VAC and converts it at the rack-level into 48V or 54V DC using a Power Supply Unit (PSU). This DC is bussed to each server, then converted 1-2x more to a lower voltage until it finally has been stepped down to the sub-1V rail a GPU/XPU core can run on.

Exhibit 8: Today, power from the grid undergoes multiple conversions from high voltage AC to medium voltage AC and then multiple low voltage DCs before reaching the accelerator
Traditional 48V/54V power distribution architecture



Source: Nvidia

BofA GLOBAL RESEARCH

The problem: this architecture is beginning to hit its physical limits

There are multiple problems with this method of distribution. (1) **Space constraints**: a GB300 NVL72 rack possesses up to eight power shelves to support all the computing and switching shelves. Using the same 54 VDC distribution would occupy 64 U of rack space for a Kyber rack (Rubin Ultra and beyond) leaving limited room for compute. (2) **Excessive copper**: using 54 VDC for a single 1 MW rack requires up to 200kg of bulky copper busbar to shuttle electricity which isn't sustainable at the GW-scale; (3) **Conversion inefficiency:** multiple AC/DC conversions and step downs in voltages are not energy efficient (roughly 1-2% lost per transformation) and increase the number of failure points. An entirely new power distribution architecture is necessary in the future.

800 VDC: power delivery for the GW era

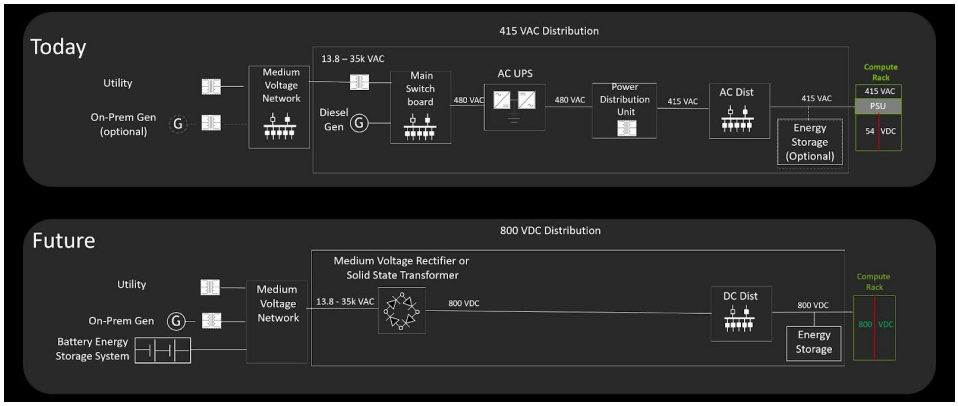
As data centers evolve to support GW-scale capacity, an entirely new architecture for the power flow is vital that is purpose-built to address the limitations of traditional 48V/54V infrastructure while being future-proofed for inevitable migrations to racks with higher power density long-term. This is where 800-volt direct current (800 VDC) comes in.

800V architecture: upgraded infrastructure from grid-to-core

The crux of the issue for traditional power distribution is the multiple voltage conversions and inherent mechanical complexity. For 800V, the 13.8kV AC power that arrives from the grid is directly converted to 800 VDC when it arrives at the campus (eliminating the conversion to medium voltage) using rectifiers (device that converts AC to DC). This transformation that happens in the power room cuts down on multiple AC to DC and DC to DC conversions that minimizes inefficiency losses. System complexity is also reduced as fewer power supply units with fans are needed helping reliability, lower heat dissipation, and improving energy efficiency.

At the row-level, 800V power delivery units (PDUs)/busways carry the electricity rather than 415V AC in the previous model which enables up to 85% more power to be transmitted through the same conductor size. The lower current demands allow the conductor to be thinner thus also drastically cutting the amount of copper wiring needed (by 45% per Nvidia). Using DC rather than AC also helps circumvent AC-specific flaws like “skin effect” and reactive power losses.

Exhibit 9: 800V DC architecture minimizes energy losses and supports higher reliability by making fewer AC to DC and DC to DC conversions along with reducing system component complexity
Existing power flow architecture vs. next generation 800V DC power distribution



Source: Nvidia

BofA GLOBAL RESEARCH

From the power source or separate power rack, two distinct conductor 800V feeds use DC to DC conversion via multiple power supply units (PSUs) to 54v in the IT rack which is then distributed to compute trays for additional DC to DC conversions before reaching the GPU device. By eliminating prior intermediary conversions, there is more space for compute resources (i.e. what actually generates revenue) within the rack itself.

Benefit #1 of 800 VDC: Scalability and futureproofing

800V architecture supports futureproofing for roadmap power density scaling as racks increase from 100kW capacity to over 1MW and beyond while using the same infrastructure. It also supports higher-density GPU clusters and enables higher performance per accelerator driving greater compute throughput and revenue potential.

Benefit #2 of 800 VDC: Improved energy efficiency

Compared to traditional systems where efficiency can be under 90%, end-to-end efficiency could be boosted up to 5% when compared to 54V systems, per Nvidia, by eliminating redundant conversion steps and reducing waste heat in the process

Benefit #3 of 800 VDC: Lower costs and copper material

The same wire gauge can carry up to 157% more power using 800 VDC vs. 415 VAC. To illustrate this, imagine that the current draw for legacy 48V architecture is 12,500 amps (600kW/48V). Leaving 1% of the voltage budget for resistive losses only leaves 38 micro-ohms across the delivery path which means a standard copper busbar which supports 1,000 amps per square inch would need 12 square inches of cross section extending to the full rack height and weighing up to 200kG per rack. However, upgrading to 800 VDC entails a current draw of only 750 amps, a resistance budget of milli-ohms (1000x better), power losses drop by over 10x, and copper material per rack is significantly lower. Further, fewer conductors and connectors would be required for 800 VDC as delivery would move to a simpler three-wire setup instead of the four needed for AC. Thus, copper use, overall material, and installation costs are reduced.

Benefit #4 of 800 VDC: simpler and more reliable architecture, TCO

Centralizing and simplifying power conversion also enhances system reliability and serviceability when compared to traditional IT racks as it lowers the number of critical components such as transformers and phase-balancing equipment . Maintenance costs may be reduced by up to 70% due to fewer PSU failures and savings on labor costs for upkeep. These improvements throughout the revamped power flow accumulate to up to a 30% improvement in total cost of ownership (TCO) according to Nvidia.

Exhibit 10: The 800 VDC power conversion flow is given fewer current and voltage transformation steps along with a simplified distribution path
Power conversion chain: 800 VDC vs. legacy 54V and why the new architecture is superior at any given stage

Stage	800 VDC architecture	Legacy 54 V architecture	Why 800 VDC is better
Medium Voltage (MV) Grid Feed	Conversion: None V / I: 13.8 kVAC 3ϕ / ~42 A Device class: MV switchgear	Conversion: None V / I: 13.8 kVAC 3ϕ / ~42 A Device class: MV switchgear	No difference at this point; both start from the same MV AC utility feed.
Front-end conversion	Conversion: 13.8 kVAC to 800 VDC V / I: 800 VDC / ~1,250 A Device class: Centralized AC/DC rectifier, SST, front-end power block	Conversion: 13.8 kVAC to 480 VAC V / I: 480 VAC / ~1,203 A Device class: Facility transformer	Starts DC distribution earlier and reduces downstream AC transformation / conversion stages.
Backup / energy buffer	Conversion: DC-coupled backup on 800 V bus V / I: 800 VDC / ~1,250 A Device class: DC breaker, solid-state protection, battery / supercap interface	Conversion: 480 VAC to DC link to 480 VAC V / I: 480 VAC / ~1,203 A Device class: Double-conversion UPS	Avoids repeated AC/DC/AC conversion for backup; fewer loss points and simpler power path.
Facility distribution	Conversion: None after front-end rectification V / I: 800 VDC / ~1,250 A Device class: DC switchgear, busway, switchboard, DC protection	Conversion: 480 VAC to 208 VAC V / I: 208 VAC / ~2,776 A Device class: PDU transformer, AC branch distribution	Higher voltage lowers current, reducing cable bulk, copper mass, and I^2R losses.
Rack / sidecar input	Conversion: None at rack input V / I: 800 VDC / ~1,250 A Device class: HVDC input, pre-charge, hot-swap, disconnect, protection	Conversion: 208 VAC to 54 VDC V / I: 54 VDC / ~18,519 A Device class: AC/DC rack PSU shelf, 54 V busbar	Avoids creating a massive ~18.5 kA rack-level 54 V bus, reducing busbar and connector burden.
Intermediate Bus Conversion (IBC) stage — option A: direct 12 V bus	Conversion: 800 VDC to 12 VDC V / I: 12 VDC / ~83,333 A Device class: Isolated HV DC/DC, fixed-ratio or resonant IBC	Conversion: 54 VDC to 12 VDC V / I: 12 VDC / ~83,333 A Device class: LV high-current DC/DC IBC	Keeps the 12 V high-current domain closer to compute instead of distributing 54 V at very high current across the rack.
IBC stage — option B: 50/54 V bridge bus	Conversion: 800 VDC to 50/54 VDC V / I: 50/54 VDC / ~18.5–20.0 kA Device class: Isolated HV DC/DC or power delivery board; may feed 12 V or 6 V next	Conversion: Already at 54 VDC after rack PSU V / I: 54 VDC / ~18,519 A Device class: 54 V busbar feeding downstream DC/DC	Preserves 48/54 V ecosystem compatibility while moving the high-current 54 V domain closer to the load.
IBC stage — option C: final-evolution 6 V bus	Conversion: 800 VDC to 6 VDC V / I: 6 VDC / ~166,667 A Device class: Isolated high-ratio bus converter located very close to compute	Conversion: Typically 54 VDC to 12 VDC before VRM V / I: 12 VDC / ~83,333 A Device class: LV high-current IBC plus downstream VRM	Collapses the downstream chain to 800 V to 6 V to sub-1 V, delaying ultra-high current until very close to silicon.
Voltage Regulator Module (VRM) / point-of-load (POL)	Conversion: 12 VDC or 6 VDC to <1 V V / I: ~0.8 V / ~1.25M A aggregate Device class: Multiphase buck VRM / POL power stage	Conversion: 12 VDC to <1 V V / I: ~0.8 V / ~1.25M A aggregate Device class: Multiphase buck VRM / POL power stage	Final rail is still sub-1 V; the advantage is upstream, where 800 V reduces high-current distribution before the VRM zone.

Source: BofA Global Research, Nvidia, Infineon, Schneider; 3ϕ = 3-phase power, V = Voltage, I = Current

BofA GLOBAL RESEARCH

Transition timing: occurs in phases to give ecosystem time to develop

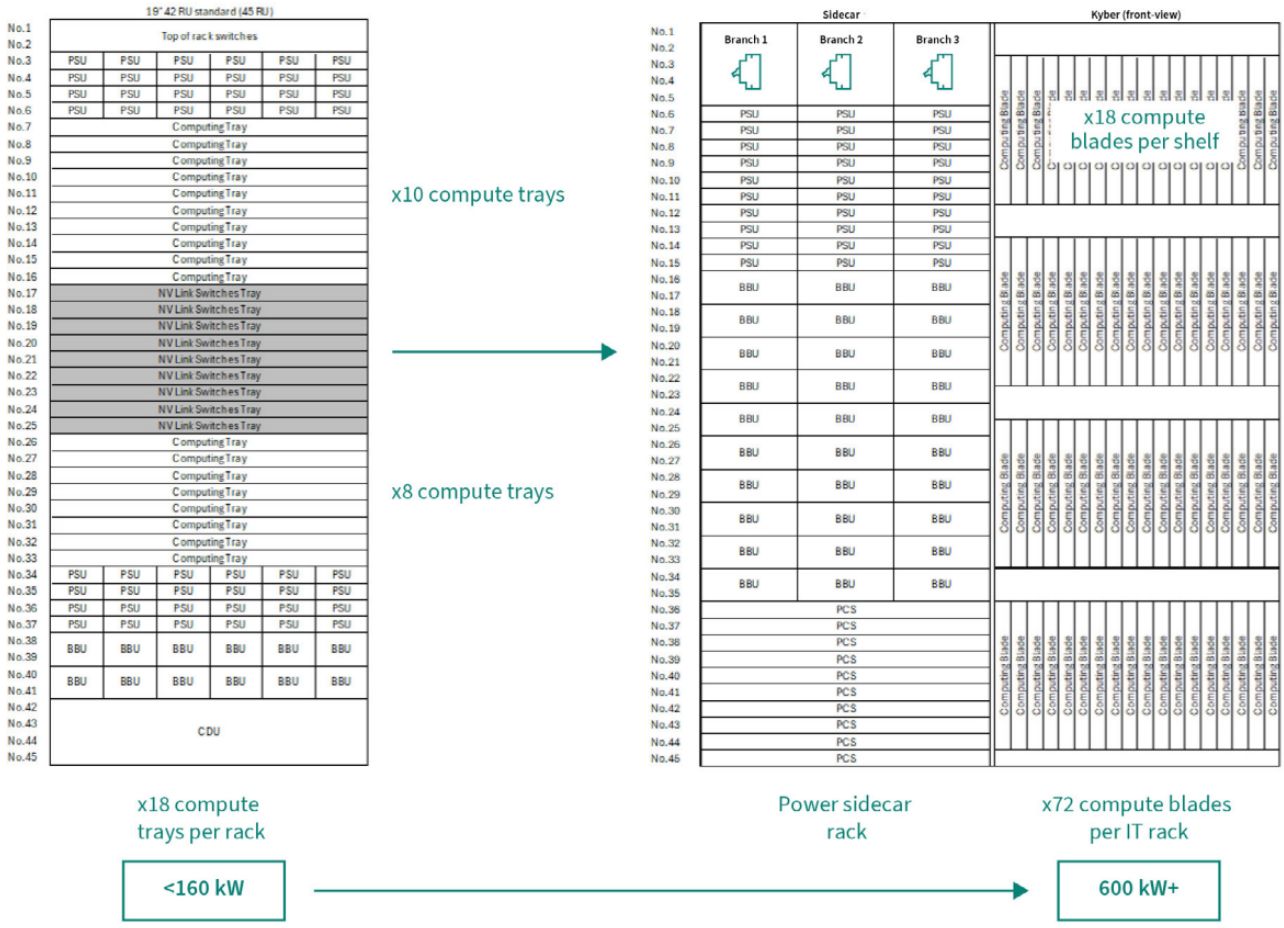
800 VDC architecture is the optimal path forward as it provides scalable end-to-end integration, can already leverage the learning cycles from deploying 800V in the electric vehicle and utility-scale solar industries (hence why auto/industrial semis well positioned), and benefit from maturing silicon carbide (SiC) and gallium nitride (GaN) ecosystems. However, there might be multiple stages before data centers end up at the final state of 800 VDC systems in the future (exhibit 12).

Existing 415 VAC infrastructure involves medium voltage step-down transformers, low voltage switchboards, AC uninterruptible power supply (UPS), PDUs, and AC distribution panels (i.e. busways). Power arrives at each compute rack in 415 VAC form where a PSU integrated in the rack converts it to the 54 VDC for the various compute trays. External energy storage could be leveraged to manage power swings from GPU workloads.

The first transitional architecture is white space retrofit (second example in exhibit 11) where AC to DC conversion is placed outside the rack. This flow utilizes an 800 VDC side power rack which contains rectifiers and PDUs to convert from AC to DC local to each rack. While not the optimal solution, this is a necessary component until 800 VDC grade equipment is ready to deploy and allows compute rack density to risk by offloading power conversion from the IT rack.

Exhibit 11: By offloading power conversion and distribution to a dedicated side car that integrates power supply units (PSUs), battery back up units (BBUs), and power conversion system (PCS), Nvidia's Kyber configuration is able to prioritize more space for compute in its IT rack

Example of a disaggregated Kyber IT rack with a dedicated power side car (vs. a standard 160kW rack)

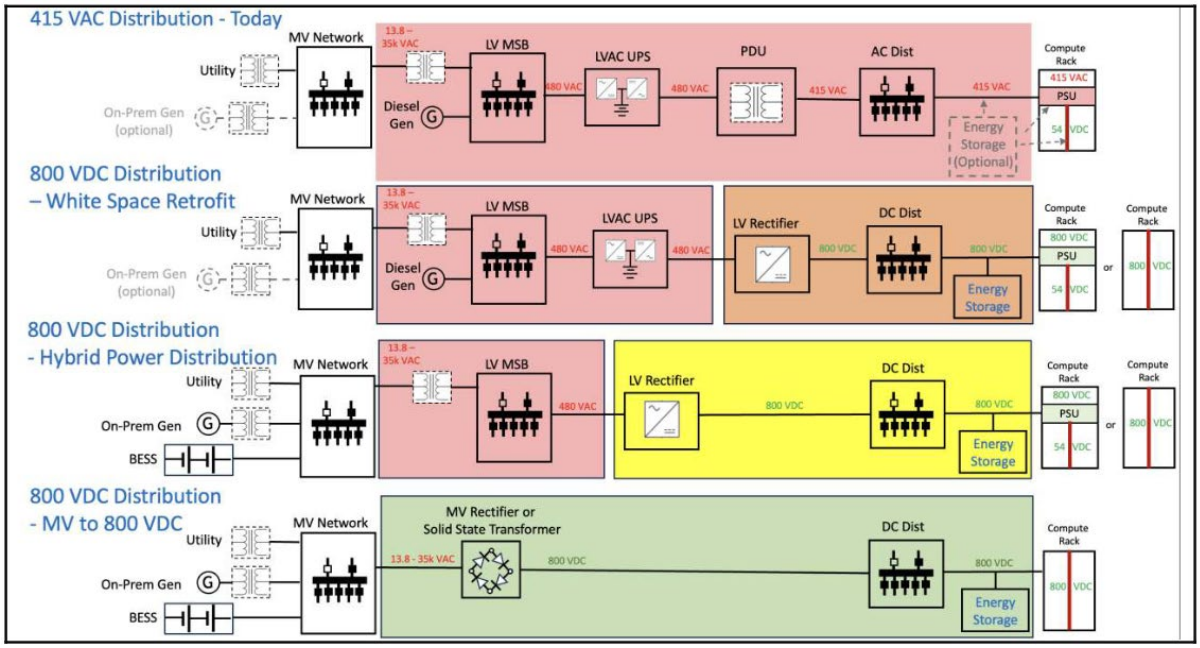


Source: Infineon "Powering the future of AI" white paper

BofA GLOBAL RESEARCH

In the third evolution (hybrid power distribution in exhibit 12), AC to DC power rectifiers at the facility-level deliver 800 VDC from low-voltage AC (<600 V). These reduce rack-level complexity and help integrate energy storage at the DC output. While a low voltage transformer is still needed, this flow works as it relies on existing building blocks (MV/LV transformers, low-voltage switchboards, while using mature large rectifiers common in battery storage and renewables technology.

Exhibit 12: As existing facilities migrate to high voltage distribution, we expect there to be several intermediary stages before "true" 800 VDC architecture is architecture including using a power side car
Evolution of power delivery architecture towards the final end state of true 800 VDC distribution



Source: Nvidia 800 VDC white paper

BofA GLOBAL RESEARCH

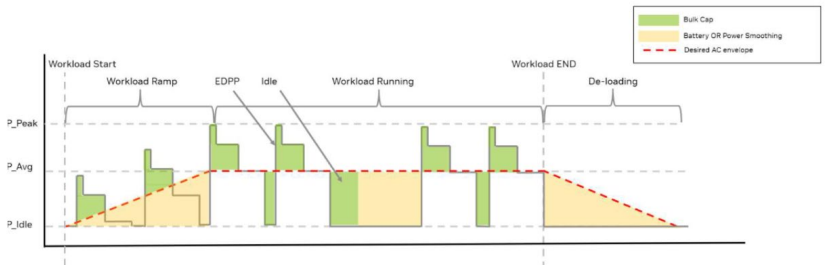
Next generation: MV AC straight to 800 VDC using MV rectifiers or SSTs

The next generation of architecture could incorporate a medium voltage rectifier which converts 35 kvAC MV AC directly to 800 VDC which is then distributed to the racks (MV to 800 VDC example in exhibit 12). This is ideal as it eliminates multiple middle layers (no 480 VAC layer, less low-voltage AC switch gear, fewer conversion stages, simpler infrastructure flow). As an alternative to the MV rectifier which is a proven, high-capacity system widely used today in various industrial markets, Solid State Transformers (SST) could emerge as another strong technology to facilitate this conversion to MV. SSTs factor in a MV AC intake, power rectification modules, and DC distribution in a highly compact and dense design which provides significant space savings. There are challenges to scaling SSTs at the multi-MW scale today due to reliability, thermal issues, and electrical stress which is why MV rectifiers would likely be adopted first. However, SSTs represent a critical element of the "end state" for 800 VDC as the tech matures.

Workload volatility: solved by moving to multi-timescale energy storage

There is another challenge when scaling AI data centers that 800V alone cannot solve which is workload volatility. Training an LLM involves thousands of GPUs executing multiple cycles of intensive computation followed by periods of data exchange in almost perfect synchronization as a single system. These dynamics lead to massive and rapid load swings across the facility which can lead to grid-scale oscillations in the power profile. In milliseconds, a rack can oscillate from 30% utilization to 100% and back again forcing components to be engineered as oversized to handle peak current needs which drives up the cost and physical footprint and threatening the overall stability of the grid.

Exhibit 13: Power needs to be managed at multiple timescales when dealing with AI workloads including grid level fluctuation requirements and GPU load demand requirements
Power envelope over a full compute cycle



Source: Nvidia 800 VDC white paper

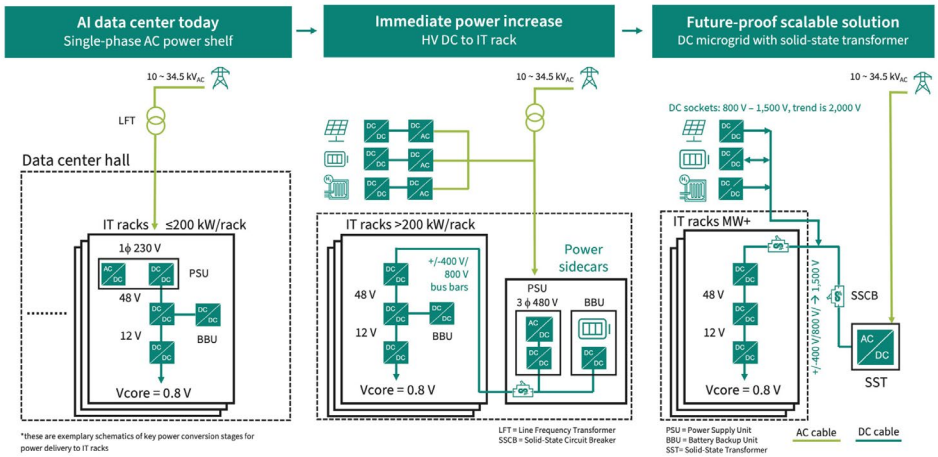
BofA GLOBAL RESEARCH

To solve this problem, a buffer needs to be created that separates the volatile power demands of the AI infrastructure from the stability needs of the grid. This is where multi-timescale energy storage comes in as power fluctuations occur across multiple timescales. For short-duration storage (milliseconds to seconds) high power capacitors and supercapacitors sit next to the compute racks and absorb high-frequency power spikes while filling brief dips created by idle workload periods. Long-duration storage (seconds to minutes) is large, facility-level battery energy storage systems (BESS) located at the interconnection with the grid. These manage more gradual but larger-scale power shifts as entire workloads ramp up or down

Hybrid microgrid: the final frontier for HV facility infrastructure

To holistically optimize the power distribution from energy generation to consumption at the GW scale, a radically different set of infrastructure is needed. DC microgrids removes AC to DC power flow and distributes high voltage straight from the grid via an HV bus that connects to the server racks. This is seen as a future-proof and "end state" type of solution where new technologies like SSTs, solid-state circuit breakers (SSCB), and high-grade energy storage systems/uninterruptible power supplies (ESS/UPS) will play an important role. By simplifying the flow and replacing bulky components with more modular tech, reliability is also enhanced while boosting overall efficiency. Retrofits are difficult, so we would expect microgrids to come from greenfield site development commencing towards the end of the decade (CY28-30).

