Morgan Stanley | RESEARCH

ASIA PACIFIC INSIGHT

October 6, 2026 12:44 AM GMT

Greater China Technology Hardware

Integrated Power Solutions Driving Secular Demand

Increasing power consumption, along with the deployment of BBU/CBU/800/±400VDC/SST in AI clusters, could push the technology barrier higher and present growth opportunities for players with superior capabilities. Fuel cells could also emerge as a new opportunity amid power constraints.

Secular demand for power amid AI capex: AI-related power consumption will continue to increase, and we believe this will lead to secular demand for power-related components, such as power supply units (PSUs), battery backup units (BBUs), capacitor bank units (CBUs), coolant distribution units (CDUs), and so on. Moreover, the requirement for efficiency will inevitably necessitate the adoption of 800/±400 voltage direct current (VDC) in the power rack/center/block configuration (initially from late 2026 and broadly in 2028, we estimate), eventually extending to the deployment of solid state transformers (SST) in the longer term – all pointing to a more centralized power delivery architecture in the future. This trend of aggregation could raise the technology barriers amid the fusion of white/grey space, favoring players with better capabilities and execution. Overall, we forecast the datacenter PSU market size to grow at a 75% three-year CAGR to reach US$33bn in 2028.

Fuel cells to improve flexibility and resilience of AI power delivery: Despite the strong AI capex outlook (MSe global cloud capex +97%/20% YoY in 2026/2027), power has become a bottleneck, indicated by MS's forecast power shortfall of 5GW in 2026, scaling to a total 72GW through 2029 (see report). Fuel cells, among other alternative energy sources, could emerge as a good solution to pick up some slack and enhance efficiency through active micro grid management. Those who are able to offer comprehensive solutions from power sources to power delivery could be increasingly more competitive in the market.

Assuming coverage of Delta and Lite-On:

- Delta (OW, PT NT$2,500, based on 30x 2027/2028 P/E) should maintain its dominant position in the industry with its integrated power/thermal related solutions while fuel cells could add incremental growth opportunities in the longer term. We think a 30x P/E target multiple is justified by our estimated 63% earnings CAGR in 2025-28, driven by a 44% revenue CAGR and operating margin expansion to 22.1% in 2028 vs 15.1% in 2025,
- Lite-On (EW, PT NT$315, based on 17x 2027/2028 P/E), the second-largest player in AI power-related products, should also ride the tailwinds despite its slower progress in HVDC power racks and liquid cooling solutions. We think a 17x P/E target multiple is justified by our estimated 47% earnings CAGR in 2025-28 (29% revenue CAGR and 6.8ppt OPM expansion to 16.8% in 2028).

Derrick Yang
Equity Analyst
Derrick.Yang@morganstanley.com +886 2 2730-2862

Kazuo Yoshikawa, CFA
Equity Analyst
Kazuo.Yoshikawa@morganstanleymufg.com +81 3 6836-8408

Andy Meng, CFA
Equity Analyst
Andy.Meng@morganstanley.com +852 2239-7689

Vivi Huang
Research Associate
Vivi.Huang@morganstanley.com +886 2 2730-2860

Hidetaka Suzuki
Research Associate
Hidetaka.Suzuki@morganstanleymufg.com +81 3 6836-8402

GREATER CHINA TECHNOLOGY HARDWARE

Asia Pacific
Industry View **In-Line**

WHAT'S CHANGED

Delta Electronics Inc. (2308.TW)	**From**	**To**
Price Target	NT$2,700.00	**NT$2,500.00**
Lite-On Technology (2301.TW)	**From**	**To**
Price Target	NT$176.00	**NT$315.00**

investors should be aware that the firm may have a conflict of

Research as only a single factor in making their investment decision.

For analyst certification and other important disclosures, refer to the Disclosure Section, located at the end of this report.

+= Analysts employed by non-U.S. affiliates are not registered with FINRA, may not be associated persons of the member and may not be subject to FINRA restrictions on communications with a subject company, public appearances and trading securities held by a research analyst account.

