BERNSTEIN
SOCIETE GENERALE GROUP

Global Semiconductors and Asia Hardware

Global Semis & Asia Hardware: Artificial Intelligence... Sooner or Later - Energizing the Future of AI Data Centers


David Dai, CFA
+852 2918 5704
david.dai@bernsteinsg.com


Alex Wang, CFA
+852 2123 2613
alex.wang@bernsteinsg.com


Stacy A. Rasgon, Ph.D.
+1 213 559 5917
stacy.rasgon@bernsteinsg.com


Carmine Milano
+44 20 7762 1857
carmine.milano@bernsteinsg.com


Jack Lin
+852 2123 2683
jack.lin@bernsteinsg.com


Juho Hwang
+852 2123 2632
juho.hwang@bernsteinsg.com


Shirley Yang, CFA
+852 2123 2660
shirley.yang@bernsteinsg.com


Ethan Xu
+852 2123 2634
ethan.xu@bernsteinsg.com


Alrick Shaw
+1 917 344 8454
alrick.shaw@bernsteinsg.com


Arpad von Nemes
+1 917 344 8461
arpad.vonnemes@bernsteinsg.com

This report is part of a global series in which our analysts consider what their sector will look like by the early 2030s when AI has been fully integrated and commercialized. We believe the power semis and components will be essential to power the AI GPUs and data centers.

AI accelerator power (GWs) grows 2.7x in 3 years. Based on our forecasts of AI accelerator shipments, we anticipate the equivalence of 49.5 GW of AI accelerators to be shipped in 2028, which is a big increase from 18.6 GW equivalent of accelerator shipment in 2025. The growth is mainly driven by the exponential rise of power per accelerator: Nvidia H100 (2022) has TDP (total design power) of up to 700W, while Rubin Ultra (2027) will be at ~3500W, a 5x increase in 5 years. As the accelerators get deployed, the rest of the supply chain, especially power and data center infrastructure, have to catch up and meet the big capacity requirement.

AI power semi content grows 4x in 3 years. Following the strong momentum in AI-related demand, we now forecast 2028 AI power semi content to be US$ 7.7bn, vs. 2025 market of $1.9bn. The AI power semi content growth (4x) is ahead of GW growth (2.7x). In 2025, power semi content is $0.10/W; in 2028, we expect this to become $0.16/W, as driven by several upgrades to improve power efficiency: The transition to 800V architectures, the migration of DC–DC Stage 1 and Stage 2 in terms of materials (from silicon to SiC/GaN) and topology (from lateral to vertical power), and the adoption of SSTs in the future. We see additional upside potential for power $/W beyond this period, as these transitions are still gaining momentum toward the end of the decade. Our AI Power Semi market model is available here: **AI Server power Model**

Infineon followed by Renesas are the big winners in this space. We estimate that **Infineon (OP)** has ~45% share in AI power semi, as Infineon offers the most complete product portfolio for every stage (AC-DC, IBC and POL) and every material (Si, SiC, and GaN), and its strong production capacity and quality ensures its leadership position. **Renesas (OP)** has ~11% share in 2025, having Nvidia as a major customer and is gaining share in AWS, and its AI revenue contribution is doubling to 8% this year. **TXN (MP)** has relatively low current exposure (with overall datacenter at ~9% of revenue, ~$1.5B across both AI and non-AI applications) though it is growing quickly (up >60% YoY in 2025). **ADI (MP)** had ~$2B in overall datacenter business in 2025 though split across ATE, optical, and power; power (including delivery and control, hot-swap, and high-performance protections solutions) represented ~30% of this (~$600M), with the overall datacenter segment growing ~50% last year.

Server component is showing similar growth trajectory. We revise up NVIDIA server power component TAM, and expect it to grow 60-70% YoY each year in 2026-27, and another c.45% in 2028, reaching US$14B. Adding demands from ASICs and China market, overall AI power component TAM will be US$23B in three years. **Delta Electronics (OP)** is the clear leader as the market moving to power rack architecture starting from 2H26.

First Published: 16 Apr 2026 20:30 UTC Completion Date: 16 Apr 2026 16:00 UTC

BERNSTEIN TICKER TABLE

			15 Apr 2026		TTM		Adjusted EPS			Adjusted P/E (x)		
			Closing	Price	Rel.							
Ticker	Rating	Cur	Price	Target	Perf.	Cur	2025A	2026E	2027E	2025A	2026E	2027E
IFX.GR (Infineon)	O	EUR	44.58	52.00	50.6%	EUR	1.38	1.65	2.33	57.2	44.3	26.2
6723.JP (Renesas)	O	JPY	2,765.50	3,300.00	20.5%	JPY	181.61	235.36	254.83	15.2	11.8	10.9
2308.TT (Delta)	O	TWD	1,845.00	1,830.00	431.6%	TWD	23.09	35.75	53.78	79.9	51.6	34.3
ADI (Analog Devices)	M	USD	347.94	375.00	67.6%	USD	7.79	11.62	13.38	44.7	29.9	26.0
TXN (Texas Instruments)	M	USD	216.29	205.00	17.5%	USD	5.45	6.26	6.88	39.7	34.5	31.4
NVDA (NVIDIA)	O	USD	198.87	300.00	60.2%	USD	4.77	8.88	11.94	41.7	22.4	16.7
AMD (Advanced Micro)	M	USD	258.12	265.00	162.2%	USD	4.17	6.53	13.23	61.8	39.5	19.5
EDME			1,528.28									
JPL			2,434.83									
ASIAX			1,798.81									
SPX			7,022.95									

O - Outperform, M - Market-Perform, U - Underperform, NR - Not Rated, CS - Coverage Suspended

2308.TT, TXN estimate is Reported EPS; IFX.GR, 2308.TT, TXN valuation is Reported P/E (x); NVDA base year is 2026;

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

Infineon: We rate Infineon Outperform with a price target of EUR 52.00

Renesas: We rate Renesas Outperform with a price target of JPY 3,300

Delta: We rate Delta Outperform with a price target of NT$1,830.00.

ADI (Market-Perform, $375): ADI is executing well with both cyclical and idiosyncratic drivers, though the shares remain expensive.

TXN (Market-Perform, $205): TXN shares feel fully valued in the current environment.

NVDA (Outperform, $300.00): The datacenter opportunity is enormous, and still early, with material upside still possible.

AMD (MP, $265): AI expectations remain high, but server strength may help, and the new Meta deal seems supportive to FY27 numbers.

DETAILS

Following the strong momentum in AI-related demand as Hyperscalers continue to expand Capex, **we have revised our AI server power semiconductor forecast for 2027 to US$ 5.5Bn, up from our previous estimate of US$ 5Bn. We also introduce a 2028 forecast of US$ 7.7Bn, representing 40% YoY growth**. We expect the AI power semiconductor market to expand fourfold between 2025 and 2028, reflecting a 60% CAGR. For our near-term projections, we have made modest upward adjustments: 2026 increases by US$ 150Mn and 2025 by less than US$ 100Mn (Exhibit 1).

Nvidia will remain the dominant customer in the AI power semiconductor market. However, its share of the total market is expected to decline from 73% in 2025, or US$ 1.4Bn, to 56% in 2028, equivalent to US$ 4.3Bn—still implying a robust 45% CAGR over the next three years. We expect the strongest growth to come from ASICs, driven primarily by Google TPUs, which are projected to double annually, reaching approximately US$ 2Bn in 2028 from just US$ 200Mn in 2025 (Exhibit 2).

Our AI Power Semi model is available for download here: AI Server power Model.