Exhibit 14: DC microgrids optimize system efficiency through use of new equipment like SSTs, SSCBs, and more, representing an end state for the shift to 800 VDC architecture
Data center architecture evolution



Source: Infineon

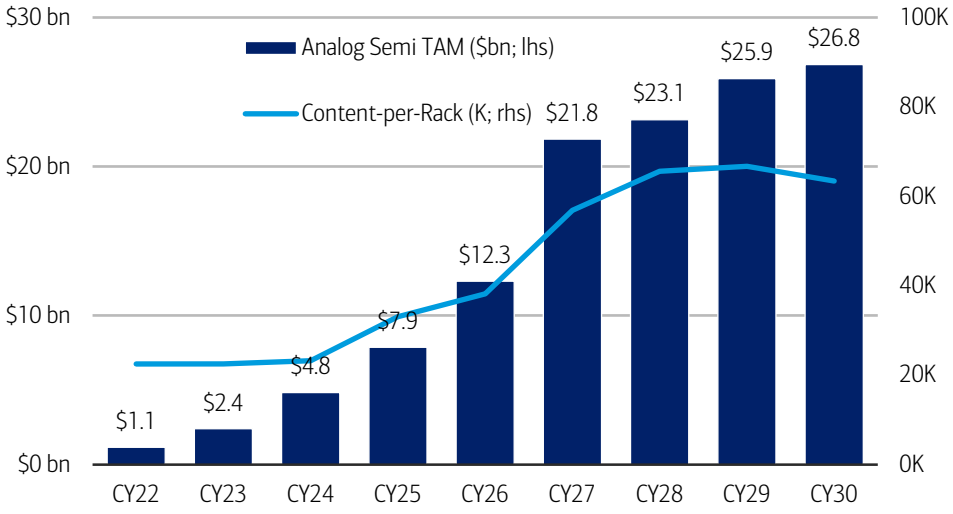
BofA GLOBAL RESEARCH

Analog semis: key beneficiaries in high power

We believe the analog semi space will see transformational growth in their AI-related sales as the power tree becomes much more semiconductor intensive. In order for data centers to enjoy the benefits of higher efficiency, greater power density, and the maximization of physical space for compute, we think value will begin to accrue towards suppliers of active power conversion, protection, sensing, control, and high-voltage DC components within the data center rather than just bulky transformers, PSU shelves, or busbars. Major suppliers should enjoy content gains from unit growth, new socket/market expansion, and rising ASPs (sometimes 2-3x higher) as every layer of the power flow needs to be upgraded from facility power the processor rail.

We believe the ultimate winners in the space will (1) boast system-level expertise that can tackle the entire power tree (portfolio breadth) rather than single-point product vendors, (2) offer products that can adequately handle the demands of high voltages in an exceptionally reliable and safe way, (3) offer flexibility regardless of what architecture and topology wins the day, (4) support multiple device types (especially wide-band gap semis), and (5) be those that can provide reference designs and co-design support while being anchored to the main leaders in the ecosystem (e.g. Nvidia).

Exhibit 15: We expect the AI TAM for analog semis to grow 28% 5-year CAGR through CY30 to $27bn from $7.9bn in CY25. Average content per rack could grow from mid-$20K to over $60K by CY30
Analog semi data center and strategic power infrastructure TAM ($bn) and average content-per-rack ($)



Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia

BofA GLOBAL RESEARCH

Our model: $7.8bn TAM in CY25 surging to $27bn in CY30 led by high power

We developed a bottoms-up industry demand model that sizes the analog semi TAM for AI across company revenue/share, device type, and component type within the data center for both low power (100-160kW, <100kW) and high power (>600kW, 1MW+) racks. Underpinning these forecasts are estimates on implied GWs deployed for AI based on expectations for XPU/GPU rack demand over the next few years.

Our TAM is segmented into two key categories: (1) data center (i.e. within the data hall at the server/rack level) and (2) power infrastructure (grid, utilities, core infra). We flag that our power infra-TAM focuses on the most "strategic" parts of the market which are still emerging (e.g. SSTs, solid-state circuit breakers) and likely underestimates the true sales related to infrastructure benefitting from AI (e.g. solar, transmission, etc.)

We estimate the analog AI market to be $7.9bn as of CY25, expect it to grow to $12bn in CY26, and predict a surge to $28bn by CY30 (28% CAGR) as high power racks become a larger part of the mix. Average content-per-rack grows from mid-$20K up to $650K during this time as both unit attach and ASPs rise considerably.

As AI racks scale toward 1MW, the role of various device types enhances. Silicon remains the lower-voltage workhorse, SiC moves into high-voltage conversion/protection, GaN enables dense DC/DC near compute, analog adds safety/control, MCUs coordinate distributed power, and sensors provide real-time visibility. While analog ICs are the largest market, growing to $14bn by CY30, we think that silicon carbide (SiC) and gallium nitride (GaN) should gain material share as power density becomes critical.

Per exhibit 18 and 19 (full model in back), we estimate TXN has the highest share in AI semis currently and should retain that title through CY30. Infineon sees the biggest share jump and should become the second largest AI vendor by the end of the decade. Close behind is analog leader ADI while we also note ON's sizable share expansion.

Exhibit 16: Every device type likely plays a larger role in the move to 1MW class racks. While analog ICs represent the largest market today and by CY30, we note that wide-bandgap semis like SiC and GaN are the fastest growing at 63%/69% CY25-30 CAGR, respectively
Comparing roles in the data center, market size, and CAGRs for major semiconductor device types

Semiconductor type	Role in today's data center	Role in 1MW/800VDC racks	Key vendors/ companies	AI TAM size CY25 ($mn)	AI TAM size CY30 ($mn)	% of TAM CY25	% of TAM CY30	CY25-30 CAGR
Silicon discrete power	Workhorse for MOSFETs, synchronous rectification, load switching, protection, lower-/mid-voltage DC/DC, and many VRM power stages.	Remains the cost-effective backbone in mature lower-voltage stages; some premium high-voltage / high-density sockets shift to SiC or GaN.	onsemi; STMicroelectronics; Infineon; Alpha and Omega Semiconductor; ROHM; Monolithic Power Systems; Texas Instruments; Analog Devices	$1,791	$4,780	23%	18%	22%
Silicon carbide / SiC	More common in EV, industrial, renewables, UPS, and high-voltage power than in traditional server boards.	Moves upstream into high-voltage conversion, 800V protection, power-rack / sidecar stages, MV rectifier/SST concepts, and rugged hot-swap or pre-charge.	onsemi; STMicroelectronics; Infineon; ROHM; Wolfspeed	$183	$2,098	2%	6%	63%
Gallium nitride / GaN	Used where high frequency, small magnetics, high efficiency, and power density matter.	Becomes important for dense DC/DC conversion close to compute, including 800V to 50/54V, 800V to 12V, and 800V to 6V paths.	Navitas; STMicroelectronics; Texas Instruments; Infineon; Power Integrations; Innoscience	$118	$1,612	1%	6%	69%
Analog ICs	Controllers, drivers, hot-swap, current/voltage sense, isolation, telemetry, PMICs, supervisors, and digital power interfaces.	Becomes the safety and control layer: inrush control, fault detection, isolation, hot-swap, diagnostics, telemetry, and fail-safe operation.	Analog Devices; Texas Instruments; Microchip; Infineon; Renesas; Monolithic Power Systems	$5,198	$15,919	66%	59%	25%
MCUs	Local control for PSUs, fan/power management, telemetry, fault logging, and digital power loops.	Becomes distributed power orchestration across rectifiers, sidecars, IBCs, VRMs, backup energy storage, and thermal systems.	STMicroelectronics; Texas Instruments; Microchip; NXP; Renesas; Infineon	$332	$1,516	4%	6%	36%
Sensors	Current, voltage, temperature, airflow, pressure, leak, and power/thermal monitoring across servers and infrastructure.	Becomes the real-time visibility layer for MW-class racks: current, voltage, temperature, transient detection, fault prediction, and cooling health.	Allegro MicroSystems; STMicroelectronics; Infineon; Renesas;	$242	$867	3%	3%	29%

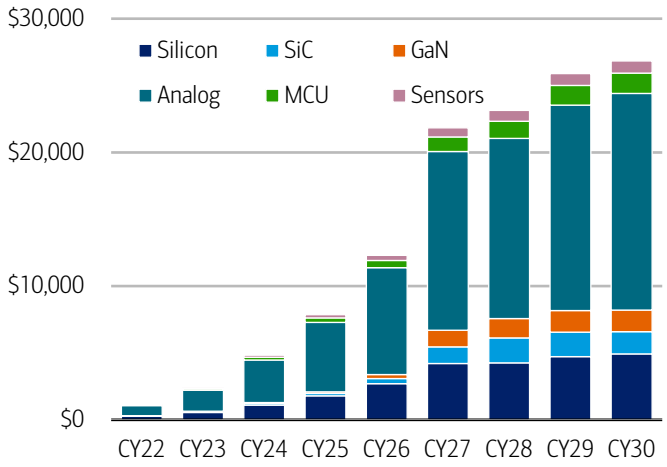
Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia

BofA GLOBAL RESEARCH

Exhibit 16: Every device type likely plays a larger role in the move to 1MW class racks. While analog ICs represent the largest market today and by CY30, we note that wide-bandgap semis like SiC and GaN are the fastest growing at 63%/69% CY25-30 CAGR, respectively
Comparing roles in the data center, market size, and CAGRs for major semiconductor device types

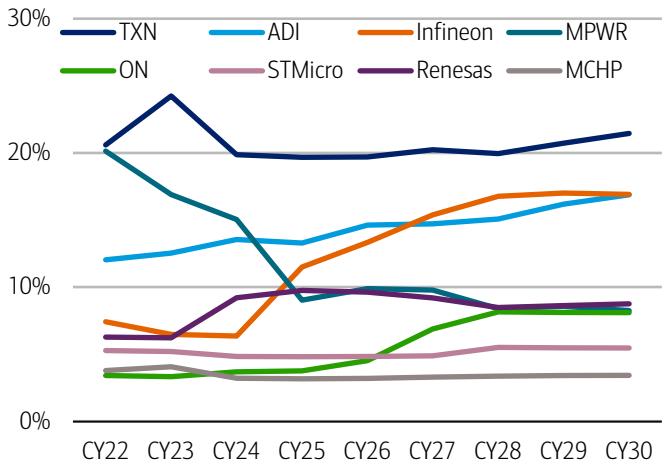
Semiconductor type	Role in today's data center	Role in 1MW/800VDC racks	Key vendors/ companies	AI TAM size CY25 ($mn)	AI TAM size CY30 ($mn)	% of TAM CY25	% of TAM CY30	CY25-30 CAGR

Exhibit 17: While analog ICs will continue to be the largest market, we believe SiC and GaN will enjoy the most share gains
AI analog semi TAM by device type



Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia
BofA GLOBAL RESEARCH

Exhibit 18: TXN enjoys dominant share in the AI TAM but we flag that discrete suppliers ON and Infineon gain significant share
Analog semi vendor AI market share



Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia
BofA GLOBAL RESEARCH

Exhibit 19: We estimate that AI-related (data center/power infra) revenue across the analog space grew +63% YoY in CY25 to $7.9bn and could grow +56% YoY in CY26 to $12bn. We believe the market will accelerate again in CY27 when high-power Kyber racks for Rubin Ultra become available on the market
Estimated analog semi supplier sales across data center and strategic power infrastructure CY22-30

Company Revenue Summary ($; mn)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total	**$1,147**	**$2,395**	**$4,829**	**$7,864**	**$12,302**	**$21,837**	**$23,094**	**$25,847**	**$26,793**
TXN	$236	$580	$959	$1,547	$2,416	$4,368	$4,549	$5,277	$5,658
ADI	$138	$300	$654	$1,045	$1,794	$3,173	$3,424	$4,106	$4,439
Infineon	$85	$155	$307	$904	$1,647	$3,415	$3,949	$4,490	$4,630
MPWR	$231	$404	$726	$710	$1,212	$2,110	$1,915	$2,179	$2,176
ON	$39	$80	$178	$295	$566	$1,577	$1,979	$2,219	$2,298
STMicro	$60	$124	$233	$379	$597	$1,090	$1,301	$1,452	$1,504
Renesas	$72	$149	$445	$767	$1,181	$1,986	$1,931	$2,194	$2,306
MCHP	$43	$97	$155	$249	$393	$716	$764	$865	$901
Other	$242	$504	$1,172	$1,969	$2,496	$3,402	$3,282	$3,065	$2,880
Low Power/Voltage	**$1,147**	**$2,394**	**$4,829**	**$7,864**	**$10,862**	**$10,790**	**$8,530**	**$9,527**	**$10,739**
TXN	$236	$580	$959	$1,547	$2,172	$2,226	$1,799	$2,055	$2,369
ADI	$138	$300	$654	$1,045	$1,617	$1,718	$1,449	$1,713	$1,983
Infineon	$85	$155	$307	$904	$1,413	$1,498	$1,246	$1,456	$1,665
MPWR	$231	$404	$726	$710	$1,090	$1,131	$805	$899	$959
ON	$39	$80	$178	$295	$475	$568	$492	$593	$716
STMicro	$60	$124	$233	$379	$525	$527	$425	$478	$543
Renesas	$72	$149	$445	$767	$1,058	$1,048	$784	$875	$986
MCHP	$43	$97	$155	$249	$344	$342	$271	$302	$341
Other	$242	$504	$1,172	$1,969	$2,169	$1,733	$1,259	$1,156	$1,176
High Power/Voltage	**$0**	**$0**	**$0**	**$0**	**$1,439**	**$11,047**	**$14,564**	**$16,320**	**$16,054**
TXN	$0	$0	$0	$0	$244	$2,142	$2,750	$3,222	$3,289
ADI	$0	$0	$0	$0	$177	$1,455	$1,976	$2,393	$2,455
Infineon	$0	$0	$0	$0	$234	$1,917	$2,702	$3,035	$2,964
MPWR	$0	$0	$0	$0	$122	$979	$1,109	$1,280	$1,217
ON	$0	$0	$0	$0	$91	$1,009	$1,486	$1,626	$1,583
STMicro	$0	$0	$0	$0	$72	$563	$876	$974	$961
Renesas	$0	$0	$0	$0	$123	$939	$1,147	$1,319	$1,321

Exhibit 19: We estimate that AI-related (data center/power infra) revenue across the analog space grew +63% YoY in CY25 to $7.9bn and could grow +56% YoY in CY26 to $12bn. We believe the market will accelerate again in CY27 when high-power Kyber racks for Rubin Ultra become available on the market

Estimated analog semi supplier sales across data center and strategic power infrastructure CY22-30

Company Revenue Summary ($; mn)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
MCHP	$0	$0	$0	$0	$49	$374	$493	$563	$561
Other	$0	$0	$0	$0	$326	$1,669	$2,023	$1,909	$1,703
Company Revenue Summary YoY Growth	**CY22**	**CY23**	**CY24**	**CY25**	**CY26**	**CY27**	**CY28**	**CY29**	**CY30**
Total	**-**	**108.8%**	**101.7%**	**62.9%**	**56.4%**	**77.5%**	**5.8%**	**11.9%**	**3.7%**
TXN	-	145.8%	65.2%	61.3%	56.2%	80.8%	4.1%	16.0%	7.2%
ADI	-	117.4%	118.0%	59.7%	71.7%	76.9%	7.9%	19.9%	8.1%
Infineon	-	82.4%	97.7%	194.5%	82.3%	107.3%	15.6%	13.7%	3.1%
MPWR	-	75.2%	79.5%	-2.2%	70.7%	74.0%	-9.2%	13.8%	-0.1%
ON	-	103.3%	123.7%	65.7%	91.7%	178.7%	25.4%	12.1%	3.6%
STMicro	-	106.0%	87.5%	62.2%	57.7%	82.6%	19.4%	11.6%	3.6%
Renesas	-	107.0%	198.5%	72.5%	54.0%	68.2%	-2.8%	13.6%	5.1%
MCHP	-	124.6%	59.5%	60.5%	57.6%	82.3%	6.7%	13.1%	4.1%
Other	-	108.4%	132.5%	68.0%	26.8%	36.3%	-3.5%	-6.6%	-6.1%
Low Power/Voltage	**-**	**108.7%**	**101.7%**	**62.9%**	**38.1%**	**-0.7%**	**-20.9%**	**11.7%**	**12.7%**
High Power/Voltage	**-**	**-**	**-**	**-**	**-**	**667.6%**	**31.8%**	**12.1%**	**-1.6%**
Company Share (%)	**CY22**	**CY23**	**CY24**	**CY25**	**CY26**	**CY27**	**CY28**	**CY29**	**CY30**
Total	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**
TXN	20.6%	24.2%	19.9%	19.7%	19.6%	20.0%	19.7%	20.4%	21.1%
ADI	12.0%	12.5%	13.5%	13.3%	14.6%	14.5%	14.8%	15.9%	16.6%
Infineon	7.4%	6.5%	6.4%	11.5%	13.4%	15.6%	17.1%	17.3%	17.3%
MPWR	20.1%	16.9%	15.0%	9.0%	9.9%	9.7%	8.3%	8.4%	8.1%
ON	3.4%	3.3%	3.7%	3.8%	4.6%	7.2%	8.6%	8.6%	8.6%
STMicro	5.3%	5.2%	4.8%	4.8%	4.9%	5.0%	5.6%	5.6%	5.6%
Renesas	6.3%	6.2%	9.2%	9.8%	9.6%	9.1%	8.4%	8.5%	8.6%
MCHP	3.8%	4.1%	3.2%	3.2%	3.2%	3.3%	3.3%	3.4%	3.4%
Other	21.1%	21.1%	24.3%	25.0%	20.3%	15.6%	14.2%	11.9%	10.8%

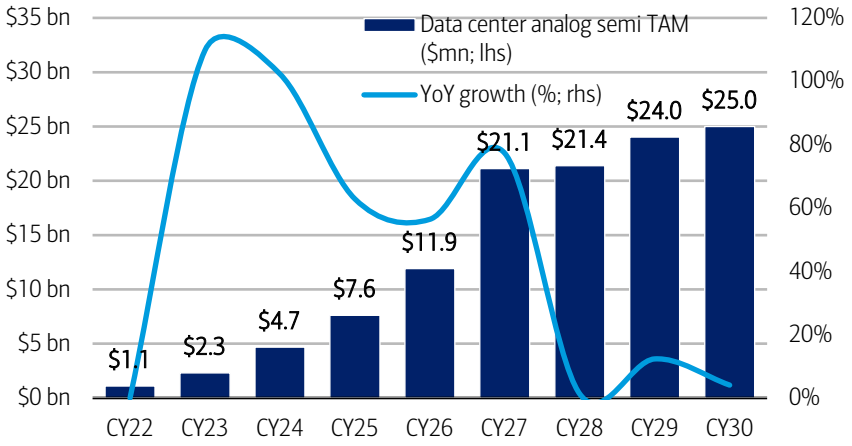
Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia

BofA GLOBAL RESEARCH

Data center: $25bn TAM from rack to core

We estimate the data-center analog semi opportunity grows from $7.6bn in CY25 to $25.0bn in CY30, a ~26% CAGR, with growth near +60-80% CY26/27 as high-power AI racks ramp into CY27. As AI racks scale from sub-100kW systems to 600kW+ sidecar designs and eventually 1MW+ hybrid-microgrid architectures, overall content per rack grows, but the component value mix migrates. In lower power racks, content is spread across conventional rack PSUs, 48V distribution, switching, CPU/GPU board power, optics, sensing, and timing. As racks move into the 600kW+ class, this value shifts towards sidecar power delivery, high voltage intermediate conversion, optics, protection, PCS/transient buffering, and much larger compute board power content. In this phase, the rack is still power supply-heavy, but the PSU evolves from a simple in-rack box to a higher power sidecar system feeding increasingly dense GPU tray.

Exhibit 20: The data center analog TAM at $7.6bn CY25 growing to $25bn by CY30, a ~28% CAGR
Analog semi AI data center TAM



Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia

BofA GLOBAL RESEARCH

While there are no official specs from Nvidia, we assume rack density for 1MW class racks continues toward a 576 GPU package architecture, while centralized SST/HVDC microgrid infrastructure is captured separately. Rack-local PSU, BBU, and PCS content therefore grow less, while incremental value concentrates in HV/MV/LV IBC conversion, GPU-board VRM and vertical power delivery, CPU-complex attached power, optics, HBM/memory power, protection, sensing, and telemetry.