The Story in Charts

Exhibit 1: AI servers continuing to see higher TDP...

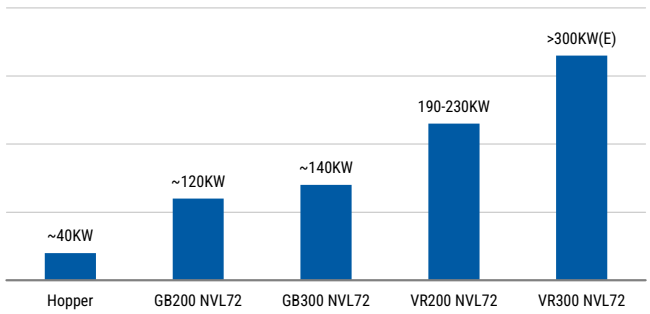


Exhibit 2: ...which drives data center power supply demand

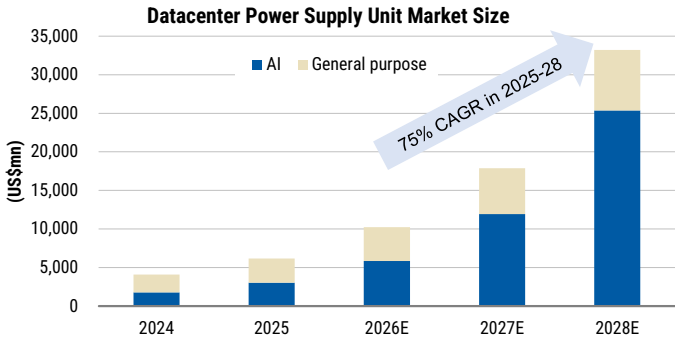
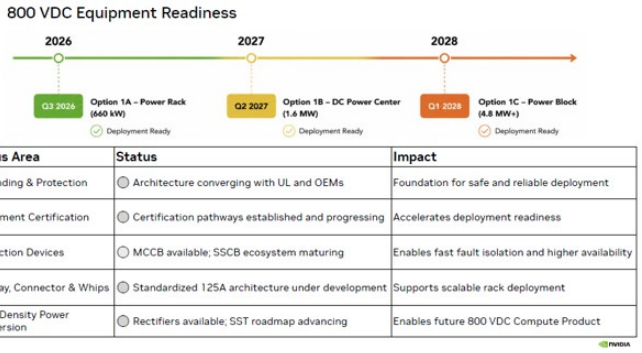


Exhibit 3: HVDC to see early adoption from late 2026...



Focus Area	Status	Impact
Grounding & Protection	Architecture converging with UL and OEMs	Foundation for safe and reliable deployment
Equipment Certification	Certification pathways established and progressing	Accelerates deployment readiness
Protection Devices	MCCB available; SSCB ecosystem maturing	Enables fast fault isolation and higher availability
Busway, Connector & Whips	Standardized 125A architecture under development	Supports scalable rack deployment
High-Density Power Conversion	Rectifiers available; SST roadmap advancing	Enables future 800 VDC Compute Product