EXHIBIT 1: **We revised our estimates for the AI power semis market. While CY25-26 changed only marginally, we raised our CY27 forecast by half billion. We also introduced a new estimate for CY28, projecting the market to reach US$ 7.7Bn.**

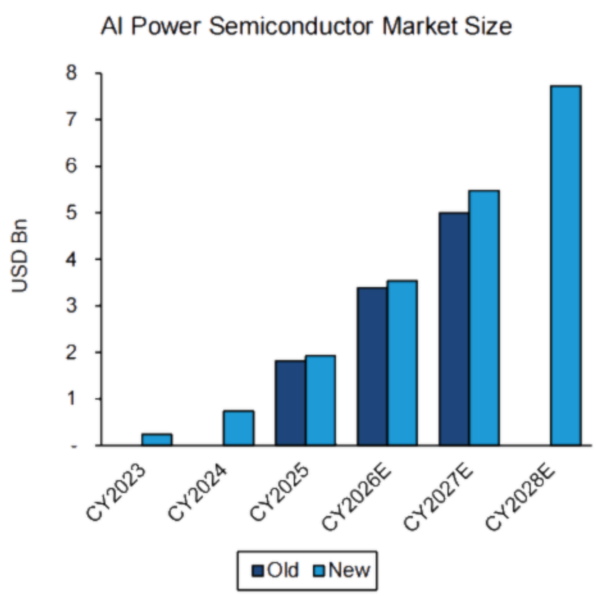

Source: Company reports, Bernstein analysis and estimates

EXHIBIT 2: **While we expect Nvidia to remain the dominant customer in the AI power semiconductor market, we anticipate that ASICs will grow at a faster rate, driven primarily by Google.**

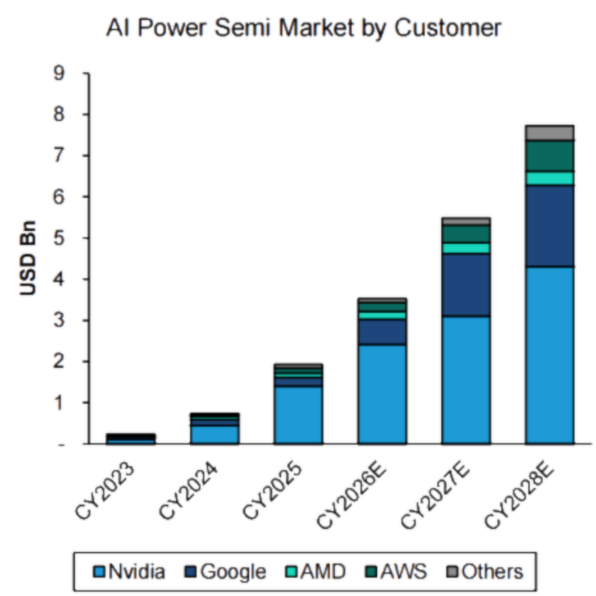

Source: Company reports, Bernstein analysis and estimates

RISING TDP DRIVES THE BULK OF AI POWER SEMIS EXPANSION

While we anticipate AI accelerator unit growth may slow—peaking in 2027 and easing in 2028 as performance per XPU improves faster than deployments—we still expect the AI power market to grow strongly, primarily driven by rising GPU TDP that more than offsets lower unit volumes. AI demand will remain structurally strong even as unit growth tapers, because hyperscalers increasingly meet capacity needs through higher per-chip throughput rather than more chips (Exhibit 3). Nvidia's next platforms materially lift compute per GPU/system: the GB300 NVL72 (Blackwell Ultra) delivers 1.5x more dense FP4 Tensor Core FLOPS and 2x higher attention performance versus Blackwell GPUs. Looking ahead, the Rubin platform targets a 4x reduction in the number of GPUs needed to train MoE models versus Blackwell, alongside up to a 10x reduction in inference token cost—supporting more output with fewer accelerators. Google's TPU roadmap shows a similar step-change: TPU v6 provides a 4.7x increase in peak compute per chip versus TPU v5e.

Higher compute density also drives higher power density: server PSUs must deliver more power within fixed rack dimensions as AI processor power requirements rise. For this reason, we anticipate TDP will increase materially

across next-generation accelerators (Exhibit 4):

- **Nvidia:** TDP has scaled from 300W in V100 to 400W in A100 and to the 700W-class with H100. For Blackwell, NVIDIA discloses B200 as configurable around 1,200W per GPU. We expect Rubin TDP to increase by 50% to 2.1 kW, then rise a further 65% with Rubin Ultra to ~3.5 kW. We anticipate Feynman TDP to be ~5.2 kW.

- **Google:** Google disclosed TPU v4 at ~340W, rising to ~1 kW with TPU v8. We expect this rapid growth to continue, with TPU v9 nearly tripling to ~2.6 kW and TPU v10 increasing further to ~3.6 kW.

- **AMD:** AMD discloses MI250/MI250X at 500W TDP, MI300X at 750W, MI325X at 1,000W, and MI355X at 1,400W. We expect MI455X at ~1,960W and MI500X/MI550X at ~3,136W.

- **AWS:** We estimate Trainium 2 TDP at ~500W, increasing to ~700W for Trainium3, with Trainium4 reaching ~1,050W.

EXHIBIT 3: **AI accelerator unit growth is expected to peak around 2027 and slow in 2028, but overall AI demand will stay strong because hyperscalers will rely more on higher performance per chip rather than deploying more chips to meet capacity needs.**

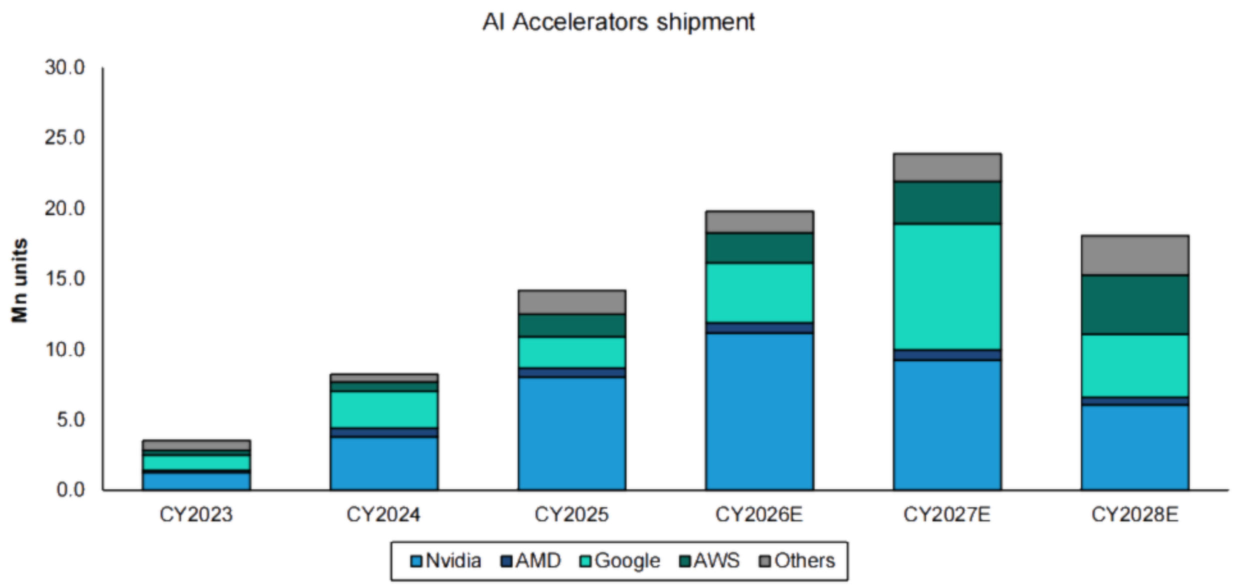

Source: Company reports, Bernstein analysis and estimates

EXHIBIT 4: **We anticipate AI accelerator TDP will rise rapidly with each new generation as systems scale to meet sharply higher compute and memory bandwidth demands.**

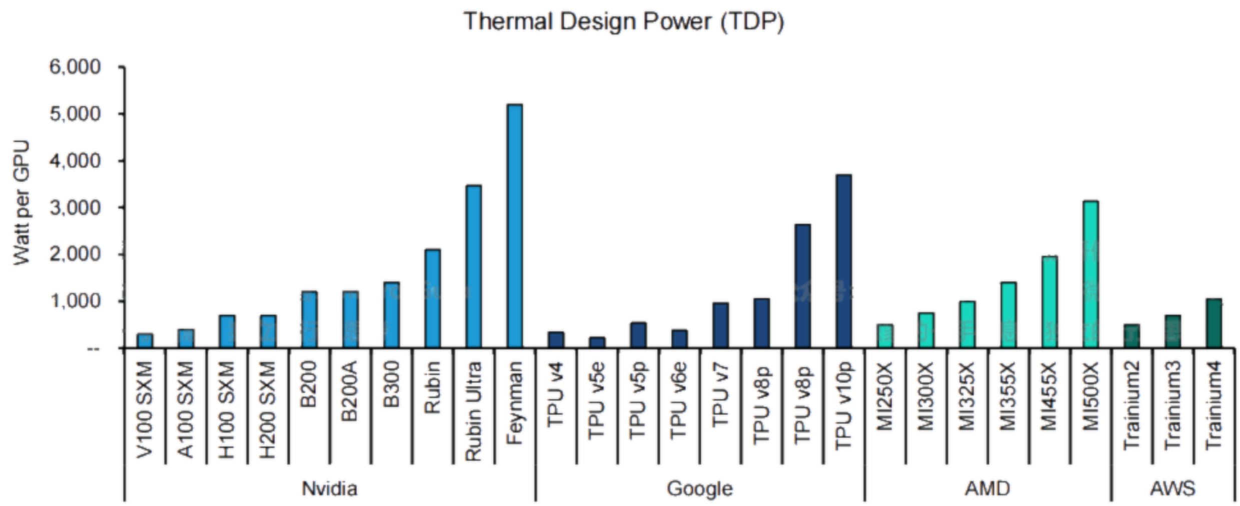

Source: Company reports, Bernstein analysis and estimates

Looking at overall AI power demand, computed as watts per AI accelerators multiplied by accelerators shipments, **we estimate total AI accelerator power demand equate to ~37GW in 2028, up from 12.8GW in 2025, implying a 43% CAGR over 2025–2028E.** This suggests that TDP growth accounts for the vast majority of AI power semis growth, which is expanding at ~60% CAGR 25-28E (Exhibit 5).