Exhibit 21: We estimate analog semi content will surge from $36K per 100-160kW rack today to almost $300K per rack at 600kW and potentially up to $917K after 1MW
Evolution of component content as racks increase in power capacity

	<100kW per rack		<160kW per rack		600+kW per rack		1MW class racks	
Components	Content	%	Content	%	Content	%	Content	%
Switching Infrastructure	$1,589	9%	$3,284	9%	$9,364	3%	$18,727	2%
PSU	$2,938	17%	$5,760	16%	$29,250	10%	$29,250	3%
BBU	$395	2%	$840	2%	$6,300	2%	$6,300	1%
48V Bus Converter | HV IBC	$1,271	7%	$2,700	7%	$43,200	15%	$216,000	24%
GPU Board	$4,320	25%	$9,720	27%	$68,040	23%	$272,160	30%
Smart NIC cards	$691	4%	$864	2%	$3,600	1%	$10,368	1%
Protection	$424	2%	$900	2%	$7,200	2%	$28,800	3%
Optical Infrastructure	$1,728	10%	$4,401	12%	$43,963	15%	$131,888	14%
HBM PMIC/memory power	$813	5%	$1,728	5%	$8,640	3%	$34,560	4%
CPU complex and power	$1,382	8%	$3,240	9%	$45,360	16%	$136,080	15%
Clocks and timing	$518	3%	$792	2%	$2,592	1%	$5,184	1%
Signal Chain Sensing	$1,210	7%	$2,088	6%	$4,176	1%	$8,352	1%
PCS	$0	0%	$0	0%	$19,500	7%	$19,500	2%
Total	$17,280	100%	$36,317	100%	$291,184	100%	$917,170	100%

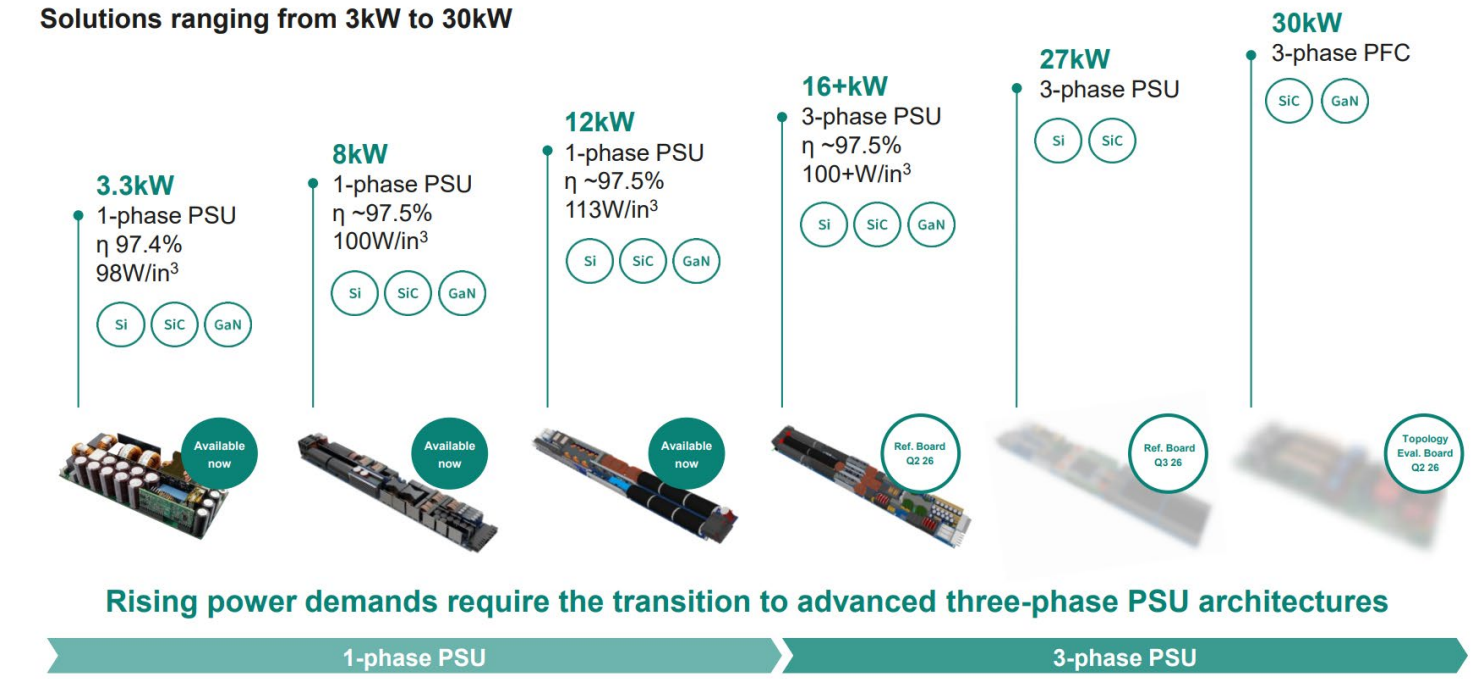
Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia

BofA GLOBAL RESEARCH

Power Supply Units (PSU)

A power supply unit (PSU) is the rack's AC to DC conversion block. It takes data center AC power and produces a stable 48V/50V DC rail that feeds the rack busbar, server trays, and backup units. In today's architecture, multiple PSUs sit in a power shelf inside the rack (exhibit 22). Each PSU includes a PFC (power factor correction) front end which helps to draw power cleanly from the AC grid, an isolated DC to DC stage to create the regulated output, auxiliary/standby power, protection so multiple PSUs can share the bus safely, and control/communications for monitoring and current sharing. Today, 3kW/5.5kW 48V PSUs are the workhorses of 160kW racks, with 12kW single-phase PSUs as an intermediate step as rack power rises.

Exhibit 22: PSUs are evolving from single-phase 3.3kW architectures towards three-phase 30kW devices disaggregated into sidecars
Infineon PSU roadmap



Source: Infineon

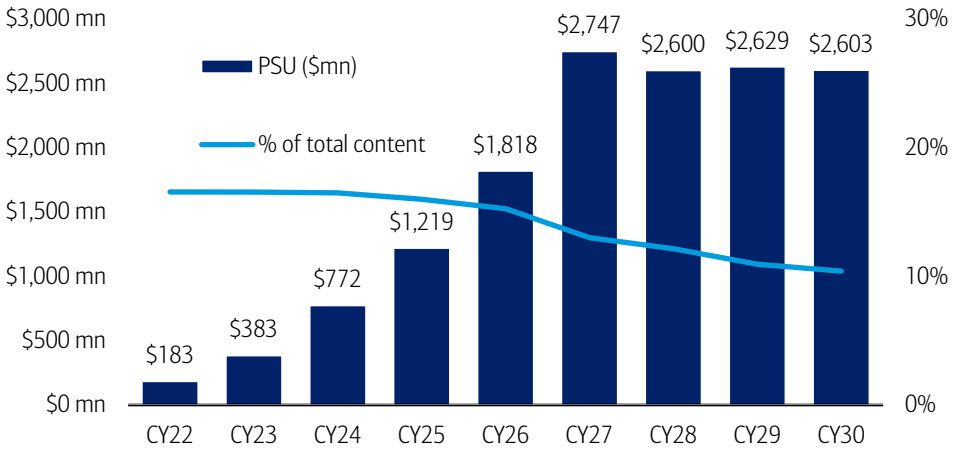
BofA GLOBAL RESEARCH

In higher-power AI racks, the PSU's role changes from a relatively contained in-rack 48V power box into part of a sidecar or high-voltage power system. Along with BBUs, PSUs move out of the IT rack and into sidecars as racks approach 1MW, with three-phase PSUs (for higher power delivery) taking 400V/480V AC and producing 400V or 800V DC. The sidecar PSU can be seen as a transition architecture that converts legacy 480V AC to 800V DC (essentially helping data centers be able to adopt Rubin Ultra), while the longer-term end state moves toward DC distribution where compact IBCs in the IT tray replace much of the traditional rack PSU function. Moving to the three-phase PSUs requires several enhancements to the circuitry. The PFC (power intake) is adapted for three-phase power, DC to DC stages are upgraded for higher voltages, higher performance protection (e.g. hot swap controllers), and faster transient control (being able to react quickly if GPUs suddenly jump from low power to high power or vice versa).

Whereas traditional PSUs rely heavily on advanced silicon MOSFETs, we see a shift towards SiC and GaN (40-50% of content each). SiC is best suited for high voltage front end of larger PSUs (e.g. PFC and AC to DC stages) and high voltage DC to DC stages given its ability to deal with higher voltages with better efficiency and lower losses. GaN is applied where switching speed and power density are critical, including high-frequency PFC and certain DC to DC conversion stages.

Currently, PSUs, make up 16-18% of the TAM but this could decline as this stage of power delivery is moved upstream as microgrid architectures gain traction. Suppliers that are well positioned include Infineon, ON, STMicro (broad discrete platform) and analog vendors ADI and TXN in key sockets like protection.

Exhibit 23: PSU total content could grow from ~$1.2bn today to $2.6bn by CY30
Power Supply Unit analog semi TAM estimates ($ mn; lhs) and % of total content (%; rhs)



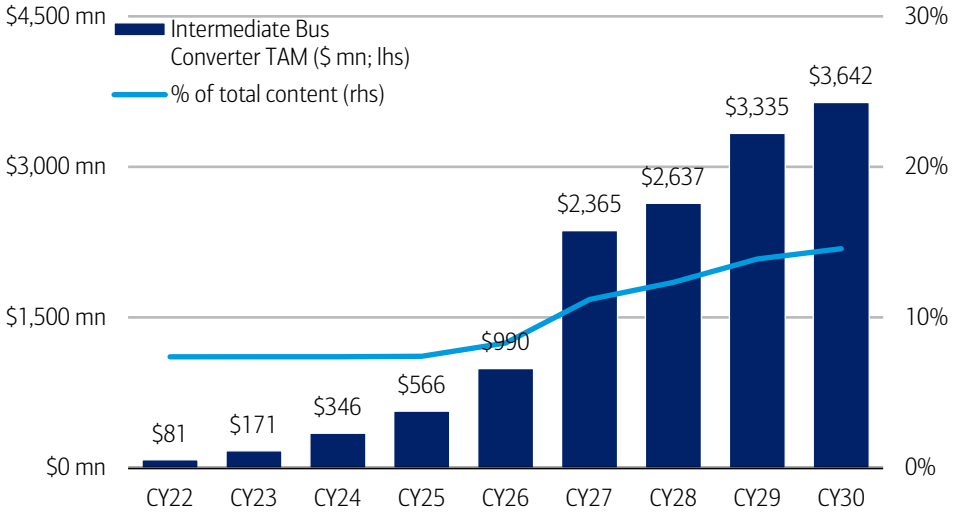
Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia

BofA GLOBAL RESEARCH

Intermediate Bus Conversion (IBC)

An intermediate bus converter (IBC) is the conversion stage between rack-level distribution and the final processor VRM. It creates an intermediate rail the compute tray can use before the sub-1V GPU core rail stage, such as 50V, 12V, or 6V. The legacy 48V/54V IBC architecture was a practical low-voltage rack bus for lower-density systems, but it becomes inefficient as rack power moves from 100–160kW toward 600kW and 1MW+.

Exhibit 24: IBC experiences among the most dramatic increases of content as the component becomes an essential bridge between the high voltage and low voltage flow. We see a $600mn market today growing quickly to $3.6bn by CY30 with content % rising towards >20%.
Intermediate Bus Conversion (IBC) TAM and % of total content



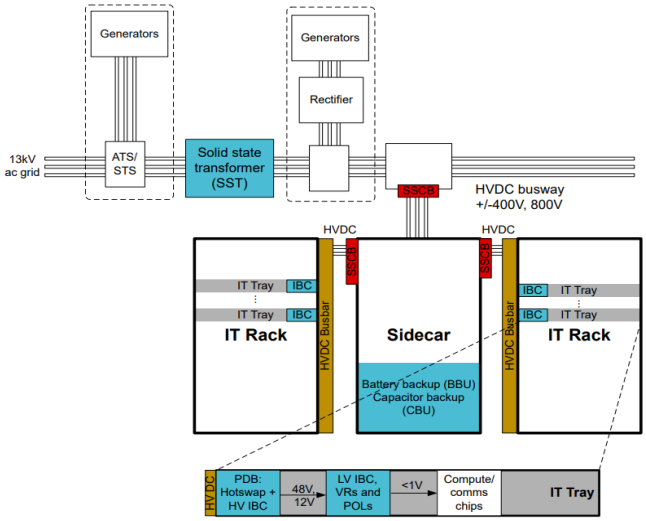
Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia

BofA GLOBAL RESEARCH

In a 100–160kW rack, the IBC function is still mostly a 48V/54V bus-conversion problem, so the value is concentrated in the lower-voltage ecosystem and content is relatively modest ($2.7K/rack per our estimates). In a 600kW+ sidecar rack, value moves into the high voltage (HV) and medium voltage (MV) portions because the rack now needs to receive 800V/400V, protect it, and step it down into usable tray rails, leading to a large step up in value (15% of rack TAM). In a 1MW+ rack, the value shifts even more toward the MV conversion bridge between 800V distribution and the processor tray. At 1MW+, the power per tray and current density are so high that the conversion moves even closer to the compute board with more of an emphasis on 800V to 12V, 800V to 6V, or 50V to 6V stages sitting close to the GPU tray.

Exhibit 25: IBC's role only increases in prominence as we move to next generation power distribution flows, still retaining the important function as high to low conversion (see placed in IT racks below) even as other components move upstream

IBC's role in a final generation AI computing SST and DC distribution

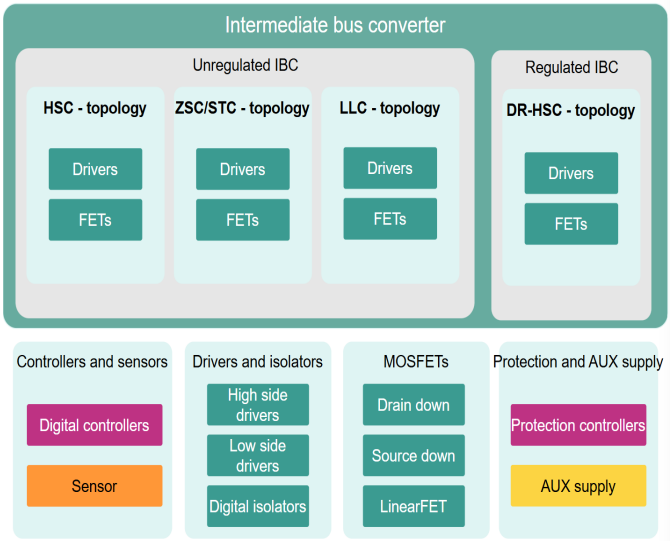


Source: Texas Instruments

BofA GLOBAL RESEARCH

Exhibit 26: While legacy 48V IBC components disappear, these are replaced by higher value HV/MV/LV IBCs that carry more wide-bandgap semi content such as SiC and GaN

48V Intermediate Bus Converter block diagram

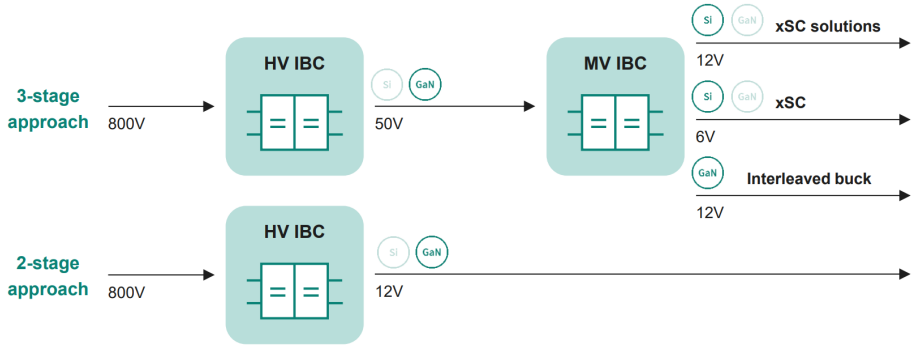


Source: Infineon

BofA GLOBAL RESEARCH

There are several potential architecture paths for the conversion flow, each with tradeoffs. 800V to 12V to core keeps the chain shorter and preserves a familiar board rail, but it puts a large step-down ratio into one IBC which can add stress. 800V to 50V to 6V to core spreads the work across stages, keeps part of the 48V/50V ecosystem alive, and lets the 6V stage sit closer to the processor. 800V to 6V to core is the most direct which reduces stages and lowers the conversion ratio handled by the final VRM, but it forces the 6V converter very close to the GPU board because current is much higher. The location of value therefore moves from rack-level 48V bus conversion toward HVDC sidecar or rack distribution, then into denser conversion modules near the compute tray, and in the most aggressive architecture into ultra-dense 800V-to-6V conversion close to the GPU board.

Exhibit 27: Suppliers (like Infineon below) often offer solutions that support alternative IBC flows such as 3-stage approaches (800V to 50V to 12V/6V) or 2-stage (800V to 12V). Each architecture comes with distinct tradeoffs and benefits.
Potential approaches for IBC

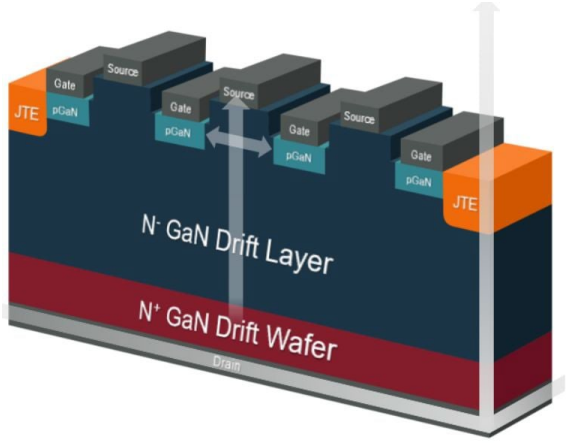


Source: Infineon

BofA GLOBAL RESEARCH

SiC and GaN are specifically useful in IBCs because the converter must combine high ratio conversion, high efficiency, density, and thermal reliability. SiC is most relevant in the high-voltage input stage of an 800V-to-50V or 800V-to-12V IBC, where the devices must withstand the 800V bus, handle transients, and switch efficiently without creating excess heat. GaN is more relevant where the design is optimizing for very high switching frequency, smaller magnetics, and compact high-density conversion closer to the compute tray. Vertical GaN is an emerging GaN-on-GaN power device architecture where current flows through the device vertically rather than laterally across the surface. In AI racks, the relevant sockets IBCs such as 800V-to-50V, 800V-to-12V, or potentially 800V-to-6V, where designers want GaN-like fast switching and density but more voltage headroom and ruggedness than lateral GaN. Silicon remains relevant in lower-voltage secondary-side MOSFETs, synchronous rectification, controllers, drivers, sensing, and protection where cost and maturity matter. We expect GaN to be the leading IBC material for high power racks.

Exhibit 28: ON's vGaN transistors are fabricated on bulk GaN substrates, enabling current to flow vertically through the crystal lattice rather than laterally. Using the same substrate for fabrication (Rather than GaN-on-Si) can lead to improved yield and ruggedness, a sort of cross between SiC and standard GaN. However, manufacturing vGaN is technically challenging.
Vertical GaN example



Source: Onsemi

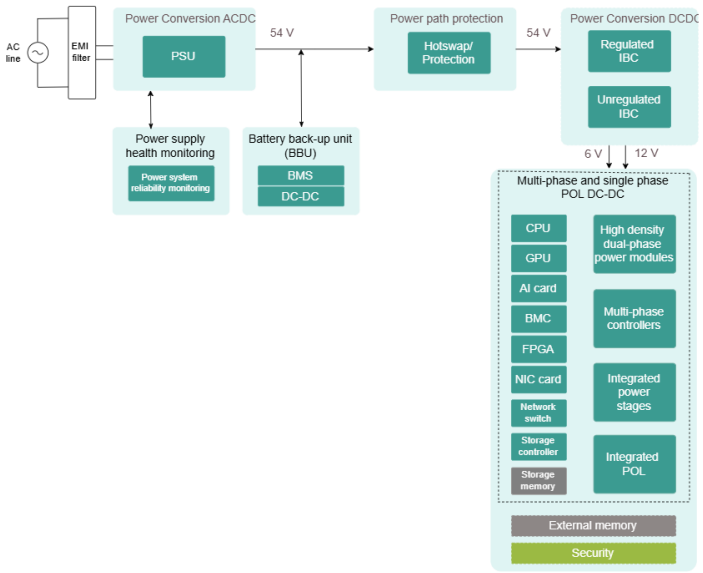
BofA GLOBAL RESEARCH

The competitive landscape is dynamic for these sockets. Infineon is competing across the full IBC stack, including power devices, protection, sensing, control, and downstream power. Renesas is focused on GaN-based high-density DC/DC conversion, including 48V-to-400V architectures with stacking to 800V. TXN is advocating for a direct jump to an 800V-to-6V isolated bus converter followed by 6V-to-sub-1V multiphase power, plus 800V hot-swap protection and 800V-to-12V options. MPS remains relevant where the intermediate rail feeds tightly integrated processor power and where control and density near the accelerator matter. Navitas is a wide-bandgap supplier targeting both 800V-to-50V and 800V-to-6V conversion, with the latter framed around reducing conversion stages and freeing board space near the GPU board.

Server Board (GPU/XPU power, CPU, VRM, multi-phase)

On the GPU/compute board, the biggest power challenge is delivering very high current into very low-voltage rails. As accelerators integrate more compute silicon, HBM, and high-speed I/O, total board power rises while GPU core voltage stays below 1V. The voltage regulator module (VRM) is the local power stage that converts an intermediate board rail into the sub-1V GPU/ASIC core rail. These VRMs sit close to the accelerator package because long power paths create loss and heat. As future processors move toward 2–4kW per GPU, high-density DC to DC stages supply <1V GPU core power where currents rise into the thousands of amps.

Exhibit 29: Multi-phase and single phase voltage regulator modules are among the most valuable sockets on the server board (bottom right box), helping to power GPUs, CPUs, and more
Block diagram of power delivery to the sever board and key components

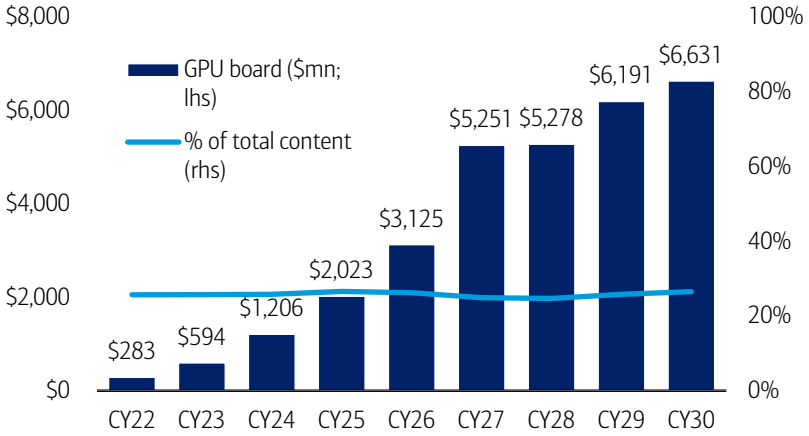


Source: Infineon

BofA GLOBAL RESEARCH

VRMs use multiphase power because the GPU core rail needs more current than a single power stage can practically deliver. Multiphase spreads the load across many smaller phases, which improves thermal spreading, current sharing, transient response, and board layout. This becomes more important as GPU load steps get faster and larger. Analog multiphase control is proven, but digital multiphase control becomes more valuable as phase counts rise because it adds telemetry, programmable tuning, current balancing, fault reporting, and easier reuse across processor platforms.

Exhibit 30: Power delivery to the processor (GPU/XPU) is among the most valuable components (mid-20% of content TAM) and grows from $2bn in CY25 to over $6.6bn by CY30
TAM for GPU related components on the server board

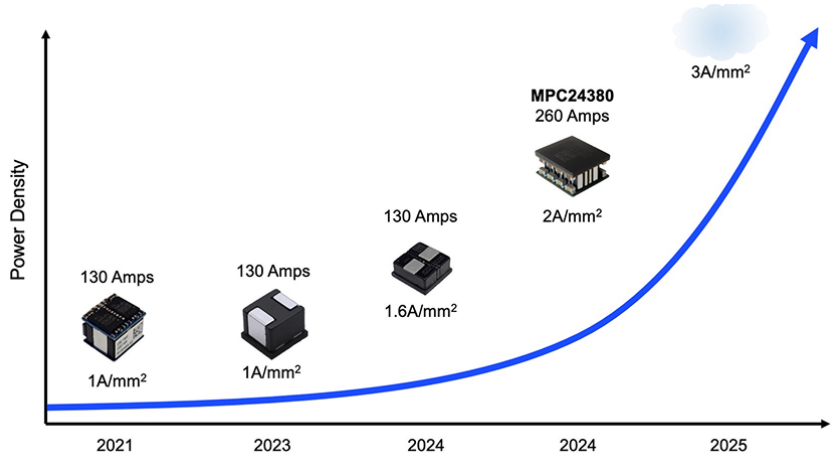


Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia

BofA GLOBAL RESEARCH

The compute tray power architecture is also changing at the processor level. In lower-power systems, the final VRM typically steps down from a familiar board rail such as 12V. In higher-power trays, more architectures are looking at lower intermediate rails closer to the processor, such as 6V, to reduce the burden on the final VRM. The tradeoff is that lower intermediate voltage means higher current, so the conversion stage must sit closer to the load. This is why the opportunity is shifting toward high-density VRMs, integrated power stages, better capacitors, tighter sensing, and advanced packaging around the processor rather than just more discrete regulators across the board.

Exhibit 31: VRMs increasingly have to channel more current through a smaller physical footprint as power density rises with the roadmap tripling amps per millimeter squared in several years via smaller packages, multiphase integration, and innovations in advanced packaging
Voltage Regulator Module roadmap



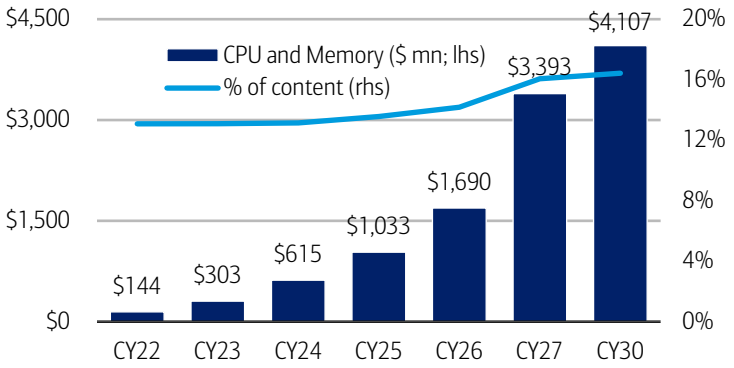
Source: Monolithic Power Systems

BofA GLOBAL RESEARCH

GPU-board content remains the largest socket, and in our model GPU-board content rises from roughly $9.7k in a 100–160kW rack to $272k in a 1MW+ rack, reflecting more GPUs, higher current per accelerator, more phases, more vertical power, and more package-adjacent regulation. However, the CPU and HBM broaden the analog opportunity beyond the GPU core VRM.

CPU content is smaller per socket, but it grows as we push deeper into the agentic AI. If NVIDIA deploys separate CPU racks, the opportunity expands from CPU attach inside the GPU tray to a standalone pool of CPU VRMs, memory power, board-level point of loads, sensing, timing, protection, NIC/DPU power, and platform-management rails. HBM also adds content as memory capacity, bandwidth, and stack complexity rise, driving more PMICs, low-noise rails, sequencing, sensing, and thermal monitoring.

Exhibit 32: The board opportunity is expanding as content tied to CPU and HBM continues to rise in conjunction with rising processor power as well as with new agentic workloads
CPU complex and attach content along with HBM/memory attach content

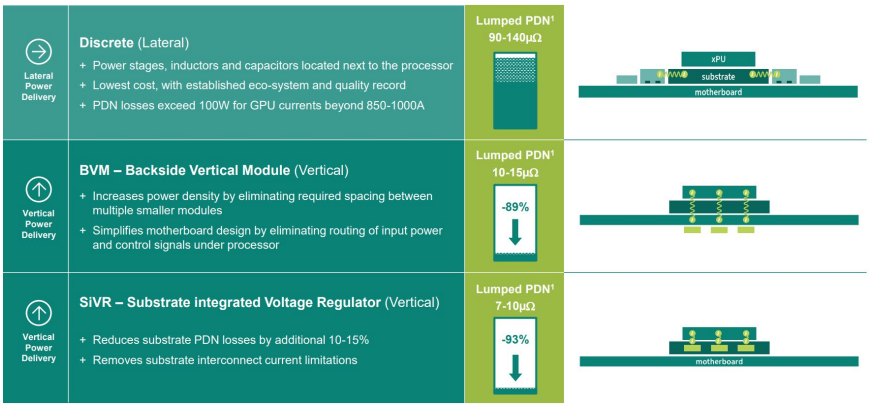


Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia

BofA GLOBAL RESEARCH

Power delivery is also shifting physically. Lateral power delivery places regulators beside the processor and routes current across the PCB and package. It is mature and lower cost, but losses rise at very high current, and it consumes valuable real estate around the package that can be used for other functions. Vertical power delivery moves power under the processor or through the backside of the board/package, shortening the current path and reducing losses. Integrated power delivery goes further by moving regulation and capacitance into or next to the package. ADI's acquisition of Empower fits directly here because ADI says Empower adds integrated voltage regulator and silicon capacitor technology, enabling conversion closer to the processor and improving efficiency for higher-density systems.

Exhibit 33: Lateral power delivery today is increasingly making way for vertical power delivery in order to reduce losses. Integrated voltage regulators extend the roadmap further LT
Power delivery options to the core



Source: Infineon

BofA GLOBAL RESEARCH

The competitive landscape is shifting from selling discrete power parts to owning more of the compute power delivery architecture. MPS is competing through density, integration, digital control, and proximity to the accelerator. Infineon is competing with breadth across processor power, vertical power, protection, sensing, and control. ADI is moving from high-performance analog power into package-adjacent power through Empower, which gives it a stronger position in the highest-value last-inch socket. TI is positioned around broad analog power management and simplified processor power paths. Renesas is competing around high-density DC/DC conversion, controllers, drivers, and 800V architecture support. This is such a valuable socket because units and ASPs both rise. Units rise as GPU count, HBM content, rails, phases, sensors, and protection points increase per tray. ASPs rise because the solution is moving from commodity discretes toward digital multiphase controllers, integrated power stages, vertical modules, IVRs, silicon capacitors, telemetry, and advanced packaging. Discrete component count may consolidate, but content dollars per GPU/compute board rise materially as power delivery becomes a limiting factor for AI performance.

Other components (protection, sensors, optics, etc.)

We have covered the major and most valuable sockets as data centers move towards high voltage architecture, but next we will cover the smaller categories of components.

Protection (i.e. hot swap controllers)

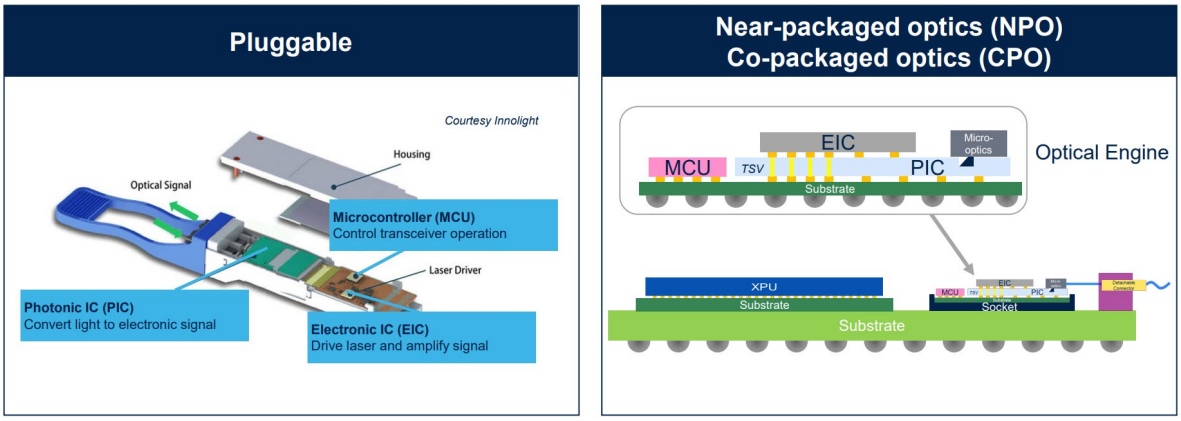
Protection is the circuitry that lets trays, boards, and power modules connect to a live DC bus without damaging the rack. In 48V systems, this is mostly eFuse, hot-swap, overcurrent, short-circuit, and telemetry around the rack bus. In 400V/800V systems, it becomes a higher-ASP socket because the controller has to manage pre-charge, inrush current, discharge, fault isolation, live serviceability, and real-time telemetry at much higher stored energy. ADI has an established hot-swap portfolio, while Infineon and TI are also pushing 400V/800V protection and hot-swap for AI racks. We think protection can grow to become a $700mn market by CY30 but note Infineon believes the AI protection SAM could reach $800mn during the same time led by 400V/800V products.

Optical infrastructure

Optical analog content is one of the largest categories in our model because AI rack bandwidth scales with GPU count, cluster size, and the move from 800G to 1.6T links. In our content build, optical infrastructure rises from roughly $4.4k in a 100–160kW rack to $44k in a 600kW+ rack and $132k in a 1MW+ rack because there are more optical ports, more lanes per port, higher-speed analog front ends, and more power/thermal management around each module. An 800G 200G/lane module can include roughly one TIA, four photodiodes, and zero to one laser drivers, depending on whether the architecture uses external lasers, linear drive, or a more integrated optical engine. Moving to 1.6T can roughly double that analog footprint, and in some architectures the analog content can more than double as control, CDR, power, TEC, monitoring, and SiPho interface content increases. Regardless of whether the industry pursues retimed transceivers, LPO/LRO, or even CPO, analog content should remain constant.

ADI sells optical control and power solutions including precision controllers, converters, and TEC controllers making up roughly half their data center revenue. Renesas sells VCSEL drivers and TIA receiver arrays. MTSI sells CDRs, TIAs, laser drivers, CW lasers, photodiodes, and analog CDR-based optical components. STMicro's optical angle is a full module level stack. Its PIC100 platform is the photonic IC (PIC) layer for silicon photonics optical interconnects, supporting 200G-per-lane capability for 800G/1.6T-class modules, while its SiGe BiCMOS technology supports the electronic IC (EIC) layer, including high-speed TIAs and laser drivers. ST also emphasizes the MCU another major component inside the transceiver, used to control module operations

Exhibit 34: PICs, EICs, and MCUs are examples of analog products that vendors like STMicro can sell into both an optical transceiver (pluggable) or a co-packaged optics (CPO) solution
Analog content in optics



Source: STMicro, Innolight

BofA GLOBAL RESEARCH

Clocking and timing

Clocks and timing ICs generate and distribute clean reference signals, while retimers regenerate PCIe/CXL links so GPUs, CPUs, NICs, switches, and accelerators can communicate across longer traces and denser rack topologies. Content rises because PCIe 5/6, CXL, 800G networking, optics, and pooled accelerator architectures need more lanes, lower jitter, longer reach, and better telemetry. MCHP sells PCIe oscillators, clock generators, and buffers through Gen7, and its 3nm Gen 6 PCIe switch adds a higher-value adjacency with 64GT/s PCIe 6.0, up to 160 lanes, low latency, diagnostics, and CPU/GPU/accelerator/storage connectivity. MCHP also is involved in retimers but Marvell and Astera Labs are stronger here. ASPs rise as the socket moves from basic clocks toward low-jitter timing, PCIe switching, retiming, security, and fleet reliability

Sensors/Signal chain

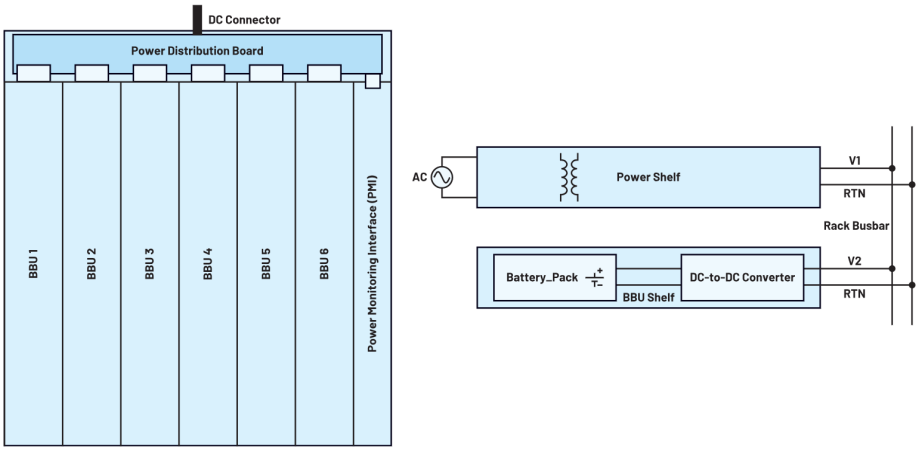
Current sensors measure current flow for regulation, protection, telemetry, and fault response. In today's 48V racks, the job is monitoring busbars, PSUs, BBUs, VRMs, and protection paths. In 800V/high-power racks, current sensing becomes more important because GaN/SiC stages switch faster, fault windows shrink, and racks need real-time visibility across IBCs, hot swap, BBU, PCS, and tray power. TMR sensors are attractive because they can provide faster, lower-noise, contactless current measurement versus traditional magnetic sensors. Key players include ALGM, Infineon, TDK, AKM, Melexis, and TXN/ADI in adjacent current-sense amplifiers and isolation. Signal-chain sensing is the broader analog layer around measurement, isolation, conversion, and monitoring. It includes voltage/temperature sensing, isolated amplifiers, ADCs, DACs, digital isolators, bias supplies, telemetry, and MCU control. In lower-voltage racks, this is a support socket around PSUs, BBUs, fans, VRMs, and thermal monitoring. In 400V/800V racks, it becomes more strategic because the system needs isolated measurement, faster fault detection, higher common-mode robustness, and tighter feedback loops. Key players include ADI, TXN, Infineon, Renesas, and Silicon Laboratories.

Battery Backup Unit (BBU)

A battery backup unit provides short-duration backup and power conditioning so the rack can ride through grid events, shut down safely, or bridge to other backup sources. In lower-voltage racks, BBUs can sit in or near the rack and support the 48V/50V bus. In higher-power racks, the BBU function increasingly shifts toward sidecar or centralized UPS-ESS/battery infrastructure, while rack-local BBU content grows more slowly if the central microgrid table captures the larger energy-storage opportunity (as we capture in our model) although content generally improves. Materials include lithium-ion batteries, silicon MOSFETs, GaN in high density partial-power conversion, gate drivers, MCUs, current sensors, and protection devices. Infineon, ADI, and TXN have the most comprehensive product portfolio to address BBUs.

Exhibit 35: The BBU shelf provides backup power if an AC power outage occurs for a defined period of time, giving time for the rack to be moved between power sources without disrupting the IT gear. OCP's architecture below shows a shelf that sends power across a common bus bar to IT gear but for Rubin Ultra we think BBU functions could relocate to the power sidecar

BBU shelf configuration in an Open Compute Project ORV3 open rack power architecture



Source: Analog Devices, Open Compute Project

BofA GLOBAL RESEARCH

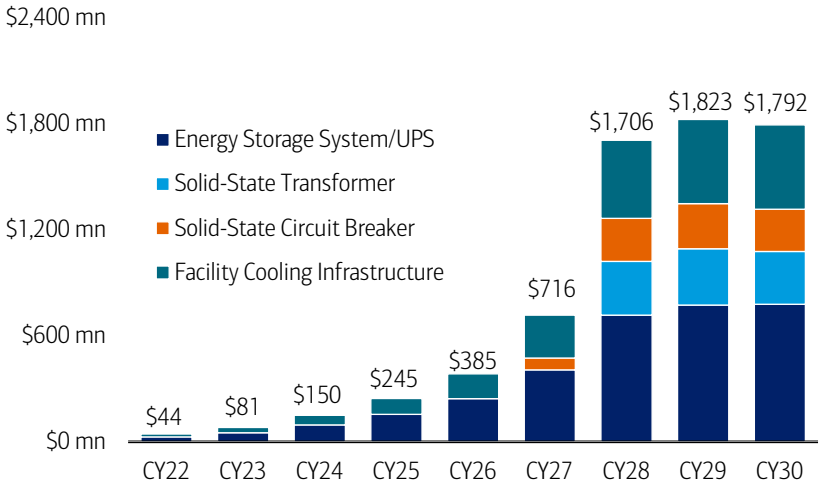
Power Conversion System (PCS)

PCS is the power conversion system that buffers fast rack or tray-level transients. It is not long-duration backup like a BBU but instead is a fast response energy buffer that helps keep the DC bus stable when GPUs ramp power quickly. In lower power racks, this function can be handled by bulk capacitance inside PSUs, trays, or VRM input networks. In 800V/1MW racks, it becomes a more visible socket because the rack needs to manage sharper load steps, higher energy on the DC bus, and more serviceable power shelves or sidecars (hence why content begins after 600kw/rack). Content depends on whether buffering sits centrally at 800V, inside a sidecar, or closer to trays. TXN and Infineon are both closely involved here. In 1MW racks, PCS content likely moves upstream as the buffering could move into a centralized HVDC bus or microgrid energy storage layer but does not disappear from the rack entirely.

Power Infra: $2bn TAM from grid to hall

We estimate the strategic power infrastructure analog semi opportunity grows from $300mn in CY25 to $1.8bn in CY30, a ~49% CAGR, with growth near +90%/140% CY26/27 as Solid-State Transformers (SST) and Solid-State Circuit Breakers (SSCB) gain traction. Our model drives content on a MW basis and we note the inflection in CY28 begins with the 1MW level racks and the introduction of the hybrid microgrid architecture which moves power content upstream. This is not an exhaustive TAM, and analog suppliers are likely shipping more to energy/power infrastructure that is seeing a "halo effect" from AI, but we focus on the more prominent categories. These sales would likely be reported in the industrial or infrastructure segments of the various analog vendors rather than in data center.

Exhibit 36: We forecast the power infra-analog semi TAM to grow from <$300mn today to $1.8bn by the end of the decade, a 49% CAGR
Strategic power infrastructure TAM ($mn)



Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia

BofA GLOBAL RESEARCH

The end state of 800V turns AI power delivery into a site-level microgrid architecture, where the rack becomes more compute-dense and DC-native while more value moves upstream into HVDC conversion, storage, protection, and cooling. In our model, content rises from roughly $12.4k to $38.9k because 1MW+ racks need an 800V backbone, centralized ESS/UPS for ride-through, SST or high-voltage conversion to create the DC bus, SSCBs and hot-swap protection to make the bus safe and serviceable, and facility liquid-cooling infrastructure to remove the heat. SiC becomes more important in high-voltage conversion and protection where blocking voltage, thermal robustness, and fault handling matter. GaN fits higher-frequency, higher-density conversion stages where smaller magnetics and efficiency matter. Silicon remains important in controllers, drivers, sensing, telemetry, lower-voltage power stages, and cooling controls.

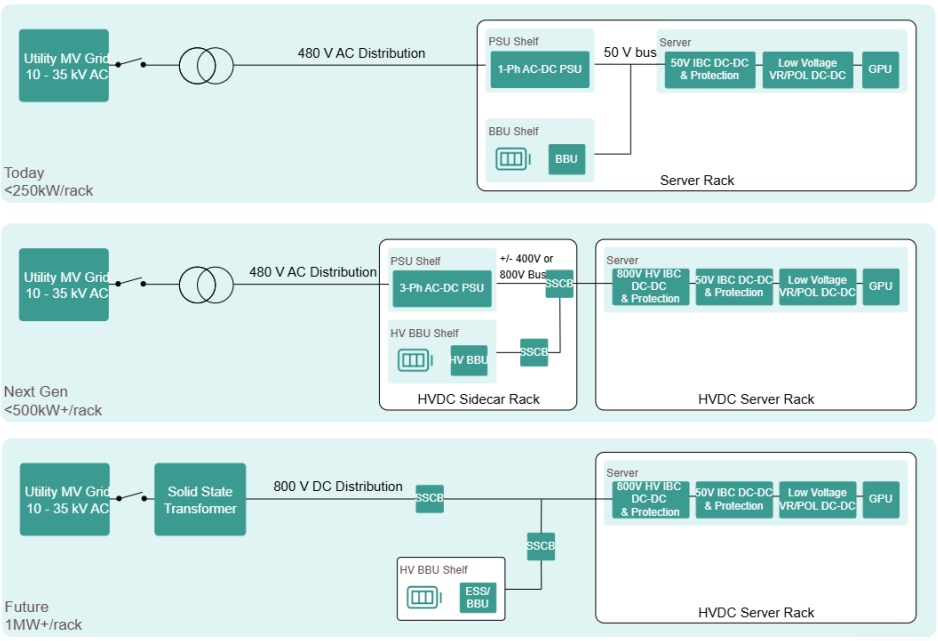
Exhibit 37: 1MW racks will introduce new infrastructure-level technologies including SSTs and SSCBs. Content per MW could rise over 3x to $38.9K by the end of the decade.
Strategic power infra-analog semi content evolution from low voltage to high voltage

	Per MW (Low Voltage Era)		Per MW (High Voltage Era)	
Components	Sales	%	Sales	%
Energy Storage System/UPS	$7,850	64%	$14,130	36%
Solid-State Transformer	$0	0%	$8,772	23%
Solid-State Circuit Breaker	$0	0%	$7,018	18%
Facility Cooling Infrastructure	$4,500	36%	$9,000	23%
Total	$12,350	100%	$38,920	100%

Source: BofA Global Research estimates, company reports, Gartner, Omdia, Infineon, Nvidia

BofA GLOBAL RESEARCH

Exhibit 38: Previous functions like the HVDC sidecar rack for 500-600kW+ racks could be disaggregated and content can move up the chain as new technologies like solid-state transformers, enhanced ESS, and solid-state circuit breakers bear more of the load for 1MW+ class racks
Grid-to-core block diagram evolution



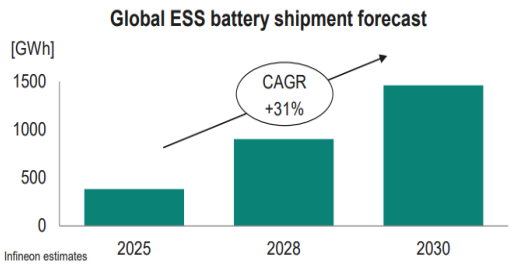
Source: Infineon

BofA GLOBAL RESEARCH

Energy Storage System (ESS)/Uninterruptible Power Supply (UPS)

Historically, ESS was mainly a residential, commercial, and utility-scale power asset. In residential and commercial settings, batteries stored excess solar or low-cost grid power and provided backup during outages. At utility scale, ESS helped buffer renewable intermittency, shift power from low-demand to high-demand periods, and stabilize the grid by smoothing supply/demand imbalances. The semiconductor opportunity was formerly concentrated around the battery-management system, inverter/PCS, sensing, isolation, MCUs, and protection needed to safely monitor battery health and move energy in and out of the pack. For AI data centers, the role expands because the load is not just larger, it is also more concentrated and more dynamic. Rather than just acting as a standby reserve, energy storage becomes a core part of the power delivery architecture by helping provide ride-through, smooth GPU-driven load swings, support the HV sidecars, and ultimately bridge to a hybrid microgrid style facility design.