Assuming AI accelerators represent 70% of total server power, we estimate total AI data center power demand of ~18GW in 2025, rising to ~53GW by 2028. Note that, due to inventory builds and deployment lead times, actual data center GW deployment likely lags these figures by ~6 months (Exhibit 6).

EXHIBIT 5: **We estimate AI accelerator power reaches ~37GW in 2028 versus 12.8GW in 2025, a 43% CAGR over 2025-2028E. Most of the growth in AI power is coming from higher TDP.**

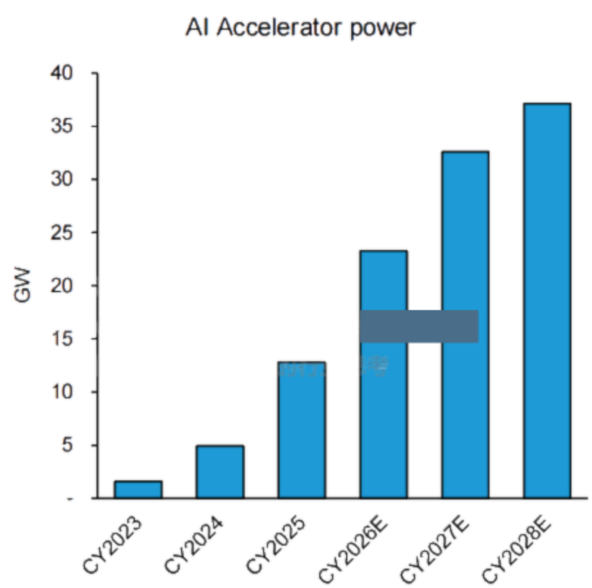

Source: Company reports, Bernstein analysis and estimates

EXHIBIT 6: **If we assume GPUs to account for ~70% of server power, total AI data center demand scales from ~18GW in 2025 to ~53GW in 2028.**

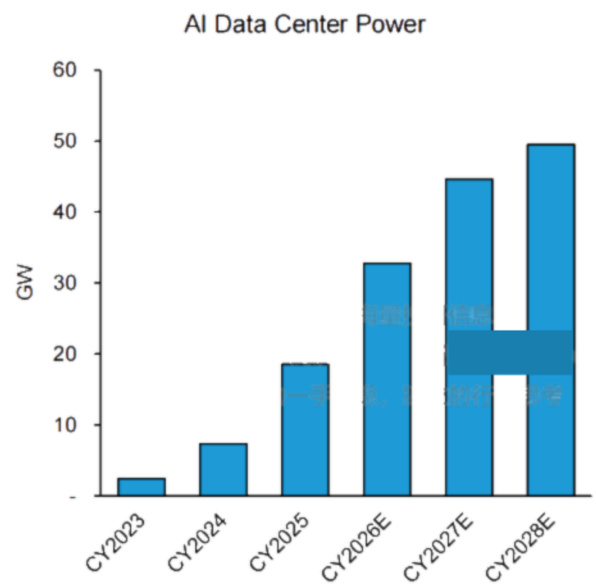

Source: Company reports, Bernstein analysis and estimates

POWER $ CONTENT PER WATT IS ALSO RISING, ADDING ANOTHER TAILWIND TO AI POWER GROWTH

The power flow in an AI data center typically has three stages:

1. **PSU/BBU (AC-DC):** Alternating current (AC) from the grid is converted to direct current (DC) (50v now, 800v in the future) by the server's power supply unit (PSU).

2. **DC-DC stage 1 (Intermediate Bus Controller, IBC):** The high voltage or mid voltage DC is stepped down to an intermediate bus voltage, which may be 12 V in some systems or higher voltages such as 48 V or 54 V in newer designs.

3. **DC-DC stage 2 (Point-of-load, POL):** a local DC-to-DC converter near the chip reduces the voltage again to about 1 V at the point of load, which is the level the GPU core needs to operate efficiently. This layered approach matters because the GPU draws very high current at very low voltage, so the final converter must sit close to the chip to minimize losses and voltage drop.

We estimate that the final step (DC–DC Stage 2) accounts for the largest share of total power semiconductor content per server, at 40%. This is followed by the PSU/BBU stage, which represents 32%. The second step in the process, DC–DC Stage 1, accounts for only 15% of total power semiconductor content per server. The remaining 13% is attributed to other components such as hot-swap controllers, eFuses, and related devices (Exhibit 7).

TrendForce estimates AI racks will keep increasing in power density, moving from 200 kW in 2025 with Nvidia NVL72 GB300 toward 1 MW from 2028 with Feynman. The key change is the move from 54 VDC rack buses toward 800 VDC, because higher voltage reduces current and therefore cuts losses, cable size, and heat. In parallel, the architecture shifts from 1-phase AC to 3-phase AC, and then to high-voltage DC (HV-DC), which fits much larger loads more efficiently. The adoption of solid-state transformers (SSTs) is the next step, since they can help convert power at much higher scale for future multi-megawatt racks (Exhibit 9).

The transition to 800V architectures, the adoption of SSTs, and the migration of both DC–DC Stage 1 and Stage 2 in terms of materials (from silicon to GaN) and topology (from lateral to vertical power) are all driving an increase in power semiconductor cost per watt. We will discuss each of these transitions in more detail below. **We anticipate that power semiconductor cost per watt will rise rapidly, from $0.10 in 2025 to $0.12 in 2027, before further accelerating to $0.16 in 2028, representing a total increase of 50%. We see additional upside potential beyond this period, as these transitions are only gaining momentum toward the end of the decade (Exhibit 8).**

EXHIBIT 7: **The DC-DC Stage 2, which converts current to 1V and directly powers the GPU, represents 40% of total power content per server. The AC-DC stage (PSU) is the second-largest component, accounting for 32%.**

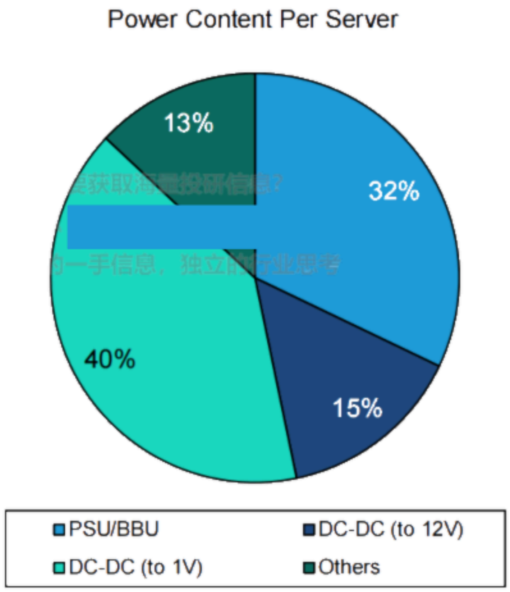

Source: Company reports, Bernstein analysis and estimates

EXHIBIT 8: **Power semiconductor cost per watt is expected to increase, growing by approximately 50% from $0.10 in 2025 to $0.16 in 2028.**