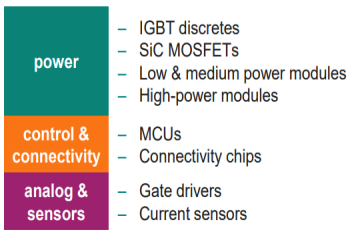
Exhibit 39: ESS battery shipments grow +31% CAGR CY25 to CY30
Global ESS battery shipment forecasts 2025/2026/2028 (Infineon)



Source: Infineon

BofA GLOBAL RESEARCH

Exhibit 40: Most of the ESS content is concentrated in power, but control (e.g. MCUs), and analog (gate drivers sensors) are valuable too
Devices used in ESS



Source: Infineon

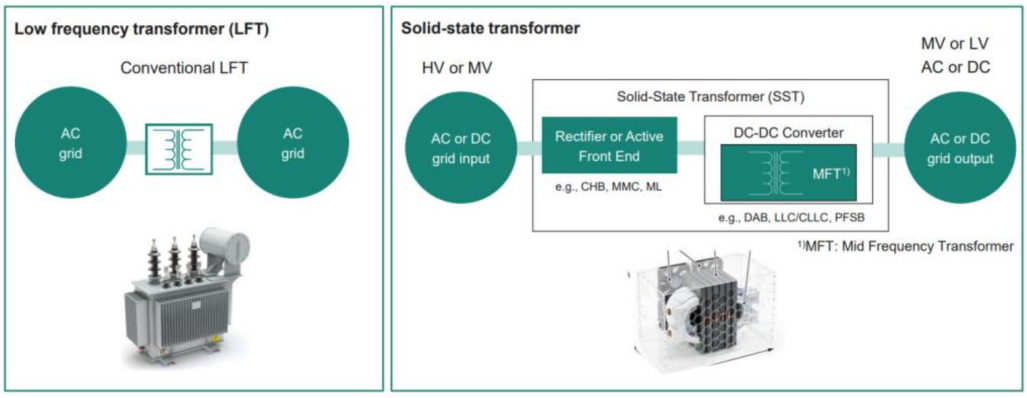
BofA GLOBAL RESEARCH

We assume ESS/UPS content per MW rises from $7.9k in the low-voltage era to $14.1k in the high-voltage era. The uplift comes from the electronics around the storage system as it must interface with higher voltage buses and higher power transient loads, which adds more battery monitoring, current and voltage sensing, isolation, gate drivers, MCUs, and protection. SiC should benefit most as ESS connects to higher-voltage PCS/UPS and protection stages. GaN should be more selective, showing up where high-frequency DC/DC conversion improves density rather than as the main ESS content. By company, Infineon looks best positioned where ESS becomes part of the broader grid-to-core HVDC architecture because it participates across Si, SiC, GaN, MCUs, sensors, drivers, and protection. TI should benefit from the breadth of its power path portfolio across ESS/UPS, sensing, conversion, eFuse, hot-swap, and MCU content. Renesas is relevant where backup power shifts toward higher-voltage BBU and distributed storage designs, specifically arguing that wide-bandgap devices can help HV BBU systems address weight, volume, and efficiency constraints. We estimate the ESS opportunity for AI data centers specifically today to be relatively small at $156mn but believe it can reach nearly $800mn by CY30, a +38% CAGR.

Solid-State Transformer (SST)

Solid-state transformers are the next infrastructure layer above ESS in the data center high voltage power stack. Historically, the transformer was mostly a passive grid asset that stepped medium voltage utility power down to lower voltage AC, after which the data center still needed layers of UPS, switchgear, PDUs, PSUs, and AC to DC conversion before power reached the rack. This was acceptable when rack power was lower and loads were less dynamic, but AI changes the problem because power availability, deployment speed, and conversion efficiency become larger constraints. Indeed, there are many reports of potential data center projects at risk of delay due to grid constraints and supply shortages in conventional transformers (medium voltage transformer lead times up to 3 years per Wolfspeed). This is where SSTs come in as they replace a bulky, long lead time, low-frequency transformer with a modular power-electronics system that can convert medium-voltage AC, typically 13.8kV to 35kV, directly to an 800 VDC bus for the compute rack. In doing so, this helps (1) compress the power chain by reducing downstream AC to DC stages, and (2) creates a more modular grid interconnect that can be easier to deploy. Compared with a conventional transformer, SSTs could be ~40x lighter, ~14x smaller, and have a ~50% faster construction time, per Infineon.

Exhibit 41: When compared to a conventional transformer, SSTs are superior as it replaces much of the passive transformer function typically done using copper wirings and magnetic cores to something that is more active, with high frequency switching, isolation, digital control, and protection. In a hybrid microgrid, SSTs will be a core block as part of the 800 VDC backbone.
Solid-state transformer vs. a low frequency transformer (LFT)



Source: Infineon

BofA GLOBAL RESEARCH

We assume SST content moves from essentially $0/MW in the low voltage era (virtually no deployments today) to $8.8k/MW in the high voltage era. A conventional transformer has next to zero analog semi content, but an SST is built out of high voltage power devices, gate drivers, isolation, current and voltage sensing, controllers, protection, magnetics, and digital control. SiC should be the primary material winner because the main SST challenge is high voltage, high efficiency switching at the grid interface. We expect pilot tests to begin soon if not already ongoing, with emerging wide-bandgap vendors like Wolfspeed and Navitas demonstrating solutions, but with larger suppliers like Infineon also well positioned given their portfolio breadth (not just discretes needed, also analog and MCUs). We think the real inflection should begin CY28-30 when hybrid microgrid architectures begin to see adoption. Infineon believes the entire SST market (not just semi content) could be $1bn by CY30 and we think the analog semi opportunity could reach $500mn by that same point in time.

Exhibit 42: A key debate of SSTs is whether to implement medium voltage-based designs or high voltage-based designs. The high voltage market is less established but could be ideal in the longer term due to its design simplicity and higher performance.
Tradeoffs between a medium voltage and high voltage silicon carbide SST design

Design Nuances	MV-rated SiC power devices <5kW blocking voltage	HV-rated SiC power devices >5kW blocking voltage
Architecture	Series stacking of multiple devices, often in three to five-level architectures	Wolfspeed 10kV die can enable a simple two-level cell
Modularity	Smaller building blocks are easier to handle and service	Larger and less modular therefore not as easy to handle
Insulation	Requires precise attention to insulation coordination between cells	Simpler insulation
Control Scheme	Control algo needed to gate series switches for cleaner output	Less complex control scheme due to fewer cells
Gate Driver	3.3kW rated SiC gate drivers commercially available	Very few comparable solutions available above 3.3kV requiring vendors to do custom development and qualification
Electromagnetic Interference (EMI)	Multi-level architecture requires smaller dv/dt	Switch with high dv/dt requiring ultra-low-parasitic design

Source: Wolfspeed, BofA Global Research; dv/dt = change in voltage over time

BofA GLOBAL RESEARCH

Solid-State Circuit Breakers (SSCB)

Solid-state circuit breakers are the protection layer that becomes more important as data center power shifts from lower-voltage AC/48V architectures toward HVDC distribution. Historically, protection was handled by electromechanical breakers, relays, fuses, and lower voltage eFuses, which were adequate when power levels were lower and AC systems naturally crossed zero current. The issue in high-voltage AI data centers is that DC distribution and higher rack power create faster, higher-energy fault events where mechanical breakers are too slow and arcing becomes harder to manage. Doing this with a SSCB, however, grants the ability to interrupt current on a nano/microsecond scale rather than the millisecond scale of mechanical breakers, a critical capability. Also a core part of the hybrid microgrid architecture, SSCBs enable ultra-fast overload and short-circuit interruption, integrated monitoring, and remote control. Similar to an ESS or SST, circuit protection moves from a relatively passive safety component to an active, monitored power-management node that helps make 400/800V distribution serviceable, fault-tolerant, and safe enough for 600kW to 1MW-class AI racks.

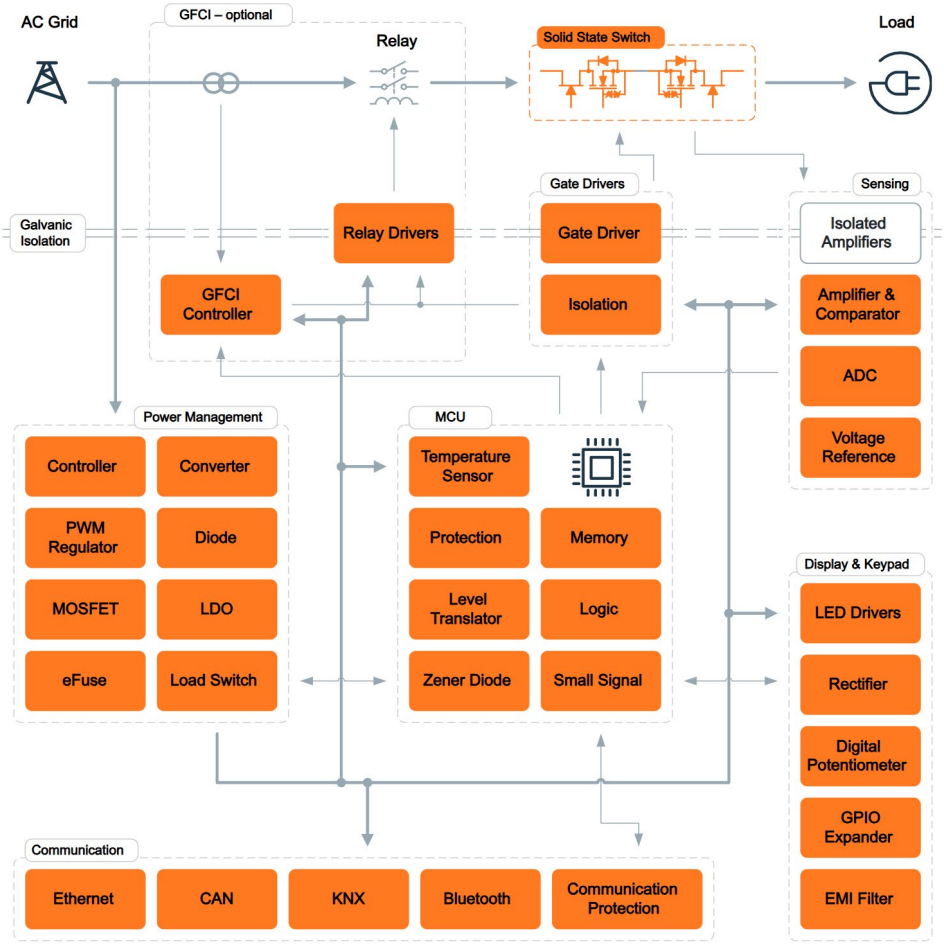
Exhibit 43: SSCBs are superior to electromechanical breakers in controllability, speed, arcing mitigation, and reliability/maintenance (no mechanical parts)
Feature comparison between solid-state circuit breakers and electromechanical breakers

Feature	Solid-state breaker (Si, SJ, SiC, IGBT, IGCT)	Electromechanical breaker
Full controllability	●●●●●	●●●
High speed	●●●●●	●●
Conduction loss	●●	●●●●●
No arcing	●●●●●	●●
Use cycles: no maintenance	●●●●●	●●
Cost per amp	●●	●●●●●
Voltage rating vs. on RDS(on)	●●●	●●●●●

Source: Onsemi, BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 44: SSCB content is higher vs. traditional circuit breakers due to the inclusion of high voltage power switches, gate drivers, sensing, MCUs and more vs. the traditional mechanical design
SSCB block diagram



Source: Onsemi

BofA GLOBAL RESEARCH

We assume SSCB content moves from essentially $0/MW in the low voltage era to $7k/MW in the high voltage era. The content uplift comes because the breaker itself becomes a semiconductor system with high voltage power switches, gate drivers, current/voltage sensing, isolation, MCUs/control, thermal monitoring, and communications replacing a mostly mechanical protection element. Silicon MOSFETs are cost-effective below 600V, but once voltage exceeds 600V, silicon resistance becomes too high. In the 600V–2000V range, SiC FETs offer the lowest resistance per unit area, while SiC devices can tolerate high instantaneous temperature rise and handle up to 4x the energy per unit area of silicon. This makes SiC the best device type in HVDC SSCBs, especially 750V/1200V/1700V-class devices. Both Infineon and ON look well positioned for SSCBs given their SiC JFET (junction field effect transistor) to MOSFET portfolio which unlocks the energy savings and supports the reliability and scalability requirements needed. We think SSCBs could see adoption into CY27, but the opportunity mainly inflects when 1MW racks become more commonplace leading to a ~$400mn opportunity for analog semi vendors by CY30 when compared to an overall $800mn SSCB module market (total not just semis) per Infineon also by decade-end.

AI Power Analog Semi Model

Exhibit 45: Inside a $7.9bn TAM today growing to $26.8bn, TXN should enjoy the largest revenue position but we note ADI and Infineon also have large scale due to portfolio breadth
Total company revenue summary across data center and power infra

Company Revenue Summary ($; mn)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total	**$1,147**	**$2,395**	**$4,829**	**$7,864**	**$12,302**	**$21,837**	**$23,094**	**$25,847**	**$26,793**
TXN	$236	$580	$959	$1,547	$2,416	$4,368	$4,549	$5,277	$5,658
ADI	$138	$300	$654	$1,045	$1,794	$3,173	$3,424	$4,106	$4,439
Infineon	$85	$155	$307	$904	$1,647	$3,415	$3,949	$4,490	$4,630
MPWR	$231	$404	$726	$710	$1,212	$2,110	$1,915	$2,179	$2,176
ON	$39	$80	$178	$295	$566	$1,577	$1,979	$2,219	$2,298
STMicro	$60	$124	$233	$379	$597	$1,090	$1,301	$1,452	$1,504
Renesas	$72	$149	$445	$767	$1,181	$1,986	$1,931	$2,194	$2,306
MCHP	$43	$97	$155	$249	$393	$716	$764	$865	$901
Other	$242	$504	$1,172	$1,969	$2,496	$3,402	$3,282	$3,065	$2,880
Low Power/Voltage	**$1,147**	**$2,394**	**$4,829**	**$7,864**	**$10,862**	**$10,790**	**$8,530**	**$9,527**	**$10,739**
TXN	$236	$580	$959	$1,547	$2,172	$2,226	$1,799	$2,055	$2,369
ADI	$138	$300	$654	$1,045	$1,617	$1,718	$1,449	$1,713	$1,983
Infineon	$85	$155	$307	$904	$1,413	$1,498	$1,246	$1,456	$1,665
MPWR	$231	$404	$726	$710	$1,090	$1,131	$805	$899	$959
ON	$39	$80	$178	$295	$475	$568	$492	$593	$716
STMicro	$60	$124	$233	$379	$525	$527	$425	$478	$543
Renesas	$72	$149	$445	$767	$1,058	$1,048	$784	$875	$986
MCHP	$43	$97	$155	$249	$344	$342	$271	$302	$341
Other	$242	$504	$1,172	$1,969	$2,169	$1,733	$1,259	$1,156	$1,176
High Power/Voltage	**$0**	**$0**	**$0**	**$0**	**$1,439**	**$11,047**	**$14,564**	**$16,320**	**$16,054**
TXN	$0	$0	$0	$0	$244	$2,142	$2,750	$3,222	$3,289
ADI	$0	$0	$0	$0	$177	$1,455	$1,976	$2,393	$2,455
Infineon	$0	$0	$0	$0	$234	$1,917	$2,702	$3,035	$2,964
MPWR	$0	$0	$0	$0	$122	$979	$1,109	$1,280	$1,217
ON	$0	$0	$0	$0	$91	$1,009	$1,486	$1,626	$1,583
STMicro	$0	$0	$0	$0	$72	$563	$876	$974	$961
Renesas	$0	$0	$0	$0	$123	$939	$1,147	$1,319	$1,321
MCHP	$0	$0	$0	$0	$49	$374	$493	$563	$561
Other	$0	$0	$0	$0	$326	$1,669	$2,023	$1,909	$1,703
Company Revenue Summary YoY Growth	**CY22**	**CY23**	**CY24**	**CY25**	**CY26**	**CY27**	**CY28**	**CY29**	**CY30**
Total	**-**	**108.8%**	**101.7%**	**62.9%**	**56.4%**	**77.5%**	**5.8%**	**11.9%**	**3.7%**
TXN	-	145.8%	65.2%	61.3%	56.2%	80.8%	4.2%	16.0%	7.2%
ADI	-	117.4%	118.0%	59.7%	71.7%	76.9%	8.2%	19.9%	8.1%
Infineon	-	82.4%	97.7%	194.5%	82.3%	107.3%	15.7%	13.7%	3.1%
MPWR	-	75.2%	79.5%	-2.2%	70.7%	74.0%	-9.2%	13.8%	-0.1%
ON	-	103.3%	123.7%	65.7%	91.7%	178.7%	25.6%	12.1%	3.6%
STMicro	-	106.0%	87.5%	62.2%	57.7%	82.6%	19.6%	11.6%	3.6%
Renesas	-	107.0%	198.5%	72.5%	54.0%	68.2%	-2.5%	13.6%	5.1%
MCHP	-	124.6%	59.5%	60.5%	57.6%	82.3%	8.2%	13.1%	4.1%
Other	-	108.4%	132.5%	68.0%	26.8%	36.3%	-3.2%	-6.6%	-6.1%
Low Power/Voltage	**-**	**108.7%**	**101.7%**	**62.9%**	**38.1%**	**-0.7%**	**-20.9%**	**11.7%**	**12.7%**
TXN	-	145.8%	65.2%	61.3%	40.4%	2.5%	-19.2%	14.3%	15.3%
ADI	-	117.4%	118.0%	59.7%	54.7%	6.2%	-15.7%	18.2%	15.8%
Infineon	-	82.4%	97.7%	194.5%	56.3%	6.0%	-16.8%	16.8%	14.4%
MPWR	-	75.2%	79.5%	-2.2%	53.5%	3.7%	-28.8%	11.6%	6.7%
ON	-	103.3%	123.7%	65.7%	60.7%	19.8%	-13.4%	20.4%	20.8%
STMicro	-	106.0%	87.5%	62.2%	38.6%	0.5%	-19.4%	12.6%	13.5%
Renesas	-	107.0%	198.5%	72.5%	38.0%	-1.0%	-25.2%	11.6%	12.7%
MCHP	-	124.6%	59.5%	60.5%	38.1%	-0.6%	-20.8%	11.7%	12.7%
Other	-	108.4%	132.5%	68.0%	10.2%	-20.1%	-27.3%	-8.2%	1.7%
High Power/Voltage	**-**	**-**	**-**	**-**	**-**	**667.6%**	**31.8%**	**12.1%**	**-1.6%**
TXN	-	-	-	-	-	778.3%	28.5%	17.1%	2.1%
ADI	-	-	-	-	-	722.0%	36.4%	21.0%	2.6%
Infineon	-	-	-	-	-	717.8%	41.0%	12.3%	-2.3%
MPWR	-	-	-	-	-	700.1%	13.3%	15.4%	-4.9%
ON	-	-	-	-	-	1003.8%	47.6%	9.4%	-2.7%
STMicro	-	-	-	-	-	678.8%	56.2%	11.2%	-1.3%
Renesas	-	-	-	-	-	664.2%	22.8%	14.9%	0.1%
MCHP	-	-	-	-	-	668.5%	34.9%	13.9%	-0.5%
Other	-	-	-	-	-	411.5%	21.8%	-5.6%	-10.8%

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 46: Infineon should gain the most share, up to 17% in CY30 (from 12% in CY25), but ON also sees impressive gains from 4% in CY25 to 9% CY30 along with ADI from 13% CY25 to 17% CY30. Discrete vendors generally gaining more share due to introduction of novel materials like SiC and GaN unlocking new markets to participate in

Market share across data center and power infra

Company Share (%)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**
TXN	20.6%	24.2%	19.9%	19.7%	19.6%	20.0%	19.7%	20.4%	21.1%
ADI	12.0%	12.5%	13.5%	13.3%	14.6%	14.5%	14.8%	15.9%	16.6%
Infineon	7.4%	6.5%	6.4%	11.5%	13.4%	15.6%	17.1%	17.3%	17.3%
MPWR	20.1%	16.9%	15.0%	9.0%	9.9%	9.7%	8.3%	8.4%	8.1%
ON	3.4%	3.3%	3.7%	3.8%	4.6%	7.2%	8.6%	8.6%	8.6%
STMicro	5.3%	5.2%	4.8%	4.8%	4.9%	5.0%	5.6%	5.6%	5.6%
Renesas	6.3%	6.2%	9.2%	9.8%	9.6%	9.1%	8.4%	8.5%	8.6%
MCHP	3.8%	4.1%	3.2%	3.2%	3.2%	3.3%	3.3%	3.4%	3.4%
Other	21.1%	21.1%	24.3%	25.0%	20.3%	15.6%	14.2%	11.9%	10.8%
Low Power/Voltage	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**
TXN	20.6%	24.2%	19.9%	19.7%	20.0%	20.6%	21.1%	21.6%	22.1%
ADI	12.0%	12.5%	13.5%	13.3%	14.9%	15.9%	17.0%	18.0%	18.5%
Infineon	7.4%	6.5%	6.4%	11.5%	13.0%	13.9%	14.6%	15.3%	15.5%
MPWR	20.1%	16.9%	15.0%	9.0%	10.0%	10.5%	9.4%	9.4%	8.9%
ON	3.4%	3.3%	3.7%	3.8%	4.4%	5.3%	5.8%	6.2%	6.7%
STMicro	5.3%	5.2%	4.8%	4.8%	4.8%	4.9%	5.0%	5.0%	5.1%
Renesas	6.3%	6.2%	9.2%	9.8%	9.7%	9.7%	9.2%	9.2%	9.2%
MCHP	3.8%	4.1%	3.2%	3.2%	3.2%	3.2%	3.2%	3.2%	3.2%
Other	21.1%	21.1%	24.3%	25.0%	20.0%	16.1%	14.8%	12.1%	11.0%
High Power/Voltage	**0.0%**	**0.0%**	**0.0%**	**0.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**
TXN	0.0%	0.0%	0.0%	0.0%	16.9%	19.4%	18.8%	19.7%	20.4%
ADI	0.0%	0.0%	0.0%	0.0%	12.3%	13.2%	13.6%	14.7%	15.3%
Infineon	0.0%	0.0%	0.0%	0.0%	16.3%	17.4%	18.5%	18.5%	18.4%
MPWR	0.0%	0.0%	0.0%	0.0%	8.5%	8.9%	7.6%	7.8%	7.6%
ON	0.0%	0.0%	0.0%	0.0%	6.4%	9.1%	10.2%	9.9%	9.8%
STMicro	0.0%	0.0%	0.0%	0.0%	5.0%	5.1%	6.0%	6.0%	6.0%
Renesas	0.0%	0.0%	0.0%	0.0%	8.5%	8.5%	7.9%	8.1%	8.2%
MCHP	0.0%	0.0%	0.0%	0.0%	3.4%	3.4%	3.4%	3.5%	3.5%
Other	0.0%	0.0%	0.0%	0.0%	22.7%	15.1%	13.9%	11.7%	10.6%

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 47: Analog boasts the biggest content pool, but SiC and GaN transform from niche device applications in the data center to major content winners by the end of the decade
Device TAM ($ mn) across both data center and power infrastructure