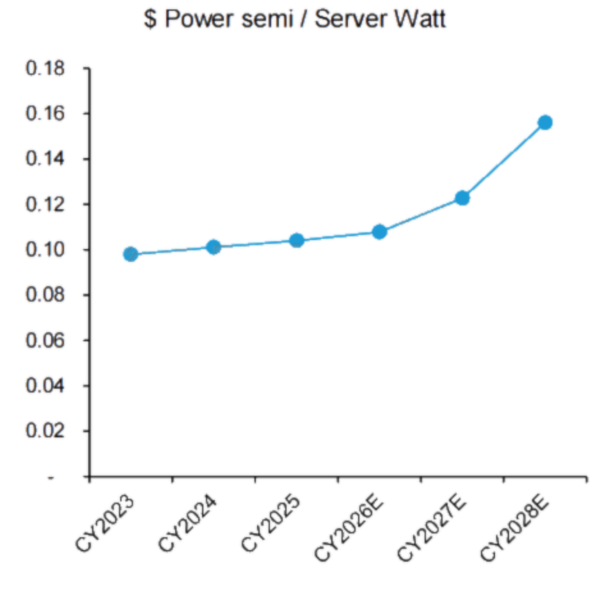

Source: Company reports, Bernstein analysis and estimates

EXHIBIT 9: **TrendForce expects that future GPU generations drive racks from tens of kilowatts toward 1 MW, forcing a shift from 54 VDC to 800 VDC buses, HV-DC distribution, and multi-megawatt SST for efficient power conversion.**

	2023	2024	2025	2026	2027	From 2028
Generation	NVL36 GH200	NVL72 GB200	NVL72 GB300	NVL72 VR200	NVL72 VR300	Feynman
Rack Power	36 kW	130kW	200 kW	250 kW	600 kW	1 MW
Rack Power Distribution	1Φ VAC	1Φ VAC	1Φ VAC	3Φ VAC	3Φ VAC	HV-DC
Power Shelf	2-3 x 18 kW	4-6 x 33 kW	6-8 x 33 kW	4 x 55 kW	8 x 72 kW	Multi MW SST
PSU	3 kW (x6)	5.5 kW (x6)	5.5 kW (x6)	8 kW (x3x2)	12 kW (x3x2)	/
PSU output / Rack voltage bus	54 VDC	54 VDC	54 VDC	54 VDC	800 VDC or +/- 400 VDC crosslink cable	800 VDC or +/- 400 VDC Busbar

Source: TrendForce, Bernstein analysis

High voltage direct current (HVDC) roadmap: Traditional AC power distribution faces inefficiencies from multiple power conversions between the grid and AI accelerators. To address this, Nvidia is promoting an 800V DC architecture, enabling higher rack power density without increased current—thus reducing resistance, heat, and copper use. In this model, AC grid power is converted once to 800V DC via a solid-state transformer (SST), eliminating multiple AC/DC stages. The 800V DC then feeds directly to racks and steps down through DC/DC modules, as shown in Delta's "Next-gen 800V server feed by SST" (Exhibit 10). Since facility-wide 800V DC requires SST deployment and stricter safety protocols, near-term adoption will likely center on standalone power racks. While the multiple steps of power conversion in infrastructure remains, power rack would free IT rack space, and enhance power reliability by adopting BBU and supercapacitors. It supplies flexible DC output (e.g., 50V or 800V).

We expect power content in Nvidia AI racks to expand sharply with the introduction of power rack. Compared to previous model (2026 Sector Outlook), we raise PSU ASP per kw assumptions to US$145/kW and keep it constant across different generation of AI servers. We expect overall power content per server rack will jump from US$50k per GB300 rack, to c.85K for VR200 Power shelf version, and c.US$140K for VR200 power rack and US$320K for Rubin Ultra rack. While we keep PSU ASP/ kW constant across different generations of AI servers, it's likely to see an ASP upside as higher spec of PSU (12kW/ 18kW) will adopt more expensive chips based on SiC. Compared to power shelf version where all power components are squeeze in the IT rack, a standalone power rack allows 100% BBU adoption and have space for supercapacitor and PDU that leads to a significantly higher US$ content for solution providers like **Delta Electronics**. Exhibit 12-Exhibit 13 indicate that Delta is the clear leader in AI power components.

Using Asia Semis team's HBM/ CoWoS model (link) and applying a discount for chips deployment, we estimate that the TAM for Nvidia's AI server power component will reach c.US$5.7B in 2026 and US$9.7B in 2027E, up from US$3.5B in 2025. This implies a 40% adoption rate for power rack in 2027 and 100% in 2028. We further extend our outlook to 2028, projecting continued strong growth of c.45%, lifting the market size to US$14B (Exhibit 14). Beyond NVIDIA, we see ASIC and China demand emerging as key incremental growth drivers for AIDC power components from 2025 onward. We model ASICs and China to represent 35%-38% of power TAM over 2025-28, and total power TAM growth of 67%/74% in 2026-27, and another 46% growth in 2028. (Exhibit 15)

EXHIBIT 10: **HVDC transition roadmap proposed by Delta**

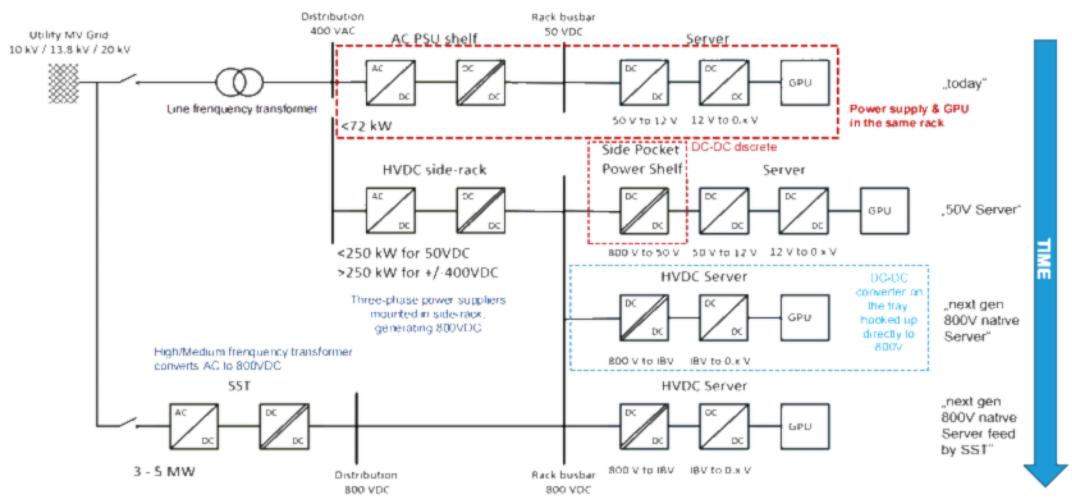

Source: Delta, Bernstein analysis

EXHIBIT 11: **Power content will increase significantly in future AI server racks, and we expect VR Ultra's power content to grow by 7.4x vs. GB 200**

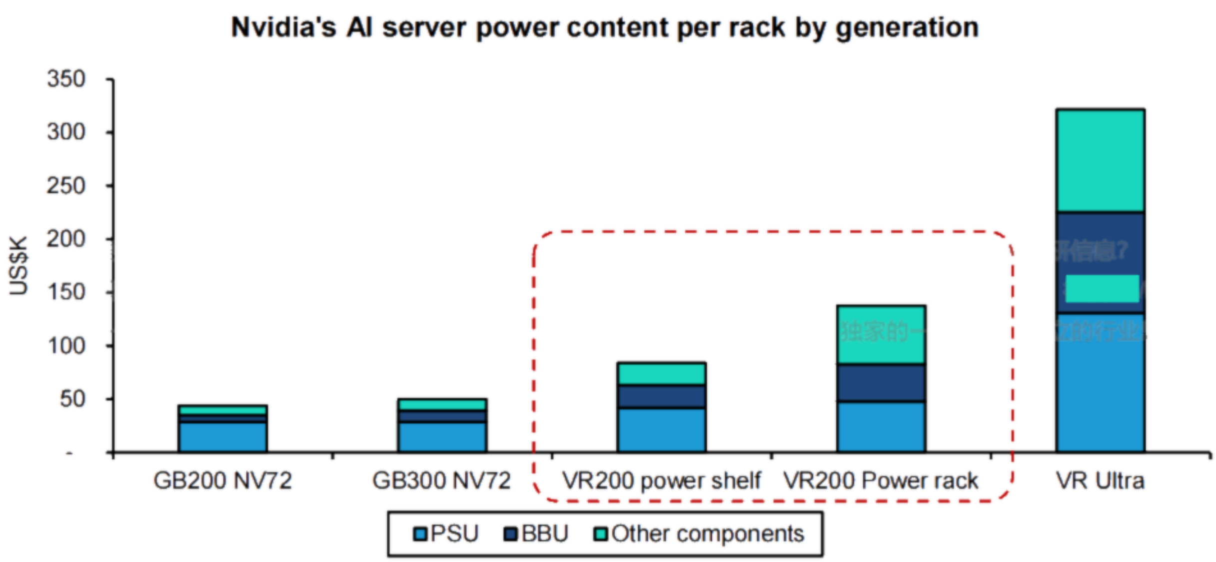

BBU's content factored in adoption rate
Source: Bernstein estimates and analysis

EXHIBIT 12: **Delta captures a lion share in AI server power market, particularly in the AC-DC market**

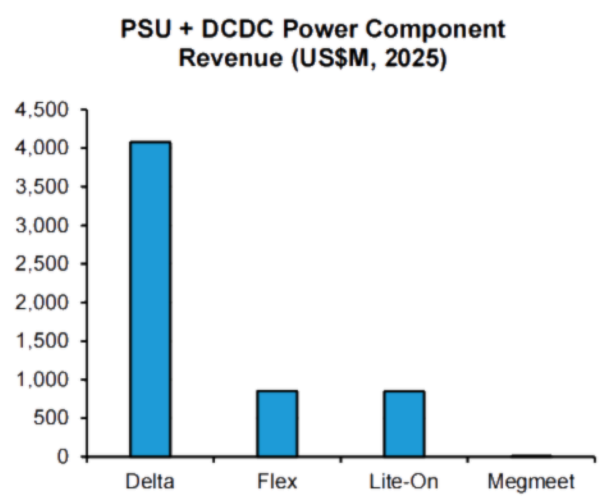

Source: Company reports, Bernstein estimates and analysis

EXHIBIT 13: **Delta's AI power revenue mix is highest among peers**

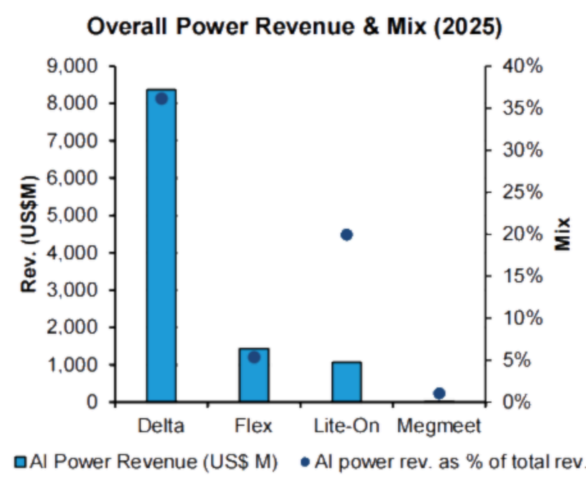

Source: Company reports, Bernstein analysis and estimates

EXHIBIT 14: **We raise our power component TAM estimates for NVIDIA servers to US$5.7B / US$9.7B in 2026 / 2027, underpinned by higher per-rack power content and stronger NVIDIA rack shipments with sustained CSP capex investment**

Power component TAM for Nvidia server

	2025	2026E	2027E	2028E
Nvidia chip shipment (K)	6,064	10,894	9,225	6,051
Total chip deployed in CSP sites	5,231	7,576	8,626	6,527
GB/VR rack and 8-GPU server shipment (K)				
Rack shipment	30	61	88	55
8-GPU server	388	402	275	69
Power content per rack/server (US$K)				
Blended power content in rack	45	51	85	242
yoy%		*14%*	*65%*	*186%*
8-GPU server	6	6	8	13
Power component TAM (US$M)				
GB/Rubin/Feynman rack	1,327	3,106	7,413	13,320
8-GPU server	2,184	2,574	2,326	929
Total power component TAM (US$M)	3,511	5,680	9,739	14,248
yoy%		*62%*	*71%*	*46%*

Source: Bernstein estimates and analysis

EXHIBIT 15: **We model total AIDC power TAM to grow at 60%+ CAGR over 2025-28E, with ~35% from ASICs/ China**

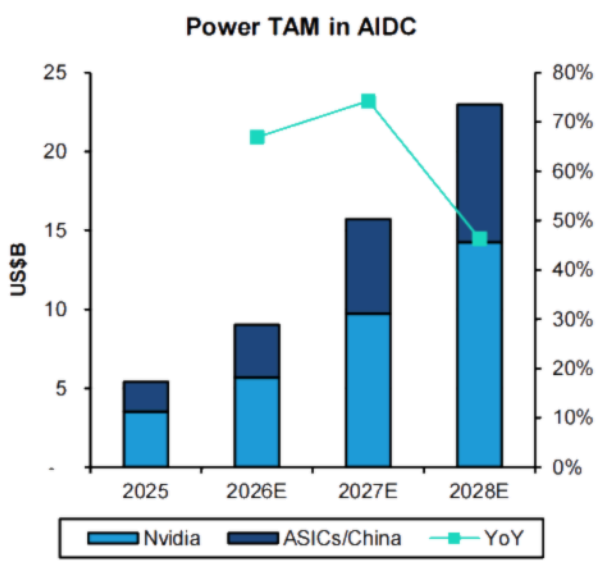

Source: Bernstein analysis and estimates

VERTICAL POWER DELIVERY: THE NEW STANDARD FOR POWERING >1 KA AI PROCESSORS

The transition from lateral power delivery (LPD) to vertical power delivery (VPD) is fundamentally about shortening the distance between the voltage regulator and the AI processor. In LPD, regulators sit around the XPU or CPU on the board, and current flows laterally through board traces before reaching the chip. That approach worked when processors drew a few hundred amps, but at AI-era currents well above 1,000A the resistance of those traces turns into large power losses and voltage drops. VPD moves the regulator directly underneath the processor so current flows vertically into the package, eliminating most of the lateral path and dramatically reducing power-delivery losses (Exhibit 16, Exhibit 17).

VPD brings clear benefits but also meaningful challenges. On the positive side, moving from lateral to true vertical power can cut the power-delivery-network resistance from 90-140μΩ down to roughly 10–15 μΩ, enabling 2,500–3,000 A GPU currents with manageable losses and improving end-to-end efficiency by a couple of percentage points at very high load. Those few percent translate into tens of watts saved per accelerator, which scale to megawatts at data-center level. The trade-offs are mechanical, thermal, and electrical. Putting regulators directly under the hottest component requires new board stack-ups, new cooling strategies, and integration of bypass capacitors into the module because the usual surface capacitors are displaced. Infineon's ultra-high current density modules can deliver about **85% reduction in PDN losses and ~55% reduction in size** versus a lateral "down" solution.

Infineon frames four main power-delivery paths. First, lateral power delivery keeps regulators on the top side of the motherboard around the processor: mature and low-cost, but limited at very high currents. Second, backside lateral-vertical power moves modules to the backside of the board or substrate, shortening the path and cutting losses but introducing new thermal management on a side of the card that historically did not dissipate heat. Third, full VPD places modules directly under the processor, making the current path almost purely vertical and minimizing PDN resistance, at the cost of more complex mechanical and thermal design. Finally, embedded or integrated vertical power goes a step further, embedding regulators into the motherboard itself or even into the processor substrate, which trims a further slice of resistance but pushes co-packaged thermal design to its limits (Exhibit 18).

Infineon has been a pioneer in this shift. Over more than a decade it has pushed module current density from legacy levels to about 2.0 A/mm² in its latest generation aimed at vertical mounting. This has been enabled by packaging innovations such as embedding the silicon die into the module substrate, integrating dozens of capacitors inside the module, and adopting fast-response topologies like TLVR to handle sharp AI load steps.

EXHIBIT 16: **In PLD a 48 V input bus converter feeds a voltage regulator on the board, which distributes 12-6 V, then 0.8 V, laterally through traces to the ASIC package.**

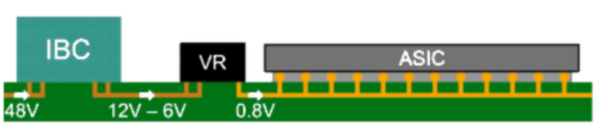

Source: Company reports, Bernstein analysis

EXHIBIT 17: **Instead VPD moves the VRM beneath the ASIC, so the board area beside the chip is freed for signals while 12-6 V is delivered vertically, reducing PDN losses and supporting much higher currents.**

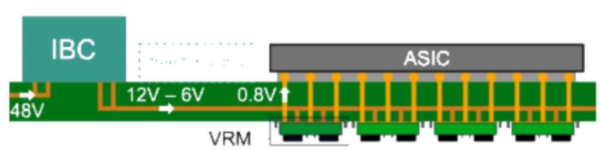

Source: Company reports, Bernstein analysis

EXHIBIT 18: **Infineon Power delivery path roadmap: lateral regulators on the board (cheap, current-limited), backside mounting (shorter paths, new thermal issues), true vertical under-chip modules (minimal PDN, complex mechanics/thermals), and fully embedded regulators in board or substrate (lowest resistance, hardest co-packaged thermal design).**