Device TAM	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Overall	**$1,147**	**$2,393**	**$4,829**	**$7,864**	**$12,301**	**$21,838**	**$23,095**	**$25,850**	**$26,795**
Silicon	$262	$546	$1,102	$1,791	$2,692	$4,128	$4,158	$4,588	$4,780
SiC	$31	$59	$114	$183	$412	$1,493	$2,175	$2,236	$2,098
GaN	$18	$37	$75	$118	$303	$1,233	$1,443	$1,618	$1,612
Analog	$751	$1,574	$3,184	$5,198	$7,974	$13,207	$13,267	$15,097	$15,919
MCU	$50	$102	$205	$332	$541	$1,092	$1,290	$1,466	$1,516
Sensors	$36	$75	$151	$242	$380	$685	$806	$891	$912
Low Power	**$1,147**	**$2,393**	**$4,829**	**$7,864**	**$10,862**	**$10,789**	**$8,529**	**$9,526**	**$10,738**
Silicon	$262	$546	$1,102	$1,791	$2,474	$2,460	$1,946	$2,173	$2,450
SiC	$31	$59	$114	$183	$261	$288	$237	$268	$305
GaN	$18	$37	$75	$118	$162	$161	$128	$143	$161
Analog	$751	$1,574	$3,184	$5,198	$7,169	$7,079	$5,581	$6,228	$7,016
MCU	$50	$102	$205	$332	$461	$466	$371	$416	$470
Sensors	$36	$75	$151	$242	$335	$336	$267	$299	$337
High Power	**$0**	**$0**	**$0**	**$0**	**$1,439**	**$11,048**	**$14,566**	**$16,323**	**$16,057**
Silicon	$0	$0	$0	$0	$218	$1,669	$2,213	$2,415	$2,330
SiC	$0	$0	$0	$0	$150	$1,205	$1,938	$1,967	$1,793
GaN	$0	$0	$0	$0	$141	$1,072	$1,315	$1,474	$1,451
Analog	$0	$0	$0	$0	$805	$6,127	$7,686	$8,870	$8,903
MCU	$0	$0	$0	$0	$80	$625	$919	$1,050	$1,047
Sensors	$0	$0	$0	$0	$45	$350	$539	$593	$575

Device TAM YoY Growth	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Overall	**-**	**108.7%**	**101.7%**	**62.9%**	**56.4%**	**77.5%**	**5.8%**	**11.9%**	**3.7%**
Silicon	-	108.7%	101.7%	62.6%	50.3%	53.3%	0.7%	10.3%	4.2%
SiC	-	93.6%	92.7%	60.8%	125.1%	262.6%	45.7%	2.8%	-6.1%
GaN	-	109.7%	101.6%	57.9%	157.2%	306.9%	17.1%	12.1%	-0.3%
Analog	-	109.6%	102.2%	63.3%	53.4%	65.6%	0.5%	13.8%	5.4%
MCU	-	106.0%	100.2%	62.2%	63.1%	101.7%	18.2%	13.6%	3.5%
Sensors	-	107.9%	101.0%	60.8%	56.6%	80.5%	17.6%	10.6%	2.3%
Low Power	**-**	**108.7%**	**101.7%**	**62.9%**	**38.1%**	**-0.7%**	**-20.9%**	**11.7%**	**12.7%**
Silicon	-	108.7%	101.7%	62.6%	38.1%	-0.6%	-20.9%	11.7%	12.7%
SiC	-	93.6%	92.7%	60.8%	42.9%	10.3%	-17.9%	13.4%	13.8%
GaN	-	109.7%	101.6%	57.9%	37.5%	-0.6%	-20.3%	11.7%	12.6%
Analog	-	109.6%	102.2%	63.3%	37.9%	-1.3%	-21.2%	11.6%	12.7%
MCU	-	106.0%	100.2%	62.2%	38.9%	1.2%	-20.4%	12.0%	12.9%
Sensors	-	107.9%	101.0%	60.8%	38.2%	0.2%	-20.5%	11.8%	12.8%
High Power	**-**	**-**	**-**	**-**	**-**	**667.6%**	**31.8%**	**12.1%**	**-1.6%**
Silicon	-	-	-	-	-	664.6%	32.6%	9.1%	-3.5%
SiC	-	-	-	-	-	701.0%	60.9%	1.5%	-8.9%
GaN	-	-	-	-	-	660.1%	22.7%	12.1%	-1.6%
Analog	-	-	-	-	-	661.5%	25.4%	15.4%	0.4%
MCU	-	-	-	-	-	677.8%	46.9%	14.3%	-0.3%
Sensors	-	-	-	-	-	684.1%	54.2%	9.9%	-3.0%

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 48: SiC and GaN share begin to inflect when high voltage racks proliferate beginning CY27 and expanding into CY30

Device TAM share aggregated for both data center and power infra

Device TAM Share (%)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Overall	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**
Silicon	22.8%	22.8%	22.8%	22.8%	21.9%	18.9%	18.0%	17.7%	17.8%
SiC	2.7%	2.5%	2.4%	2.3%	3.3%	6.8%	9.4%	8.6%	7.8%
GaN	1.5%	1.5%	1.5%	1.5%	2.5%	5.6%	6.2%	6.2%	6.0%
Analog	65.5%	65.7%	65.9%	66.1%	64.8%	60.5%	57.3%	58.3%	59.3%
MCU	4.3%	4.3%	4.2%	4.2%	4.4%	5.0%	5.6%	5.7%	5.6%
Sensors	3.1%	3.1%	3.1%	3.1%	3.1%	3.1%	3.5%	3.4%	3.4%
Low Power	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**88.3%**	**49.4%**	**36.9%**	**36.9%**	**40.1%**
Silicon	22.8%	22.8%	22.8%	22.8%	20.1%	11.3%	8.4%	8.4%	9.1%
SiC	2.7%	2.5%	2.4%	2.3%	2.1%	1.3%	1.0%	1.0%	1.1%
GaN	1.5%	1.5%	1.5%	1.5%	1.3%	0.7%	0.6%	0.6%	0.6%
Analog	65.5%	65.7%	65.9%	66.1%	58.3%	32.4%	24.1%	24.0%	26.1%
MCU	4.3%	4.3%	4.2%	4.2%	3.7%	2.1%	1.6%	1.6%	1.8%
Sensors	3.1%	3.1%	3.1%	3.1%	2.7%	1.5%	1.2%	1.2%	1.3%
High Power	**0.0%**	**0.0%**	**0.0%**	**0.0%**	**11.7%**	**50.6%**	**63.1%**	**63.2%**	**59.9%**
Silicon	0.0%	0.0%	0.0%	0.0%	1.8%	7.6%	9.6%	9.3%	8.7%
SiC	0.0%	0.0%	0.0%	0.0%	1.2%	5.5%	8.4%	7.6%	6.7%
GaN	0.0%	0.0%	0.0%	0.0%	1.1%	4.9%	5.7%	5.7%	5.4%
Analog	0.0%	0.0%	0.0%	0.0%	6.5%	28.1%	33.2%	34.3%	33.2%
MCU	0.0%	0.0%	0.0%	0.0%	0.7%	2.9%	4.0%	4.1%	3.9%
Sensors	0.0%	0.0%	0.0%	0.0%	0.4%	1.6%	2.3%	2.3%	2.1%

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 49: We see a cumulative 233 GWs to be deployed CY25-30 for AI across both Nvidia platforms and other merchant and custom ASIC alternatives

Aggregate GW deployed by accelerator platform CY22-30

GW Deployed by Platform	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30	CY25-30 Cumulative
Aggregate	**3**	**5**	**10**	**17**	**25**	**37**	**53**	**58**	**60**	**233**
NVDA Total	**0**	**2**	**4**	**8**	**12**	**17**	**30**	**30**	**28**	**118**
Rosa Feynman	-	-	-	-	-	-	1	4	7	12
Feynman	-	-	-	-	-	-	1	4	7	12
Rubin Ultra	-	-	-	-	-	2	10	18	15	45
Vera Rubin/Ultra	-	-	-	-	1	6	18	4	-	28
Rubin	-	-	-	-	1	8	1	-	-	10
GB200/GB300 (per GPU package)	-	-	0	5	7	1	-	-	-	8
B100/B200/B300 (HGX)	-	-	0	2	3	0	-	-	-	3
H100/H200	0	1	3	1	-	-	-	-	-	-
A100	0	1	0	0	-	-	-	-	-	-
non-NVDA Total	**2**	**4**	**6**	**9**	**13**	**20**	**23**	**28**	**32**	**115**
AMD	**0**	**0**	**1**	**1**	**1**	**2**	**2**	**3**	**3**	**12**
Intel	**0**	**0**	**0**	**-**	**0**	**0**	**0**	**0**	**0**	**0**
Google TPU	**2**	**2**	**3**	**2**	**4**	**8**	**10**	**11**	**12**	**45**
Broadcom 1	2	2	3	2	3	5	6	7	7	28
MediaTek	-	-	-	-	0	3	4	5	5	17
AWS Trainium	**0**	**1**	**1**	**2**	**3**	**5**	**6**	**8**	**10**	**32**
Marvell 1 (Trn)	-	0	0	1	1	1	1	1	1	4
Alchip (Trn)	0	1	1	1	3	4	5	7	9	28
Microsoft (Maia)	**-**	**-**	**-**	**-**	**-**	**0**	**0**	**0**	**0**	**1**
Marvell 2	-	-	-	-	-	0	0	0	0	1
Meta (MTIA)	**0**	**0**	**0**	**0**	**0**	**0**	**1**	**1**	**2**	**5**
Broadcom 2	0	0	0	0	0	0	1	1	2	5
OpenAI	**-**	**-**	**-**	**-**	**0**	**1**	**1**	**2**	**2**	**7**
Broadcom 3	-	-	-	-	0	1	1	2	2	7
Others	**0**	**0**	**1**	**3**	**4**	**4**	**2**	**2**	**2**	**14**

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 50: High power racks commence shipment with Rubin Ultra and Feynman

Accelerator and rack build

NVDA Accelerators	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
NVDA Units (k, per GPU package)	502	1,646	4,023	4,461	6,209	6,880	7,085	8,058	8,351
Rosa Feynman	-	-	-	-	-	-	255	1,612	2,505
Feynman	-	-	-	-	-	-	255	1,612	2,505
Rubin Ultra	-	-	-	-	-	445	2,135	4,029	3,341
Vera Rubin/Ultra	-	-	-	-	350	2,438	4,000	806	-
Rubin	-	-	-	-	450	3,232	440	-	
GB200/GB300 (per GPU package)	-	-	185	2,565	3,801	630	-	-	-
B100/B200/B300 (HGX)	-	-	145	1,239	1,608	135	-	-	-
H100/H200	12	1,110	3,495	655	-	-	-	-	-
A100	490	536	198	2	-	-	-	-	-
NVDA Racks									
NVDA rack equivalent (k)	**16**	**51**	**120**	**73**	**86**	**92**	**50**	**39**	**32**
Vera Feynman	0	0	0	0	0	0	0	3	4
Feynman	0	0	0	0	0	0	0	3	4
Rubin Ultra	0	0	0	0	0	3	15	28	23
Vera Rubin/Ultra	0	0	0	0	5	34	28	6	0
Rubin	0	0	0	0	6	45	6	0	0
GB200/GB300 (per GPU package)	0	0	3	36	53	9	0	0	0
B100/B200/B300 (HGX)	0	0	2	17	22	2	0	0	0
H100/H200	0	35	109	20	0	0	0	0	0
A100	15	17	6	0	0	0	0	0	0

Non-NVDA Accelerators	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Non-NVDA Units (AMD, INTC, TPU,etc.)	**2,331**	**3,655**	**5,944**	**11,404**	**16,301**	**19,685**	**20,066**	**22,996**	**25,724**
AMD	**55**	**93**	**433**	**374**	**575**	**819**	**922**	**1,192**	**1,367**
Intel	**10**	**59**	**2**	**-**	**1**	**3**	**9**	**21**	**29**
Google TPU	**1,469**	**1,767**	**2,060**	**1,879**	**3,128**	**6,375**	**8,018**	**9,138**	**9,965**
Broadcom 1	*1,469*	*1,767*	*2,060*	*1,879*	*2,763*	*4,240*	*5,125*	*5,443*	*5,492*
MediaTek	-	-	-	-	*365*	*2,135*	*2,893*	*3,695*	*4,473*
AWS Trainium	**192**	**566**	**920**	**1,398**	**2,107**	**2,961**	**3,845**	**5,130**	**6,415**
Marvell 1 (Trn)		*20*	*307*	*440*	*429*	*427*	*444*	*606*	*771*
Alchip (Trn)	*192*	*546*	*613*	*958*	*1,678*	*2,534*	*3,401*	*4,525*	*5,644*
Microsoft (Maia)	**-**	**-**	**-**	**-**	**-**	**98**	**153**	**196**	**237**
Marvell 2	-	-	-	-	-	*98*	*153*	*196*	*237*
Meta (MTIA)	**41**	**109**	**234**	**233**	**126**	**330**	**902**	**1,133**	**1,335**
Broadcom 2	*41*	*109*	*234*	*233*	*126*	*330*	*902*	*1,133*	*1,335*
OpenAI	**-**	**-**	**-**	**-**	**200**	**746**	**1,216**	**1,480**	**1,699**
Broadcom 3					*200*	*746*	*1,216*	*1,480*	*1,699*
Others	**564**	**1,061**	**2,295**	**7,520**	**10,164**	**8,354**	**5,001**	**4,707**	**4,676**

non-NVDA Racks	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Non-NVDA Racks (AMD, INTC, TPU,etc.)	**35**	**55**	**88**	**164**	**236**	**292**	**303**	**349**	**391**
AMD	**1**	**1**	**6**	**5**	**8**	**11**	**13**	**17**	**19**
Intel	**0**	**1**	**0**	**0**	**0**	**0**	**0**	**0**	**0**
Google TPU	**23**	**28**	**32**	**29**	**49**	**100**	**125**	**143**	**156**
Broadcom 1	23	28	32	29	43	66	80	85	86
MediaTek	0	0	0	0	6	33	45	58	70
AWS Trainium	**3**	**9**	**14**	**22**	**33**	**46**	**60**	**80**	**100**
Marvell 1 (Trn)	0	0	5	7	7	7	7	9	12
Alchip (Trn)	3	9	10	15	26	40	53	71	88
Microsoft (Maia)	**0**	**0**	**0**	**0**	**0**	**2**	**2**	**3**	**4**
Marvell 2	0	0	0	0	0	2	2	3	4
Meta (MTIA)	**1**	**2**	**4**	**4**	**2**	**5**	**14**	**18**	**21**
Broadcom 2	1	2	4	4	2	5	14	18	21
OpenAI	**0**	**0**	**0**	**0**	**3**	**12**	**19**	**23**	**27**
Broadcom 3	0	0	0	0	3	12	19	23	27
Others	**8**	**15**	**32**	**104**	**141**	**116**	**69**	**65**	**65**

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 51: We think the analog semi data center TAM grows from $7.6bn today to $25bn by CY30
Data center analog semi content per rack and GW

Total Racks (k)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total Low Power (k)	51	106	208	238	317	347	309	349	391
100-200MW Racks (k)	*12*	*25*	*57*	*184*	*263*	*229*	*148*	*162*	*184*
<100MW Racks (k)	*39*	*82*	*151*	*54*	*54*	*118*	*161*	*187*	*207*
Total High Power (>600kW; k)	0	0	0	0	5	37	43	39	32
600kW+ Racks (k)	*0*	*0*	*0*	*0*	*5*	*37*	*43*	*34*	*23*
1MW+ Racks (k)	*0*	*0*	*0*	*0*	*0*	*0*	*1*	*6*	*9*
Total Racks (k)	51	106	208	238	322	384	353	388	423
Data Center Content ($mn)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total Low Power	$1,102	$2,313	$4,679	$7,619	$10,501	$10,362	$8,172	$9,117	$10,270
100-200MW Racks	$421	$903	$2,065	$6,694	$9,568	$8,323	$5,392	$5,887	$6,689
<100MW Racks	$681	$1,410	$2,614	$925	$933	$2,039	$2,780	$3,230	$3,581
Total High Power (>600kW)	$0	$0	$0	$0	$1,415	$10,760	$13,218	$14,910	$14,734
600kW+ Racks (k)	$0	$0	$0	$0	$1,415	$10,760	$12,406	$9,777	$6,755
1MW+ Racks (k)	$0	$0	$0	$0	$0	$0	$812	$5,133	$7,979
Total Analog Data Center TAM	$1,102	$2,313	$4,679	$7,619	$11,916	$21,121	$21,389	$24,027	$25,003
Content per Rack	$21,608	$21,728	$22,481	$32,033	$36,973	$54,988	$60,623	$61,891	$59,069
GW Deployed	3	5	10	17	25	37	53	58	60
MW Deployed	2,984	5,453	10,135	16,535	25,199	36,835	52,992	57,885	59,953
Content per GW ($mn)	$369	$424	$462	$461	$473	$573	$404	$415	$417
Content per MW	$0.37	$0.42	$0.46	$0.46	$0.47	$0.57	$0.40	$0.42	$0.42
Low Power GW Deployed	3	5	10	17	24	29	24	28	32
Low Power MW Deployed	2,984	5,453	10,135	16,535	24,341	28,861	24,129	27,621	31,626
Low Power Content per GW	$369	$424	$462	$461	$431	$359	$339	$330	$325
Low Power Content per MW	$0.37	$0.42	$0.46	$0.46	$0.43	$0.36	$0.34	$0.33	$0.32
High Power GW Deployed	0	0	0	0	1	8	29	30	28
High Power MW Deployed	0	0	0	0	858	7,974	28,863	30,264	28,326
High Power Content per GW	$0	$0	$0	$0	$1,649	$1,349	$458	$493	$520
High Power Content per MW	$0.00	$0.00	$0.00	$0.00	$1.65	$1.35	$0.46	$0.49	$0.52

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 52: We think the analog semi TAM for data centers grew +63% in CY25, grows +56% in CY26, and accelerates to +77% YoY in CY27
Data center analog semi content growth

Total Racks YoY Growth (%)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total Low Power (k)	0%	109%	96%	14%	33%	9%	-11%	13%	12%
100-200MW Racks (k)	*0%*	*114%*	*129%*	*224%*	*43%*	*-13%*	*-35%*	*9%*	*14%*
<100MW Racks (k)	*0%*	*107%*	*85%*	*-65%*	*1%*	*119%*	*36%*	*16%*	*11%*
Total High Power (>600kW; k)	0%	0%	0%	0%	0%	660%	18%	-10%	-19%
600kW+ Racks (k)	*0%*	*0%*	*0%*	*0%*	*0%*	*660%*	*15%*	*-21%*	*-31%*
1MW+ Racks (k)	*0%*	*0%*	*0%*	*0%*	*0%*	*0%*	*0%*	*532%*	*55%*
Total Racks (k)	0%	109%	96%	14%	36%	19%	-8%	10%	9%
Analog Content YoY Growth (%)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total Low Power	0%	110%	102%	63%	38%	-1%	-21%	12%	13%
100-200MW Racks	0%	114%	129%	224%	43%	-13%	-35%	9%	14%
<100MW Racks	0%	107%	85%	-65%	1%	119%	36%	16%	11%
Total High Power (>600kW)	0%	0%	0%	0%	0%	660%	23%	13%	-1%
600kW+ Racks (k)	0%	0%	0%	0%	0%	660%	15%	-21%	-31%
1MW+ Racks (k)	0%	0%	0%	0%	0%	0%	0%	532%	55%
Total Racks	0%	110%	102%	63%	56%	77%	1%	12%	4%

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 53: IBC, GPU power, CPU attach, become the most important categories for content pool growth and TAM share CY25-30

Analog data center TAM by component

Detailed TAM ($; mn)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total TAM	$1,102	$2,313	$4,679	$7,619	$11,916	$21,121	$21,389	$24,027	$25,003
Switching Infrastructure	$101	$211	$427	$690	$997	$1,286	$1,159	$1,249	$1,314
PSU	$183	$383	$772	$1,219	$1,818	$2,747	$2,600	$2,629	$2,603
BBU	$25	$53	$108	$176	$273	$472	$462	$457	$438
48V Bus Converter/High Voltage IBC	$81	$171	$346	$566	$990	$2,365	$2,637	$3,335	$3,642
GPU Board	$283	$594	$1,206	$2,023	$3,125	$5,251	$5,278	$6,191	$6,631
Smart NIC Card	$37	$78	$154	$196	$282	$413	$402	$448	$476
Protection	$27	$57	$115	$189	$295	$522	$534	$628	$671
Optical Infrastructure	$119	$250	$512	$904	$1,466	$2,837	$2,921	$3,251	$3,336
HBM/Memory	$52	$109	$221	$362	$541	$811	$786	$916	$988
CPU	$92	$193	$393	$671	$1,149	$2,582	$2,756	$3,068	$3,119
Clocks and Timing	$30	$62	$123	$174	$249	$338	$316	$341	$359
Signal Chain Sensing	$72	$151	$302	$450	$636	$776	$690	$752	$805
Power Conversion System	$0	$0	$0	$0	$95	$721	$848	$764	$622
YoY Growth (%)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total		109.8%	102.3%	62.8%	56.4%	77.2%	1.3%	12.3%	4.1%
Switching Infrastructure		109.7%	102.1%	61.6%	44.3%	29.1%	-9.9%	7.8%	5.3%
PSU		109.7%	101.6%	57.9%	49.2%	51.1%	-5.4%	1.1%	-1.0%
BBU		109.8%	102.4%	63.6%	55.3%	72.7%	-2.0%	-1.2%	-4.2%
48V Bus Converter/High Voltage IBC		109.8%	102.4%	63.6%	75.0%	138.9%	11.5%	26.4%	9.2%
GPU Board		109.9%	103.0%	67.7%	54.5%	68.1%	0.5%	17.3%	7.1%
Smart NIC Card		108.9%	97.4%	27.7%	43.9%	46.1%	-2.6%	11.5%	6.2%
Protection		109.8%	102.4%	63.6%	56.4%	77.1%	2.2%	17.6%	6.9%
Optical Infrastructure		110.1%	104.3%	76.6%	62.3%	93.5%	3.0%	11.3%	2.6%
HBM/Memory		109.8%	102.4%	63.6%	49.5%	49.9%	-3.1%	16.5%	7.9%
CPU		110.0%	103.4%	70.6%	71.2%	124.7%	6.8%	11.3%	1.7%
Clocks and Timing		109.2%	99.2%	40.7%	43.5%	35.8%	-6.6%	8.0%	5.0%
Signal Chain Sensing		109.4%	100.3%	49.0%	41.4%	22.0%	-11.0%	8.9%	7.1%
Power Conversion System		0.0%	0.0%	0.0%	0.0%	660.1%	17.7%	-9.9%	-18.6%
Share of TAM (%)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Switching Infrastructure	9.1%	9.1%	9.1%	9.1%	8.4%	6.1%	5.4%	5.2%	5.3%
PSU	16.6%	16.6%	16.5%	16.0%	15.3%	13.0%	12.2%	10.9%	10.4%
BBU	2.3%	2.3%	2.3%	2.3%	2.3%	2.2%	2.2%	1.9%	1.8%
48V Bus Converter/High Voltage IBC	7.4%	7.4%	7.4%	7.4%	8.3%	11.2%	12.3%	13.9%	14.6%
GPU Board	25.7%	25.7%	25.8%	26.5%	26.2%	24.9%	24.7%	25.8%	26.5%
Smart NIC Card	3.4%	3.4%	3.3%	2.6%	2.4%	2.0%	1.9%	1.9%	1.9%
Protection	2.5%	2.5%	2.5%	2.5%	2.5%	2.5%	2.5%	2.6%	2.7%
Optical Infrastructure	10.8%	10.8%	10.9%	11.9%	12.3%	13.4%	13.7%	13.5%	13.3%
HBM/Memory	4.7%	4.7%	4.7%	4.8%	4.5%	3.8%	3.7%	3.8%	4.0%
CPU	8.4%	8.4%	8.4%	8.8%	9.6%	12.2%	12.9%	12.8%	12.5%
Clocks and Timing	2.7%	2.7%	2.6%	2.3%	2.1%	1.6%	1.5%	1.4%	1.4%
Signal Chain Sensing	6.5%	6.5%	6.4%	5.9%	5.3%	3.7%	3.2%	3.1%	3.2%
Power Conversion System	0.0%	0.0%	0.0%	0.0%	0.8%	3.4%	4.0%	3.2%	2.5%