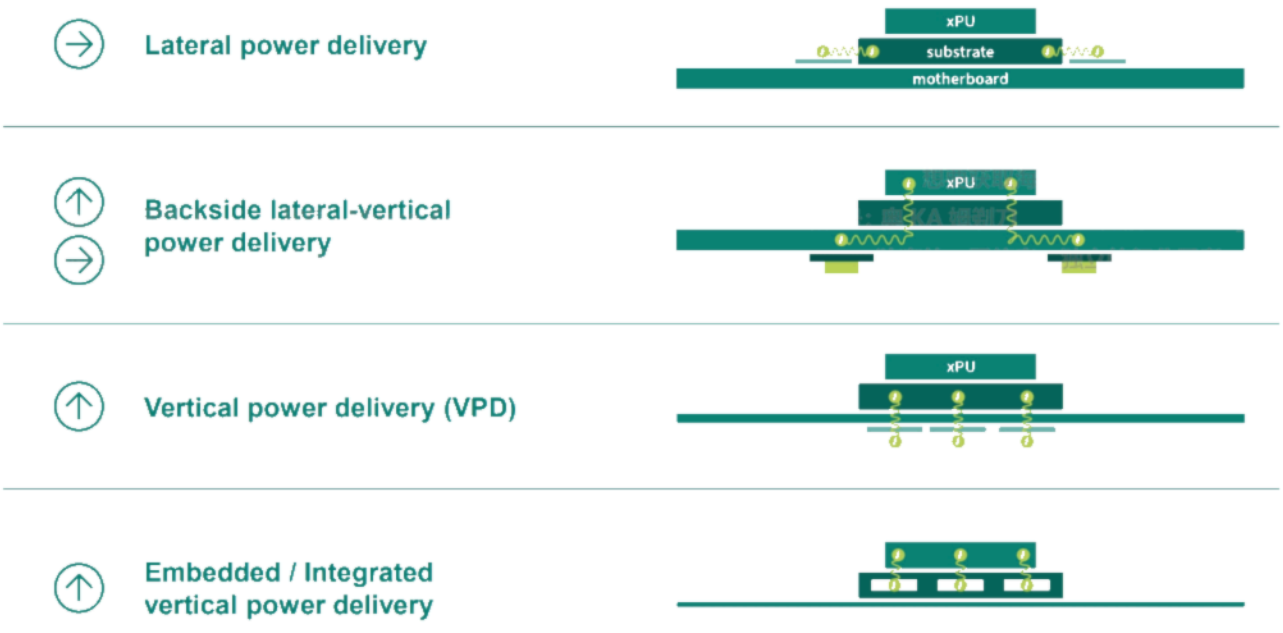

Source: Company reports, Bernstein analysis

THE DRIVE FOR HIGHER EFFICIENCY ALSO TRIGGERS THE ADOPTION OF SIC AND GAN IN DATA CENTERS.

PSUs and DC-DC stages in AI data centers increasingly **rely on a hybrid mix of Si, SiC and GaN, because no single technology optimizes cost, efficiency and power density across all conversion stages.** Si, SiC, and GaN Each technology has its own strengths and hence is positioned for power supply design as follows:

- Silicon is: Perfect for legacy or conventional designs, Optimized cost / performance solutions, suitable for traditional topologies up to 150 kHz switching

- Silicon Carbide: SiC is perfect for half-bridge (HB) hard-switching topologies like totem-pole PFC requiring high voltages (≥400V) and high temperature operating conditions, suitable for resonant topologies like LLC up to 500 kHz

- Gallium Nitride: GaN is perfect for soft-switching topologies such as LLC, for very high frequency (≥500 kHz) operation, Best in efficiency x power density FOM

AI data-center power architectures are moving from low-voltage, silicon-only designs to higher-voltage systems that strategically combine Si, SiC and GaN devices along the power path. Traditional 12 V racks at 10–20 kW rely on 50 Hz transformers and silicon-based AC-DC/DC-DC stages, achieving about 70–80% efficiency for 100–300 W CPUs. Current AI data centers instead use 48 V racks above 100 kW and >1,000 W GPUs, raising end-to-end efficiency to roughly 80–90% while introducing a first, moderate share of wide-bandgap: SiC in high-voltage PFC or rectification, GaN in some high-frequency DC-DC converters, with Si still dominant overall. Looking forward, Infineon depicts 800 V AI data centers using SiC-based solid-state transformers to create an 800 V dc bus and >1 MW racks with >3,000 W GPUs, where SiC/GaN AC-DC stages and GaN 48 V converters push rack efficiency beyond 90% and maximize power density. (Exhibit 19).

Infineon describes this opportunity under the theme "We power AI – from grid to core." Its portfolio spans solid-state transformers and HVDC distribution at the grid interface, high-efficiency PFC and PSUs at rack level, 48 V intermediate bus converters, and finally VRM and power-module solutions near the AI processor itself. Across this stack, the combination of Si, CoolSiC and CoolGaN allows Infineon to tune each stage for the required balance of cost, efficiency, power density and thermal performance, and positions the company to benefit structurally from rising WBG content in next-generation AI data-center infrastructure.

EXHIBIT 19: **AI data centers are shifting from low-voltage, silicon-only power designs to higher-voltage architectures combining silicon, SiC and GaN. Infineon frames this as "grid to core" power, spanning grid-level HVDC and SSTd, rack-level PFC and PSUs, 48-V converters, and on-chip VRMs near processors, supporting efficient delivery across the AI stack.**

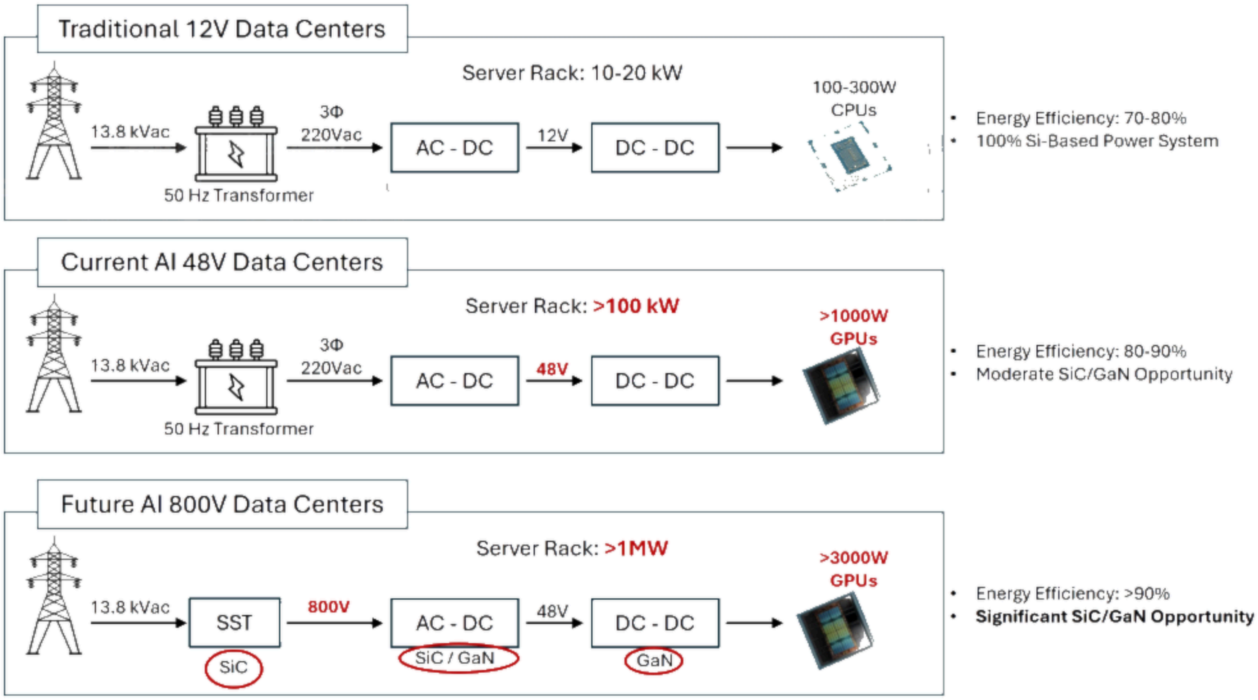

Source: TrendForce, Bernstein analysis

TrendForce anticipates wide-band gap (WBG) to increase significantly as a % of the DC power device market, rising from 11% in 2025 to more than 30% by the end of the decade, gaining share primarily in stage 1 and 2. SIC is expected to steadily increase as a % of the DC power device market, but to remain quite limited at MSD. On the other hand, they expect GaN to grow more rapidly, from 4% in 2024 to reaching 26% penetration by 2030 (Exhibit 20).

Yole estimates the GaN device market reached US$ 355Mn in CY2024, increasing 33% YoY from US$ 260Mn in 2023. They estimate Infineon represents 11% of market share, growing considerably from 2023. They estimate Innoscience is a strong market leader in 2024 with a 30% share, albeit losing some market share compared to 2023 (Exhibit 21).

We believe that suppliers which master all three technologies will benefit the most — and that is Infineon in our coverage, followed by Renesas, which has strength in GaN but not SiC (though it could tap into Wolfspeed's SiC expertise).

EXHIBIT 20: **TrendForce anticipates SIC to steadily increase as a % of the DC power device market, but to remain quite limited at MSD. On the other hand, they expect GaN to grow more rapidly, reaching 26% penetration by 2030.**

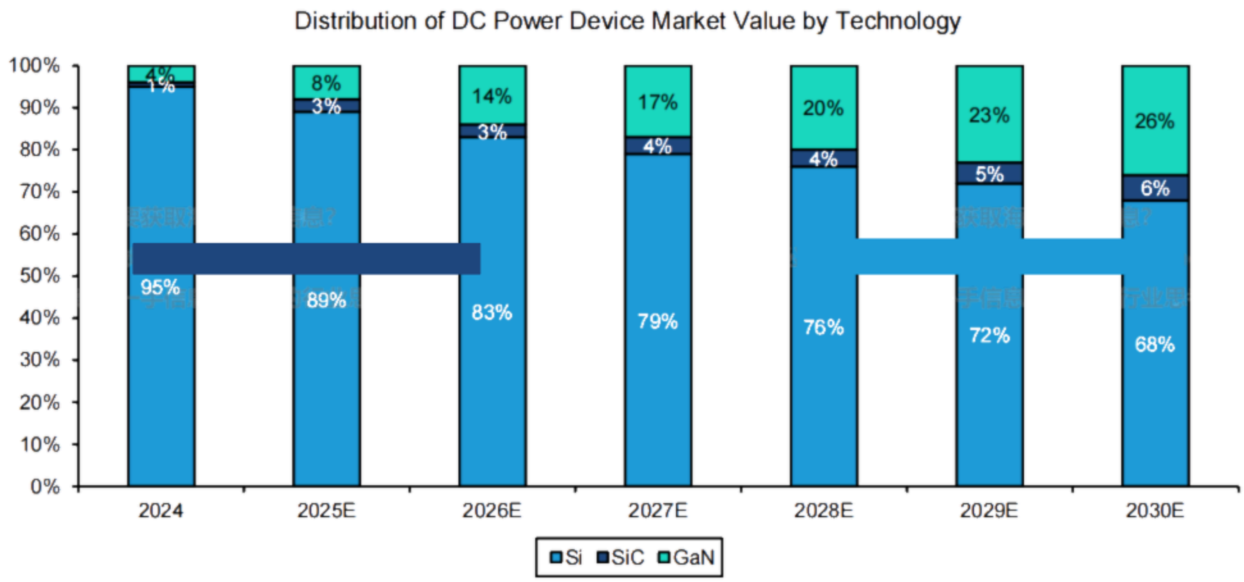

Source: TrendForce, Bernstein analysis

EXHIBIT 21: **Yole estimates the GaN device market reached US$ 355Mn in CY2024, increasing 33% YoY from US$ 260Mn in 2023. They estimate Infineon represents 11% of market share, growing considerably from 2023.**

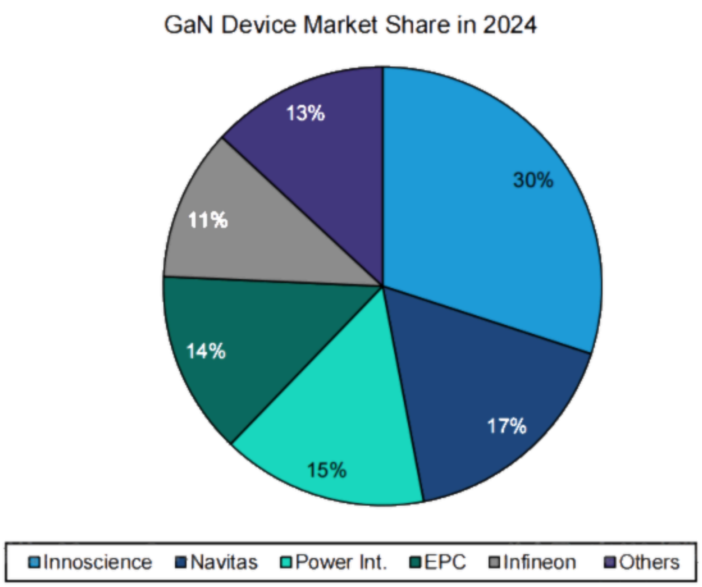

Others includes Nexperia, Texas Instr., Renesas, etc.
Source: Yole, Bernstein analysis

MAJOR PLAYERS

- **Infineon** is the leading player in the AI power market with more than 40% market share (we estimate 43%, rising to 47%), supported by its superior technology, mastery of all three key materials, and strong supply capabilities. As a pioneer in this sector, Infineon benefits from the broadest portfolio and a diversified client base, supplying virtually all AI accelerators, with IR confirming more than 10 active customers. While new entrants such as Texas Instr, ADI, STMicro (not covered), and

Onsemi (not covered) are entering the market, we believe Infineon will successfully defend its share (Exhibit 22).

Infineon entered the power data center market more than five years ago, but AI power revenue gained momentum only in FY25, when it exceeded ~EUR 750Mn. We expect this to double next year, reaching EUR 1.6Bn, slightly above guidance. Infineon disclosed an additional US$ 500Mn of revenue in non-AI DC. We estimate AI power sales could reach EUR 3.3Bn by 2028, representing around 17% of total revenue (Exhibit 23).

EXHIBIT 22: **Infineon is the market leader in the AI power market with 43% market share and we estimate it will rise to 47% in 2026.**

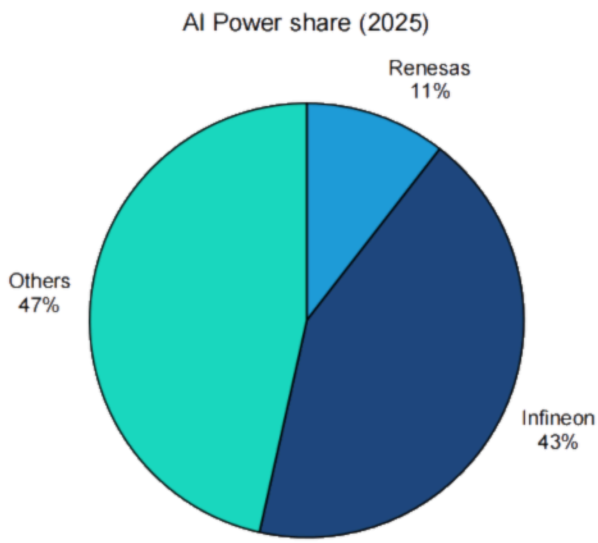

Source: Company reports, Bernstein analysis and estimates

EXHIBIT 23: **Infineon's AI power revenue exceeded ~EUR 750Mn in FY25, and we expect it to double next year, reaching EUR 1.6Bn, slightly above guidance. We estimate AI power sales could reach EUR 3.3Bn by 2028, representing around 17% of total revenue.**

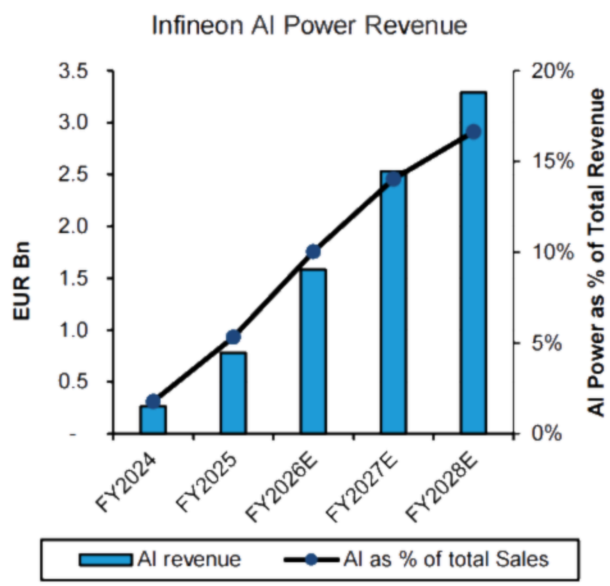

Source: Company reports, Bernstein analysis and estimates.