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 54: SiC and GaN are more prominent in shift to high power
Analog data center TAM by device type

Device TAM ($; mn)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Overall	**$1,102**	**$2,313**	**$4,679**	**$7,619**	**$11,916**	**$21,121**	**$21,389**	**$24,027**	**$25,003**
Silicon	$251	$527	$1,066	$1,733	$2,600	$3,972	$3,832	$4,237	$4,431
SiC	$12	$26	$52	$82	$253	$1,181	$1,371	$1,378	$1,261
GaN	$18	$37	$75	$118	$303	$1,233	$1,443	$1,618	$1,612
Analog	$745	$1,563	$3,163	$5,164	$7,921	$13,111	$13,055	$14,870	$15,695
MCU	$43	$90	$181	$294	$482	$984	$1,038	$1,196	$1,251
Sensors	$34	$70	$142	$229	$358	$641	$650	$725	$752
Low Power	**$1,102**	**$2,313**	**$4,679**	**$7,619**	**$10,501**	**$10,362**	**$8,172**	**$9,117**	**$10,270**
Silicon	$251	$527	$1,066	$1,733	$2,388	$2,357	$1,860	$2,075	$2,338
SiC	$12	$26	$52	$82	$113	$112	$89	$100	$112
GaN	$18	$37	$75	$118	$162	$161	$128	$143	$161
Analog	$745	$1,563	$3,163	$5,164	$7,119	$7,021	$5,532	$6,171	$6,952
MCU	$43	$90	$181	$294	$405	$400	$316	$353	$397
Sensors	$34	$70	$142	$229	$315	$312	$247	$276	$310
High Power	**$0**	**$0**	**$0**	**$0**	**$1,415**	**$10,760**	**$13,218**	**$14,910**	**$14,734**
Silicon	$0	$0	$0	$0	$212	$1,615	$1,972	$2,162	$2,093
SiC	$0	$0	$0	$0	$141	$1,069	$1,281	$1,279	$1,149
GaN	$0	$0	$0	$0	$141	$1,072	$1,315	$1,474	$1,451
Analog	$0	$0	$0	$0	$801	$6,091	$7,523	$8,699	$8,743
MCU	$0	$0	$0	$0	$77	$583	$722	$844	$854
Sensors	$0	$0	$0	$0	$43	$329	$403	$450	$441

Device TAM YoY Growth	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Overall	**-**	**109.8%**	**102.3%**	**62.8%**	**56.4%**	**77.2%**	**1.3%**	**12.3%**	**4.1%**
Silicon	-	109.8%	102.3%	62.6%	50.1%	52.8%	-3.5%	10.6%	4.6%
SiC	-	109.7%	101.6%	57.9%	209.2%	366.3%	16.0%	0.6%	-8.5%
GaN	-	109.7%	101.6%	57.9%	157.2%	306.9%	17.1%	12.1%	-0.3%
Analog	-	109.8%	102.4%	63.3%	53.4%	65.5%	-0.4%	13.9%	5.5%
MCU	-	109.8%	102.2%	62.1%	63.9%	104.2%	5.6%	15.2%	4.6%
Sensors	-	109.7%	102.0%	60.7%	56.6%	78.9%	1.5%	11.6%	3.6%
Low Power	**-**	**109.8%**	**102.3%**	**62.8%**	**37.8%**	**-1.3%**	**-21.1%**	**11.6%**	**12.6%**
Silicon	-	109.8%	102.3%	62.6%	37.8%	-1.3%	-21.1%	11.6%	12.6%
SiC	-	109.7%	101.6%	57.9%	37.5%	-0.6%	-20.3%	11.7%	12.6%
GaN	-	109.7%	101.6%	57.9%	37.5%	-0.6%	-20.3%	11.7%	12.6%
Analog	-	109.8%	102.4%	63.3%	37.9%	-1.4%	-21.2%	11.6%	12.6%
MCU	-	109.8%	102.2%	62.1%	37.8%	-1.2%	-21.0%	11.6%	12.6%
Sensors	-	109.7%	102.0%	60.7%	37.7%	-1.0%	-20.8%	11.6%	12.6%
High Power	**-**	**-**	**-**	**-**	**-**	**660.1%**	**22.8%**	**12.8%**	**-1.2%**
Silicon	-	-	-	-	-	660.1%	22.1%	9.6%	-3.2%
SiC	-	-	-	-	-	660.1%	19.8%	-0.2%	-10.2%
GaN	-	-	-	-	-	660.1%	22.7%	12.1%	-1.6%
Analog	-	-	-	-	-	660.1%	23.5%	15.6%	0.5%
MCU	-	-	-	-	-	660.1%	23.8%	16.8%	1.2%
Sensors	-	-	-	-	-	660.1%	22.6%	11.6%	-1.9%

Device TAM Share (%)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Overall	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**
Silicon	22.8%	22.8%	22.8%	22.7%	21.8%	18.8%	17.9%	17.6%	17.7%
SiC	1.1%	1.1%	1.1%	1.1%	2.1%	5.6%	6.4%	5.7%	5.0%
GaN	1.6%	1.6%	1.6%	1.5%	2.5%	5.8%	6.7%	6.7%	6.4%
Analog	67.6%	67.6%	67.6%	67.8%	66.5%	62.1%	61.0%	61.9%	62.8%
MCU	3.9%	3.9%	3.9%	3.9%	4.0%	4.7%	4.9%	5.0%	5.0%
Sensors	3.0%	3.0%	3.0%	3.0%	3.0%	3.0%	3.0%	3.0%	3.0%
Low Power	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**88.1%**	**49.1%**	**38.2%**	**37.9%**	**41.1%**
Silicon	22.8%	22.8%	22.8%	22.7%	20.0%	11.2%	8.7%	8.6%	9.3%
SiC	1.1%	1.1%	1.1%	1.1%	0.9%	0.5%	0.4%	0.4%	0.4%
GaN	1.6%	1.6%	1.6%	1.5%	1.4%	0.8%	0.6%	0.6%	0.6%
Analog	67.6%	67.6%	67.6%	67.8%	59.7%	33.2%	25.9%	25.7%	27.8%
MCU	3.9%	3.9%	3.9%	3.9%	3.4%	1.9%	1.5%	1.5%	1.6%
Sensors	3.0%	3.0%	3.0%	3.0%	2.6%	1.5%	1.2%	1.1%	1.2%
High Power	**0.0%**	**0.0%**	**0.0%**	**0.0%**	**11.9%**	**50.9%**	**61.8%**	**62.1%**	**58.9%**
Silicon	0.0%	0.0%	0.0%	0.0%	1.8%	7.6%	9.2%	9.0%	8.4%
SiC	0.0%	0.0%	0.0%	0.0%	1.2%	5.1%	6.0%	5.3%	4.6%
GaN	0.0%	0.0%	0.0%	0.0%	1.2%	5.1%	6.1%	6.1%	5.8%
Analog	0.0%	0.0%	0.0%	0.0%	6.7%	28.8%	35.2%	36.2%	35.0%
MCU	0.0%	0.0%	0.0%	0.0%	0.6%	2.8%	3.4%	3.5%	3.4%
Sensors	0.0%	0.0%	0.0%	0.0%	0.4%	1.6%	1.9%	1.9%	1.8%

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 55: Revenue opportunity inflects with high power racks. TXN, ADI, and Infineon enjoy largest revenue share

Analog semi data center revenue by supplier

Company Revenue Summary ($; mn)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total	**$1,103**	**$2,313**	**$4,679**	**$7,619**	**$11,917**	**$21,121**	**$21,389**	**$24,025**	**$25,001**
TXN	$233	$576	$951	$1,533	$2,394	$4,328	$4,460	$5,182	$5,564
ADI	$136	$296	$646	$1,030	$1,770	$3,130	$3,329	$4,004	$4,338
Infineon	$78	$143	$276	$852	$1,542	$3,210	$3,445	$3,953	$4,101
MPWR	$230	$402	$721	$705	$1,205	$2,096	$1,884	$2,146	$2,144
ON	$36	$75	$157	$257	$498	$1,444	$1,636	$1,852	$1,940
STMicro	$55	$114	$215	$348	$549	$1,001	$1,088	$1,224	$1,280
Renesas	$70	$145	$437	$755	$1,162	$1,951	$1,849	$2,106	$2,220
MCHP	$41	$94	$150	$241	$380	$692	$708	$805	$843
Other	$222	$467	$1,125	$1,898	$2,417	$3,269	$2,989	$2,752	$2,571
Low Power (<200kW racks)	**$1,103**	**$2,313**	**$4,679**	**$7,619**	**$10,502**	**$10,363**	**$8,172**	**$9,118**	**$10,270**
TXN	$233	$576	$951	$1,533	$2,151	$2,201	$1,778	$2,031	$2,342
ADI	$136	$296	$646	$1,030	$1,595	$1,691	$1,427	$1,687	$1,955
Infineon	$78	$143	$276	$852	$1,314	$1,376	$1,141	$1,336	$1,528
MPWR	$230	$402	$721	$705	$1,083	$1,122	$798	$891	$950
ON	$36	$75	$157	$257	$411	$491	$426	$517	$629
STMicro	$55	$114	$215	$348	$480	$474	$381	$428	$485
Renesas	$70	$145	$437	$755	$1,040	$1,026	$766	$854	$962
MCHP	$41	$94	$150	$241	$332	$328	$259	$289	$326
Other	$222	$467	$1,125	$1,898	$2,096	$1,652	$1,196	$1,084	$1,094
High Power (>600kW racks)	**$0**	**$0**	**$0**	**$0**	**$1,415**	**$10,759**	**$13,216**	**$14,907**	**$14,731**
TXN	$0	$0	$0	$0	$243	$2,127	$2,682	$3,150	$3,222
ADI	$0	$0	$0	$0	$176	$1,439	$1,903	$2,316	$2,383
Infineon	$0	$0	$0	$0	$228	$1,835	$2,304	$2,617	$2,573
MPWR	$0	$0	$0	$0	$122	$974	$1,086	$1,256	$1,194
ON	$0	$0	$0	$0	$87	$953	$1,209	$1,335	$1,311
STMicro	$0	$0	$0	$0	$69	$527	$707	$797	$795
Renesas	$0	$0	$0	$0	$122	$925	$1,084	$1,252	$1,258
MCHP	$0	$0	$0	$0	$48	$364	$449	$516	$517
Other	$0	$0	$0	$0	$321	$1,616	$1,793	$1,668	$1,478

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 56: We see a small $245mn opportunity today growing to $1.8bn by CY30

Strategic power infrastructure analog semi TAM

Power Infra Analog TAM ($mn)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Low Voltage Infra	**$44**	**$81**	**$150**	**$245**	**$361**	**$428**	**$358**	**$409**	**$469**
High Voltage Infra	**$0**	**$0**	**$0**	**$0**	**$24**	**$288**	**$1,348**	**$1,413**	**$1,323**
Power Infra Analog TAM	**$44**	**$81**	**$150**	**$245**	**$385**	**$716**	**$1,706**	**$1,823**	**$1,792**
Content per MW	**$12,350**	**$12,350**	**$12,350**	**$12,350**	**$12,717**	**$16,203**	**$26,822**	**$26,242**	**$24,904**
GW Deployed	3	5	10	17	25	37	53	58	60
PUE	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.2
Facility GW Deployed	4	7	12	20	30	44	64	69	72
MW Deployed	3,581	6,544	12,162	19,842	30,239	44,202	63,590	69,462	71,943
Content per GW ($mn)	$12	$12	$12	$12	$13	$16	$27	$26	$25
Content per MW	$0.01	$0.01	$0.01	$0.01	$0.01	$0.02	$0.03	$0.03	$0.02
Low Voltage GW Deployed	3	5	10	17	24	29	24	28	32
PUE	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.2
Facility GW Deployed (Low Voltage)	4	7	12	20	29	35	29	33	38
Low Voltage MW Deployed	3,581	6,544	12,162	19,842	29,209	34,634	28,955	33,145	37,952
Low Voltage Content per GW	**$12**	**$12**	**$12**	**$12**	**$12**	**$12**	**$12**	**$12**	**$12**
Low Voltage Content per MW	$0.01	$0.01	$0.01	$0.01	$0.01	$0.01	$0.01	$0.01	$0.01
High Voltage GW Deployed	0	0	0	0	1	8	29	30	28
PUE	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.2
Facility GW Deployed (High Voltage)	0	0	0	0	1	10	35	36	34
High Voltage MW Deployed	0	0	0	0	1,030	9,568	34,635	36,317	33,992
High Voltage Content per GW	**$0**	**$0**	**$0**	**$0**	**$28**	**$36**	**$47**	**$47**	**$47**
High Voltage Content per MW	$0.00	$0.00	$0.00	$0.00	$0.02	$0.03	$0.04	$0.04	$0.04

Analog Content YoY Growth (%)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total Love Voltage Infra	**0%**	**83%**	**86%**	**63%**	**47%**	**19%**	**-16%**	**14%**	**15%**
Total High Voltage Infra	**0%**	**0%**	**0%**	**0%**	**0%**	**1111%**	**367%**	**5%**	**-6%**
Power Infra Analog TAM	**0%**	**83%**	**86%**	**63%**	**57%**	**86%**	**138%**	**7%**	**-2%**
Content per MW	**0%**	**0%**	**0%**	**0%**	**3%**	**27%**	**66%**	**-2%**	**-5%**
GW Deployed	0%	83%	86%	63%	52%	46%	44%	9%	4%
Facility GW Deployed	0%	83%	86%	63%	52%	46%	44%	9%	4%
MW Deployed	0%	83%	86%	63%	52%	46%	44%	9%	4%
Content per GW ($mn)	0%	0%	0%	0%	3%	27%	66%	-2%	-5%
Content per MW	0%	0%	0%	0%	3%	27%	66%	-2%	-5%
Facility GW Deployed (Low Voltage)	0%	83%	86%	63%	47%	19%	-16%	14%	15%
Low Power MW Deployed	0%	83%	86%	63%	47%	19%	-16%	14%	15%
Low Power Content per GW	**0%**	**0%**	**0%**	**0%**	**0%**	**0%**	**0%**	**0%**	**0%**
Low Power Content per MW	0%	0%	0%	0%	0%	0%	0%	0%	0%
Facility GW Deployed (High Voltage)	0%	0%	0%	0%	0%	829%	262%	5%	-6%
High Power MW Deployed	0%	0%	0%	0%	0%	829%	262%	5%	-6%
High Power Content per GW	**0%**	**0%**	**0%**	**0%**	**0%**	**30%**	**29%**	**0%**	**0%**
High Power Content per MW	0%	0%	0%	0%	0%	30%	29%	0%	0%

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 57: SSTs and SSCBs gain traction CY28 onwards
Detailed power infrastructure analog semi TAM

Detailed TAM ($; mn)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total TAM	**$44**	**$81**	**$150**	**$245**	**$385**	**$716**	**$1,706**	**$1,823**	**$1,792**
Energy Storage System/UPS	$28	$51	$95	$156	$244	$407	$717	$773	$778
Solid-State Transformer	$0	$0	$0	$0	$0	$0	$304	$319	$298
Solid-State Circuit Breaker	$0	$0	$0	$0	$0	$67	$243	$255	$239
Facility Cooling Infrastructure	$16	$29	$55	$89	$141	$242	$442	$476	$477
YoY Growth (%)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total		**82.7%**	**85.9%**	**63.1%**	**56.9%**	**86.2%**	**138.1%**	**6.9%**	**-1.7%**
Energy Storage System/UPS		82.7%	85.9%	63.1%	56.6%	66.9%	76.1%	7.9%	0.6%
Solid-State Transformer		-	-	-	-	-	-	4.9%	-6.4%
Solid-State Circuit Breaker		-	-	-	-	-	-	4.9%	-6.4%
Facility Cooling Infrastructure		82.7%	85.9%	63.1%	57.6%	72.0%	82.7%	7.7%	0.1%
Share of TAM (%)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Energy Storage System/UPS	63.6%	63.6%	63.6%	63.6%	63.4%	56.8%	42.0%	42.4%	43.4%
Solid-State Transformer	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	17.8%	17.5%	16.6%
Solid-State Circuit Breaker	0.0%	0.0%	0.0%	0.0%	0.0%	9.4%	14.3%	14.0%	13.3%
Facility Cooling Infrastructure	36.4%	36.4%	36.4%	36.4%	36.6%	33.8%	25.9%	26.1%	26.6%

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 58: Silicon carbide should be the most valuable material in strategic power infrastructure
Analog semi device TAM for power infrastructure

Device TAM ($; mn)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Overall	**$44**	**$81**	**$150**	**$245**	**$385**	**$716**	**$1,706**	**$1,823**	**$1,792**
Silicon	$11	$19	$36	$59	$92	$156	$327	$351	$349
SiC	$18	$33	$62	$101	$158	$312	$804	$857	$838
GaN	$0	$0	$0	$0	$0	$0	$0	$0	$0
Analog	$6	$11	$21	$34	$53	$95	$212	$227	$224
MCU	$7	$12	$23	$38	$59	$108	$252	$269	$265
Sensors	$2	$5	$8	$14	$22	$45	$111	$118	$115
Low Power	**$44**	**$81**	**$150**	**$245**	**$361**	**$428**	**$358**	**$409**	**$469**
Silicon	$11	$19	$36	$59	$86	$102	$86	$98	$112
SiC	$18	$33	$62	$101	$149	$176	$147	$169	$193
GaN	$0	$0	$0	$0	$0	$0	$0	$0	$0
Analog	$6	$11	$21	$34	$50	$59	$49	$56	$65
MCU	$7	$12	$23	$38	$56	$66	$55	$63	$72
Sensors	$2	$5	$8	$14	$20	$24	$20	$23	$26
High Power	**$0**	**$0**	**$0**	**$0**	**$24**	**$288**	**$1,348**	**$1,413**	**$1,323**
Silicon	$0	$0	$0	$0	$6	$54	$241	$253	$237
SiC	$0	$0	$0	$0	$10	$135	$657	$688	$644
GaN	$0	$0	$0	$0	$0	$0	$0	$0	$0
Analog	$0	$0	$0	$0	$3	$36	$163	$171	$160
MCU	$0	$0	$0	$0	$4	$42	$196	$206	$193
Sensors	$0	$0	$0	$0	$1	$21	$91	$95	$89

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Exhibit 59: Power discrete vendors like Infineon, ON, and STMicro positioned to benefit in power infra from opportunities in SSTs, ESS, and SSCBs
Company revenue for analog semis in power infrastructure

Company Revenue Summary ($; mn)	CY22	CY23	CY24	CY25	CY26	CY27	CY28	CY29	CY30
Total	**$44**	**$82**	**$150**	**$245**	**$385**	**$716**	**$1,751**	**$1,870**	**$1,836**
TXN	$3	$4	$8	$14	$22	$40	$89	$95	$94
ADI	$2	$4	$8	$15	$24	$43	$95	$102	$101
Infineon	$7	$12	$30	$52	$105	$205	$503	$538	$528
MPWR	$1	$2	$5	$5	$8	$14	$30	$32	$32
ON	$3	$5	$21	$38	$68	$133	$343	$366	$358
STMicro	$6	$11	$18	$30	$48	$89	$213	$228	$224
Renesas	$2	$4	$7	$12	$19	$36	$82	$88	$86
MCHP	$2	$4	$5	$8	$12	$24	$56	$60	$59
Other	$19	$37	$47	$70	$79	$133	$293	$314	$309
Low Power	**$44**	**$81**	**$150**	**$245**	**$361**	**$428**	**$358**	**$409**	**$469**
TXN	$3	$4	$8	$14	$21	$25	$21	$24	$27
ADI	$2	$4	$8	$15	$22	$26	$22	$25	$29
Infineon	$7	$12	$30	$52	$99	$122	$105	$120	$137
MPWR	$1	$2	$5	$5	$7	$8	$7	$8	$9
ON	$3	$5	$21	$38	$63	$77	$66	$75	$86
STMicro	$6	$11	$18	$30	$45	$53	$44	$51	$58
Renesas	$2	$4	$7	$12	$18	$22	$18	$21	$24
MCHP	$2	$4	$5	$8	$12	$14	$12	$13	$15
Other	$19	$37	$47	$70	$74	$80	$63	$72	$83
High Power	**$0**	**$0**	**$0**	**$0**	**$24**	**$288**	**$1,393**	**$1,461**	**$1,367**
TXN	$0	$0	$0	$0	$1	$15	$68	$72	$67
ADI	$0	$0	$0	$0	$1	$16	$73	$77	$72
Infineon	$0	$0	$0	$0	$7	$83	$398	$418	$391
MPWR	$0	$0	$0	$0	$0	$5	$23	$24	$23
ON	$0	$0	$0	$0	$4	$56	$277	$291	$272
STMicro	$0	$0	$0	$0	$3	$36	$169	$177	$166
Renesas	$0	$0	$0	$0	$1	$14	$64	$67	$63
MCHP	$0	$0	$0	$0	$1	$10	$44	$47	$44
Other	$0	$0	$0	$0	$5	$53	$230	$241	$226

Source: BofA Global Research estimates, Nvidia, AMD, Intel, Marvell, Broadcom, Gartner, Omdia, company reports

BofA GLOBAL RESEARCH

Glossary

AC: Alternating Current

AC/DC: Alternating Current to Direct Current

ADC: Analog-to-Digital Converter

ADI: Analog Devices, Inc.