- **Renesas** is a key DC-DC power supplier to NVIDIA's GPUs and has been gaining content share at Amazon (AWS) since late 2025. The company's AI-related revenue doubled in both 2024 and 2025, reaching ~4% (~JPY117bn) of Renesas' revenue in 2026E (Exhibit 24). We forecast AI sales to grow a further 50% and 40% in 2026E and 2027E, respectively, lifting AI exposure to ~13% (~JPY250bn) of total revenue by 2028E (Exhibit 25). Beyond AI power solutions, Renesas also supplies memory interface ICs and timing/clock ICs to AI platforms. In 2025, digital power products accounted for ~50-60% of AI sales, memory interface ICs for 30%+, and timing/clock ICs for 10%+.

EXHIBIT 24: **Renesas' AI Revenue by Product**

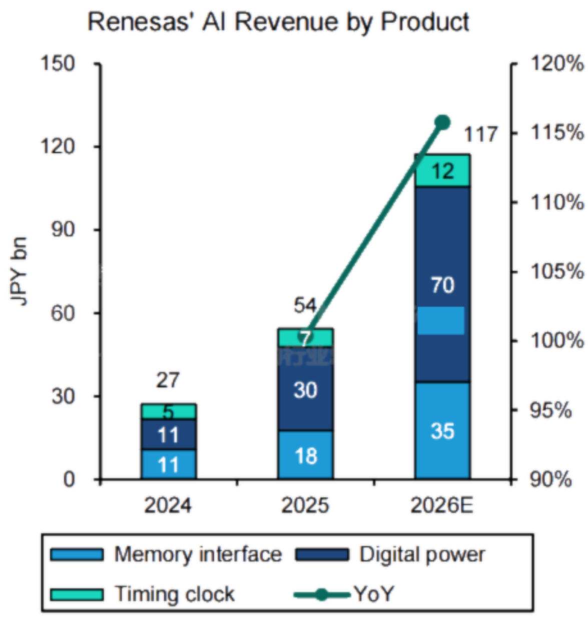

Source: Company disclosures, Bernstein estimates and analysis.

EXHIBIT 25: **Renesas's AI Sales as of % of Total Revenue**

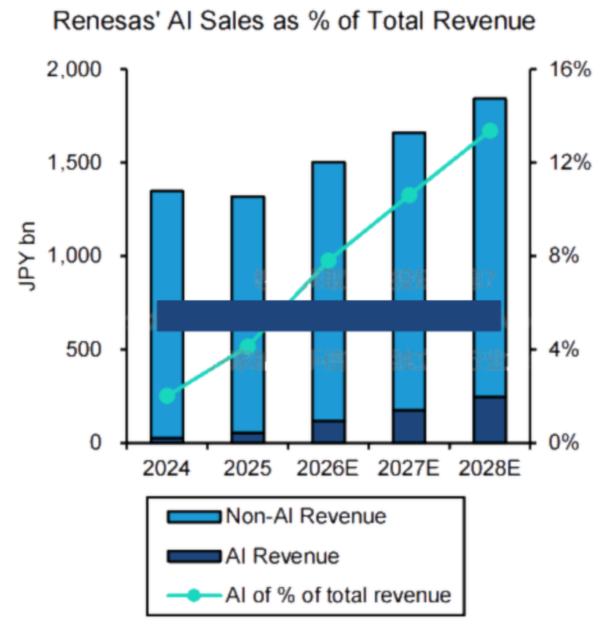

Source: Company disclosures, Bernstein estimates and analysis.

- Texas instruments currently provides mostly discretes for PSU (AC/DC) instead of power management IC / modules. Overall they have relatively low current exposure (with overall datacenter at ~9% of revenue) though it is growing quickly (up >60% YoY in 2025) (**Exhibit 26**, **Exhibit 27**)

EXHIBIT 26: **Overall datacenter (including both AI and non-AI applications) made up ~8% of TXN's revenues in 2025 (~$1.5B total)**

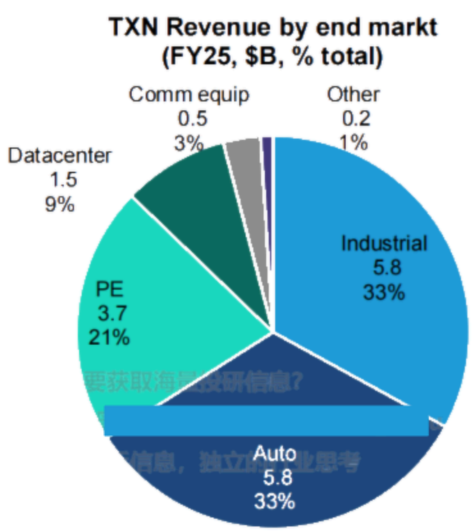

Source: Company reports, Bernstein analysis

EXHIBIT 27: **However, their datacenter segment is fast-growing, up 64% YoY last year**

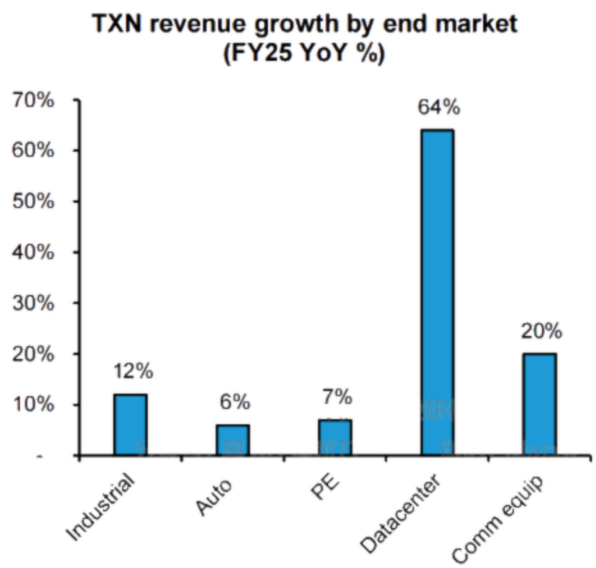

Source: Company reports, Bernstein analysis

- Analog Devices had ~$2B in overall datacenter business in 2025 though split across ATE, optical, and power. Power (including delivery and control, hot-swap, and high-performance protections solutions) represented ~30% of this (~$600M;

Exhibit 28), with the overall datacenter segment growing ~50% last year. Other datacenter offerings include optical, as well as exposure through ATE (automated test equipment) used to test datacenter semiconductors.

EXHIBIT 28: **In FY25 ADI's total datacenter exposure hit ~$2B (up ~50% YoY); power accounted for ~$600M**

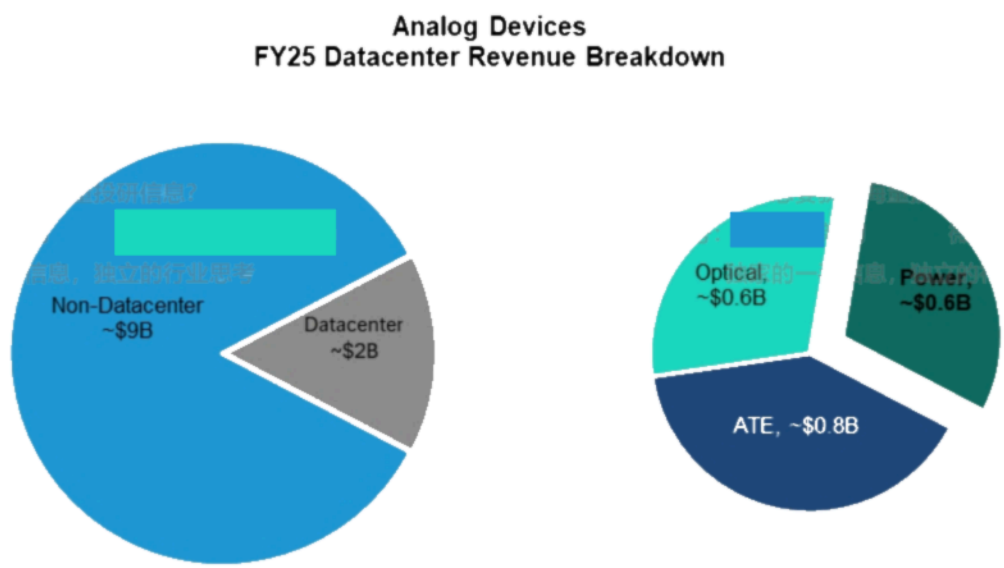

Source: Company reports, Bernstein analysis