AI: Artificial Intelligence

ALGM: Allegro MicroSystems

ASIC: Application-Specific Integrated Circuit

ASP: Average Selling Price

ATS: Automatic Transfer Switch

AWS: Amazon Web Services

BBU: Battery Backup Unit

BESS: Battery Energy Storage System

BMC: Baseboard Management Controller

BMS: Battery Management System

CAGR: Compound Annual Growth Rate

CAN: Controller Area Network

CDR: Clock and Data Recovery

CPO: Co-Packaged Optics

CPU: Central Processing Unit

CX: ConnectX

CXL: Compute Express Link

CY: Calendar Year

DAC: Digital-to-Analog Converter

DC: Direct Current

DC/DC: Direct Current to Direct Current

DPU: Data Processing Unit

dv/dt: Change in Voltage over Time

eFuse: Electronic Fuse

EIC: Electronic Integrated Circuit

EMI: Electromagnetic Interference

ESS: Energy Storage System

EV: Electric Vehicle

FET: Field-Effect Transistor

FPGA: Field-Programmable Gate Array

GaN: Gallium Nitride

GB: Grace Blackwell

GPIO: General-Purpose Input/Output

GPU: Graphics Processing Unit

GT/s: Gigatransfers per Second

GW: Gigawatt

GWh: Gigawatt-Hour

HBM: High-Bandwidth Memory

HGX: NVIDIA HGX High-Performance GPU Server Platform

HSC: Hybrid Switched-Capacitor

HV: High Voltage

HVDC: High-Voltage Direct Current

I/O: Input/Output

IBC: Intermediate Bus Converter

IC: Integrated Circuit

ICI: Inter-Chip Interconnect

IEA: International Energy Agency

IGBT: Insulated-Gate Bipolar Transistor

IGCT: Integrated Gate-Commutated Thyristor

IP: Intellectual Property

IT: Information Technology

IVR: Integrated Voltage Regulator

JFET: Junction Field-Effect Transistor

kA: Kiloampere

kG: Kilogram

kV: Kilovolt

kVAC: Kilovolt Alternating Current

kW: Kilowatt

LDO: Low-Dropout Regulator

LFT: Low-Frequency Transformer

LLC: Inductor-Inductor-Capacitor Resonant Converter

LLM: Large Language Model

LPO: Linear Pluggable Optics

LRO: Linear Receive Optics

LT: Long Term

LV: Low Voltage

MCHP: Microchip Technology Inc.

MCU: Microcontroller Unit

MOSFET: Metal-Oxide-Semiconductor Field-Effect Transistor

MPWR: Monolithic Power Systems

MTSI: MACOM Technology Solutions Holdings, Inc.

MV: Medium Voltage

MW: Megawatt

NIC: Network Interface Card

NPO: Near-Packaged Optics

NVDA: NVIDIA Corporation

NVL: NVIDIA NVLink rack-scale platform designation

NVLink: NVIDIA High-Speed GPU Interconnect

NVSwitch: NVIDIA Switch Interconnect

OCP: Open Compute Project

OCS: Optical Circuit Switch

ON: onsemi / ON Semiconductor

ORV3: Open Rack Version 3

OSFP: Octal Small Form Factor Pluggable

PCB: Printed Circuit Board

PCS: Power Conversion System; in sidecar/transient-buffer context, clarify if used as Peak Current Supplier or Power-Capacitor Shelf

PDB: Power Distribution Board

PDN: Power Distribution Network

PDU: Power Distribution Unit

PFC: Power Factor Correction

PIC: Photonic Integrated Circuit

PMIC: Power Management Integrated Circuit

POL: Point of Load

PSFB: Phase-Shifted Full Bridge

PSU: Power Supply Unit

PUE: Power Usage Effectiveness

PWM: Pulse-Width Modulation

RDS(on): On-State Drain-to-Source Resistance

RF200: Rosa Feynman 200

RoCE: RDMA over Converged Ethernet

SAM: Serviceable Available Market

Si: Silicon

SiC: Silicon Carbide

SiGe: Silicon-Germanium

SiPho: Silicon Photonics

SIVR: Substrate Integrated Voltage Regulator

SJ: Superjunction

SMT: Surface-Mount Technology

SoC: System on Chip

SSCB: Solid-State Circuit Breaker

SST: Solid-State Transformer

STMicro: STMicroelectronics

STS: Static Transfer Switch

TAM: Total Addressable Market

TCO: Total Cost of Ownership

TDP: Thermal Design Power

TEC: Thermoelectric Cooler

TIA: Transimpedance Amplifier

TMR: Tunnel Magnetoresistance

ToR: Top of Rack

TPU: Tensor Processing Unit

TXN: Texas Instruments

UALink: Ultra Accelerator Link

UPS: Uninterruptible Power Supply

V: Volt

VAC: Volts Alternating Current

VCSEL: Vertical-Cavity Surface-Emitting Laser

VDC: Volts Direct Current

vGaN: Vertical Gallium Nitride

VRM: Voltage Regulator Module

XPU: Generic Processor Unit, including CPU, GPU, ASIC, TPU, or other accelerator

YoY: Year over Year

ZSC/STC: Zero-Voltage Switched-Capacitor / Switched-Tank Converter

Disclosures

Important Disclosures

FUNDAMENTAL EQUITY OPINION KEY: Opinions include a Volatility Risk Rating, an Investment Rating and an Income Rating. VOLATILITY RISK RATINGS, indicators of potential price fluctuation, are: A - Low, B - Medium and C - High. INVESTMENT RATINGS reflect the analyst's assessment of both a stock's absolute total return potential as well as its attractiveness for investment relative to other stocks within its Coverage Cluster (defined below). Our investment ratings are: 1 - Buy stocks are expected to have a total return of at least 10% and are the most attractive stocks in the coverage cluster; 2 - Neutral stocks are expected to remain flat or increase in value and are less attractive than Buy rated stocks and 3 - Underperform stocks are the least attractive stocks in a coverage cluster. An investment rating of 6 (No Rating) indicates that a stock is no longer trading on the basis of fundamentals. Analysts assign investment ratings considering, among other things, the 0-12 month total return expectation for a stock and the firm's guidelines for ratings dispersions (shown in the table below). The current price objective for a stock should be referenced to better understand the total return expectation at any given time. The price objective reflects the analyst's view of the potential price appreciation (depreciation).

Investment rating	Total return expectation (within 12-month period of date of initial rating)	Ratings dispersion guidelines for coverage cluster[R1]
Buy	≥ 10%	≤ 70%
Neutral	≥ 0%	≤ 30%
Underperform	N/A	≥ 20%

[R1]Ratings dispersions may vary from time to time where BofA Global Research believes it better reflects the investment prospects of stocks in a Coverage Cluster.

INCOME RATINGS, indicators of potential cash dividends, are: 7 - same/higher (dividend considered to be secure), 8 - same/lower (dividend not considered to be secure) and 9 - pays no cash dividend. *Coverage Cluster* is comprised of stocks covered by a single analyst or two or more analysts sharing a common industry, sector, region or other classification(s). A stock's coverage cluster is included in the most recent BofA Global Research report referencing the stock.

BofA Global Research personnel (including the analyst(s) responsible for this report) receive compensation based upon, among other factors, the overall profitability of Bank of America Corporation, including profits derived from investment banking. The analyst(s) responsible for this report may also receive compensation based upon, among other factors, the overall profitability of the Bank's sales and trading businesses relating to the class of securities or financial instruments for which such analyst is responsible.

Other Important Disclosures

From time to time research analysts conduct site visits of covered issuers. BofA Global Research policies prohibit research analysts from accepting payment or reimbursement for travel expenses from the issuer for such visits.
Prices are indicative and for information purposes only. Except as otherwise stated in the report, for any recommendation in relation to an equity security, the price referenced is the publicly traded price of the security as of close of business on the day prior to the date of the report or, if the report is published during intraday trading, the price referenced is indicative of the traded price as of the date and time of the report and in relation to a debt security (including equity preferred and CDS), prices are indicative as of the date and time of the report and are from various sources including BofA Securities trading desks.
The date and time of completion of the production of any recommendation in this report shall be the date and time of dissemination of this report as recorded in the report timestamp.

Recipients who are not institutional investors or market professionals should seek the advice of their independent financial advisor before considering information in this report in connection with any investment decision, or for a necessary explanation of its contents.
Officers of BofAS or one or more of its affiliates (other than research analysts) may have a financial interest in securities of the issuer(s) or in related investments.
Refer to BofA Global Research policies relating to conflicts of interest.
"BofA Securities" includes BofA Securities, Inc. ("BofAS") and its affiliates. Investors should contact their BofA Securities representative or Merrill Global Wealth Management financial advisor if they have questions concerning this report or concerning the appropriateness of any investment idea described herein for such investor. "BofA Securities" is a global brand for BofA Global Research.
Information relating to Non-US affiliates of BofA Securities and Distribution of Affiliate Research Reports:
BofAS and/or Merrill Lynch, Pierce, Fenner & Smith Incorporated ("MLPF&S") may in the future distribute, information of the following non-US affiliates in the US (short name: legal name, regulator): Merrill Lynch (South Africa): Merrill Lynch South Africa (Pty) Ltd., regulated by the Financial Sector Conduct Authority; MLI (UK): Merrill Lynch International, regulated by the Financial Conduct Authority (FCA) and the Prudential Regulation Authority (PRA); BofASE (France): BofA Securities Europe SA is authorized by the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and regulated by the ACPR and the Autorité des Marchés Financiers (AMF). BofA Securities Europe SA ("BofASE") with registered address at 51, rue La Boétie, 75008 Paris is registered under no 842 602 690 RCS Paris. In accordance with the provisions of French Code Monétaire et Financier (Monetary and Financial Code), BofASE is an établissement de crédit et d'investissement (credit and investment institution) that is authorised and supervised by the European Central Bank and the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and regulated by the ACPR and the Autorité des Marchés Financiers. BofASE's share capital can be found at www.bofaml.com/BofASEdisclaimer; BofA Europe (Milan): Bank of America Europe Designated Activity Company, Milan Branch, regulated by the Bank of Italy, the European Central Bank (ECB) and the Central Bank of Ireland (CBI); BofA Europe (Frankfurt): Bank of America Europe Designated Activity Company, Frankfurt Branch regulated by BaFin, the ECB and the CBI; BofA Europe (Zurich): Bank of America Europe Designated Activity Company, Zurich Branch, regulated by the Swiss Financial Market Supervisory Authority FINMA, the ECB and CBI; BofA Europe (Madrid): Bank of America Europe Designated Activity Company, Sucursal en España, regulated by the Bank of Spain, the ECB and the CBI; Merrill Lynch (Australia): Merrill Lynch Equities (Australia) Limited, regulated by the Australian Securities and Investments Commission; Merrill Lynch (Hong Kong): Merrill Lynch (Asia Pacific) Limited, regulated by the Hong Kong Securities and Futures Commission (HKSFC); Merrill Lynch (Singapore): Merrill Lynch (Singapore) Pte Ltd, regulated by the Monetary Authority of Singapore (MAS); Merrill Lynch (Canada): Merrill Lynch Canada Inc, regulated by the Canadian Investment Regulatory Organization; Merrill Lynch (Mexico): Merrill Lynch Mexico, SA de CV, Casa de Bolsa, regulated by the Comisión Nacional Bancaria y de Valores; BofAS Japan: BofA Securities Japan Co., Ltd., regulated by the Financial Services Agency; Merrill Lynch (Seoul): Merrill Lynch International, LLC Seoul Branch, regulated by the Financial Supervisory Service; Merrill Lynch (Taiwan): Merrill Lynch Securities (Taiwan) Ltd., regulated by the Securities and Futures Bureau; BofAS India: BofA Securities India Limited, regulated by the Securities and Exchange Board of India (SEBI); Merrill Lynch (Israel): Merrill Lynch Israel Limited, regulated by Israel Securities Authority; Merrill Lynch (DIFC): Merrill Lynch International (DIFC Branch), regulated by the Dubai Financial Services Authority (DFSA); Merrill Lynch (Brazil): Merrill Lynch S.A. Corretora de Títulos e Valores Mobiliários, regulated by Comissão de Valores Mobiliários; Merrill Lynch KSA Company: Merrill Lynch Kingdom of Saudi Arabia Company, regulated by the Capital Market Authority.
This information: has been approved for publication and is distributed in the United Kingdom (UK) to professional clients and eligible counterparties (as each is defined in the rules of the FCA and the PRA) by MLI (UK), which is authorized by the PRA and regulated by the FCA and the PRA - details about the extent of our regulation by the FCA and PRA are available from us on request; has been approved for publication and is distributed in the European Economic Area (EEA) by BofASE (France), which is authorized by the ACPR and regulated by the ACPR and the AMF; has been considered and distributed in Japan by BofAS Japan, a registered securities dealer under the Financial Instruments and Exchange Act in Japan, or its permitted affiliates; is issued and distributed in Hong Kong by Merrill Lynch (Hong Kong) which is regulated by HKSFC; is issued and distributed in Taiwan by Merrill Lynch (Taiwan); is issued and distributed in India by BofAS India; and is issued and distributed in Singapore to institutional investors and/or accredited investors (each as defined under the Financial Advisers Regulations) by Merrill Lynch (Singapore) (Company Registration No 198602883D). Merrill Lynch (Singapore) is regulated by MAS. Merrill Lynch Equities (Australia) Limited (ABN 65 006 276 795), AFS License 235132 (MLEA) distributes this information in Australia only to 'Wholesale' clients as defined by s.761G of the Corporations Act 2001. With the exception of Bank of America N.A., Australia Branch, neither MLEA nor any of

its affiliates involved in preparing this information is an Authorised Deposit-Taking Institution under the Banking Act 1959 nor regulated by the Australian Prudential Regulation Authority. No approval is required for publication or distribution of this information in Brazil and its local distribution is by Merrill Lynch (Brazil) in accordance with applicable regulations. Merrill Lynch (DIFC) is authorized and regulated by the DFSA. Information prepared and issued by Merrill Lynch (DIFC) is done so in accordance with the requirements of the DFSA conduct of business rules. BofA Europe (Frankfurt) distributes this information in Germany and is regulated by BaFin, the ECB and the CBI. BofA Securities entities, including BofA Europe and BofASE (France), may outsource/delegate the marketing and/or provision of certain research services or aspects of research services to other branches or members of the BofA Securities group. You may be contacted by a different BofA Securities entity acting for and on behalf of your service provider where permitted by applicable law. This does not change your service provider. Please refer to the Electronic Communications Disclaimers for further information.

This information has been prepared and issued by BofAS and/or one or more of its non-US affiliates. The author(s) of this information may not be licensed to carry on regulated activities in your jurisdiction and, if not licensed, do not hold themselves out as being able to do so. BofAS and/or MLPF&S is the distributor of this information in the US and accepts full responsibility for information distributed to BofAS and/or MLPF&S clients in the US by its non-US affiliates. Any US person receiving this information and wishing to effect any transaction in any security discussed herein should do so through BofAS and/or MLPF&S and not such foreign affiliates. Hong Kong recipients of this information should contact Merrill Lynch (Asia Pacific) Limited in respect of any matters relating to dealing in securities or provision of specific advice on securities or any other matters arising from, or in connection with, this information. Singapore recipients of this information should contact Merrill Lynch (Singapore) Pte Ltd in respect of any matters arising from, or in connection with, this information. For clients that are not accredited investors, expert investors or institutional investors Merrill Lynch (Singapore) Pte Ltd accepts full responsibility for the contents of this information distributed to such clients in Singapore.

General Investment Related Disclosures:

Taiwan Readers: Neither the information nor any opinion expressed herein constitutes an offer or a solicitation of an offer to transact in any securities or other financial instrument. No part of this report may be used or reproduced or quoted in any manner whatsoever in Taiwan by the press or any other person without the express written consent of BofA Securities.

This document provides general information only, and has been prepared for, and is intended for general distribution to, BofA Securities clients. Neither the information nor any opinion expressed constitutes an offer or an invitation to make an offer, to buy or sell any securities or other financial instrument or any derivative related to such securities or instruments (e.g., options, futures, warrants, and contracts for differences). This document is not intended to provide personal investment advice and it does not take into account the specific investment objectives, financial situation and the particular needs of, and is not directed to, any specific person(s). This document and its content do not constitute, and should not be considered to constitute, investment advice for purposes of ERISA, the US tax code, the Investment Advisers Act or otherwise. Investors should seek financial advice regarding the appropriateness of investing in financial instruments and implementing investment strategies discussed or recommended in this document and should understand that statements regarding future prospects may not be realized. Any decision to purchase or subscribe for securities in any offering must be based solely on existing public information on such security or the information in the prospectus or other offering document issued in connection with such offering, and not on this document.

Securities and other financial instruments referred to herein, or recommended, offered or sold by BofA Securities, are not insured by the Federal Deposit Insurance Corporation and are not deposits or other obligations of any insured depository institution (including, Bank of America, N.A.). Investments in general and, derivatives, in particular, involve numerous risks, including, among others, market risk, counterparty default risk and liquidity risk. No security, financial instrument or derivative is suitable for all investors. Digital assets are extremely speculative, volatile and are largely unregulated. In some cases, securities and other financial instruments may be difficult to value or sell and reliable information about the value or risks related to the security or financial instrument may be difficult to obtain. Investors should note that income from such securities and other financial instruments, if any, may fluctuate and that price or value of such securities and instruments may rise or fall and, in some cases, investors may lose their entire principal investment. Past performance is not necessarily a guide to future performance. Levels and basis for taxation may change.

This report may contain a short-term trading idea or recommendation, which highlights a specific near-term catalyst or event impacting the issuer or the market that is anticipated to have a short-term price impact on the equity securities of the issuer. Short-term trading ideas and recommendations are different from and do not affect a stock's fundamental equity rating, which reflects both a longer term total return expectation and attractiveness for investment relative to other stocks within its Coverage Cluster. Short-term trading ideas and recommendations may be more or less positive than a stock's fundamental equity rating.

BofA Securities is aware that the implementation of the ideas expressed in this report may depend upon an investor's ability to "short" securities or other financial instruments and that such action may be limited by regulations prohibiting or restricting "shortselling" in many jurisdictions. Investors are urged to seek advice regarding the applicability of such regulations prior to executing any short idea contained in this report.

Foreign currency rates of exchange may adversely affect the value, price or income of any security or financial instrument mentioned herein. Investors in such securities and instruments, including ADRs, effectively assume currency risk.

BofAS or one of its affiliates is a regular issuer of traded financial instruments linked to securities that may have been recommended in this report. BofAS or one of its affiliates may, at any time, hold a trading position (long or short) in the securities and financial instruments discussed in this report.

BofA Securities, through business units other than BofA Global Research, may have issued and may in the future issue trading ideas or recommendations that are inconsistent with, and reach different conclusions from, the information presented herein. Such ideas or recommendations may reflect different time frames, assumptions, views and analytical methods of the persons who prepared them, and BofA Securities is under no obligation to ensure that such other trading ideas or recommendations are brought to the attention of any recipient of this information.

In the event that the recipient received this information pursuant to a contract between the recipient and BofAS for the provision of research services for a separate fee, and in connection therewith BofAS may be deemed to be acting as an investment adviser, such status relates, if at all, solely to the person with whom BofAS has contracted directly and does not extend beyond the delivery of this report (unless otherwise agreed specifically in writing by BofAS). If such recipient uses the services of BofAS in connection with the sale or purchase of a security referred to herein, BofAS may act as principal for its own account or as agent for another person. BofAS is and continues to act solely as a broker-dealer in connection with the execution of any transactions, including transactions in any securities referred to herein.

Copyright and General Information:

Copyright 2026 Bank of America Corporation. All rights reserved. iQdatabase® is a registered service mark of Bank of America Corporation. This information is prepared for the use of BofA Securities clients and may not be redistributed, retransmitted or disclosed, in whole or in part, or in any form or manner, without the express written consent of BofA Securities. This document and its content is provided solely for informational purposes and cannot be used for training or developing artificial intelligence (AI) models or as an input in any AI application (collectively, an AI tool). Any attempt to utilize this document or any of its content in connection with an AI tool without explicit written permission from BofA Global Research is strictly prohibited. BofA Global Research utilizes AI, including machine learning and other technologies, to enhance the services we provide to our clients. These technologies assist our analysts in various aspects of their work, including but not limited to data analysis, content extraction, content creation, data aggregation and summarization and identifying relevant information from diverse sources. All AI-driven processes are subject to review by BofA Global Research employees. BofA Global Research information is distributed simultaneously to internal and client websites and other portals by BofA Securities and is not publicly-available material. Any unauthorized use or disclosure is prohibited. Receipt and review of this information constitutes your agreement not to redistribute, retransmit, or disclose to others the contents, opinions, conclusion, or information contained herein (including any investment recommendations, estimates or price targets) without first obtaining express permission from an authorized officer of BofA Securities.

Materials prepared by BofA Global Research personnel are based on public information. Facts and views presented in this material have not been reviewed by, and may not reflect information known to, professionals in other business areas of BofA Securities, including investment banking personnel. BofA Securities has established information barriers between BofA Global Research and certain business groups. As a result, BofA Securities does not disclose certain client relationships with, or compensation received from, such issuers. To the extent this material discusses any legal proceeding or issues, it has not been prepared as nor is it intended to express any legal conclusion, opinion or advice. Investors should consult their own legal advisers as to issues of law relating to the subject matter of this material. BofA Global Research personnel's knowledge of legal proceedings in which any BofA Securities entity and/or its directors, officers and employees may be plaintiffs, defendants, co-defendants or co-plaintiffs with or involving issuers mentioned in this material is based on public information. Facts and views presented in this material that relate to any such proceedings have not been reviewed by, discussed with, and may not reflect information known to, professionals in other business areas of BofA Securities in connection with the legal proceedings or matters relevant to such proceedings.

This information has been prepared independently of any issuer of securities mentioned herein and not in connection with any proposed offering of securities or as agent of any issuer of any securities. None of BofAS any of its affiliates or their research analysts has any authority whatsoever to make any representation or warranty on behalf of the issuer(s). BofA Global Research policy prohibits research personnel from disclosing a recommendation, investment rating, or investment thesis for review by an issuer prior to the publication of a research report containing such rating, recommendation or investment thesis.

Any information relating to sustainability in this material is limited as discussed herein and is not intended to provide a comprehensive view on any sustainability claim with respect to any issuer or security.

Any information relating to the tax status of financial instruments discussed herein is not intended to provide tax advice or to be used by anyone to provide tax advice. Investors are urged to seek tax advice based on their particular circumstances from an independent tax professional.
The information herein (other than disclosure information relating to BofA Securities and its affiliates) was obtained from various sources and we do not guarantee its accuracy. This information may contain links to third-party websites. BofA Securities is not responsible for the content of any third-party website or any linked content contained in a third-party website. Content contained on such third-party websites is not part of this information and is not incorporated by reference. The inclusion of a link does not imply any endorsement by or any affiliation with BofA Securities. Access to any third-party website is at your own risk, and you should always review the terms and privacy policies at third-party websites before submitting any personal information to them. BofA Securities is not responsible for such terms and privacy policies and expressly disclaims any liability for them.
All opinions, projections and estimates constitute the judgment of the author as of the date of publication and are subject to change without notice. Prices also are subject to change without notice. BofA Securities is under no obligation to update this information and BofA Securities ability to publish information on the subject issuer(s) in the future is subject to applicable quiet periods. You should therefore assume that BofA Securities will not update any fact, circumstance or opinion contained herein.
Certain outstanding reports or investment opinions relating to securities, financial instruments and/or issuers may no longer be current. Always refer to the most recent research report relating to an issuer prior to making an investment decision.
In some cases, an issuer may be classified as Restricted or may be Under Review or Extended Review. In each case, investors should consider any investment opinion relating to such issuer (or its security and/or financial instruments) to be suspended or withdrawn and should not rely on the analyses and investment opinion(s) pertaining to such issuer (or its securities and/or financial instruments) nor should the analyses or opinion(s) be considered a solicitation of any kind. Sales persons and financial advisors affiliated with BofAS or any of its affiliates may not solicit purchases of securities or financial instruments that are Restricted or Under Review and may only solicit securities under Extended Review in accordance with firm policies.
Neither BofA Securities nor any officer or employee of BofA Securities accepts any liability whatsoever for any direct, indirect or consequential damages or losses arising from any use of this information.