Source: Nvidia

Exhibit 4: ...with a much higher dollar content

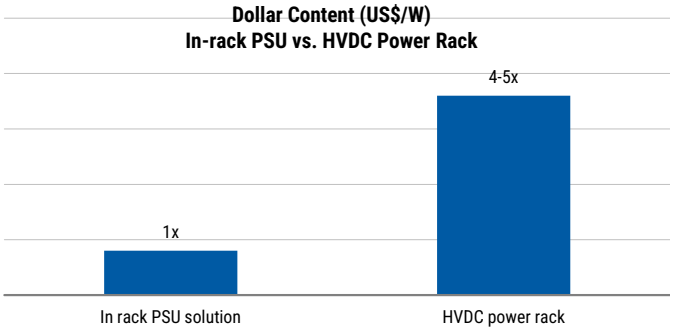
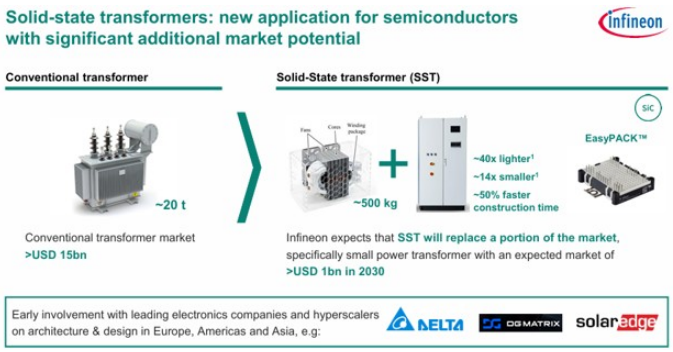
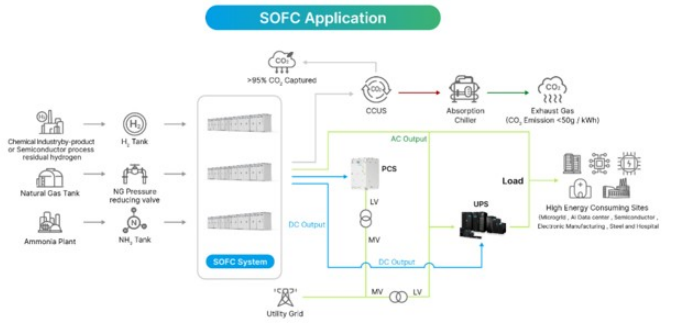


Exhibit 5: SST is also in the pipeline for more centralized power delivery



Source: Infineon

Exhibit 6: Fuel cells are a good option for better resilience in power architecture



Source: Delta

Executive Summary

AI cluster performance evolving to fulfill demand

The rapid adoption of artificial intelligence is driving a significant increase in global data center compute demand. Advances in foundation models, reasoning capabilities, agentic AI, and enterprise AI deployment are increasing the computational resources required for both training and inference. To support this growth, hyperscalers and model developers continue to deploy larger and more powerful AI clusters, driving demand for greater compute, memory, and networking capacity.

As AI infrastructure scales, cluster performance is increasingly determined by the efficiency of the entire system rather than the capabilities of individual processors alone. Compute, networking, cooling, and power infrastructure must all evolve together to support higher rack densities and larger deployments. While semiconductor innovation remains critical, supporting infrastructure is becoming an equally important factor in enabling future AI growth.

AI cluster performance is fundamentally a power problem

As AI clusters continue to scale in size and performance, power consumption is becoming one of the most important constraints on future infrastructure deployment. The power challenge is no longer limited to GPUs alone, as modern AI clusters comprise tightly integrated compute, networking, memory, storage, and cooling systems, all of which require increasing amounts of power as performance requirements rise. Consequently, the next phase of AI infrastructure investment should be viewed not only as a semiconductor upgrade cycle but also as a broad-based power infrastructure upgrade cycle.

Server power consumption has increased significantly across successive AI hardware generations. Nvidia's architecture roadmap illustrates the industry's migration from relatively modest server-level power budgets under the Ampere and Hopper generations toward rack-scale systems under Blackwell and future Rubin architectures. Estimated rack-level power consumption has increased from roughly 40kW for Hopper systems to approximately 120kW for GB200 NVL72 and 140kW for GB300 NVL72, with future Vera Rubin systems expected to approach 200-300kW per rack with current Vera Rubin systems reaching up to 330kW per rack while future Vera Rubin generations are expected to surpass these levels. As compute density continues to increase, power delivery infrastructure is becoming a critical enabler of system performance.

The increase in power consumption is not limited to compute hardware. Networking systems are also becoming increasingly power intensive as switch bandwidth scales from 100G architectures toward 800G and 1.6T generations. Despite substantial process-node improvements, switch ASIC power consumption has continued to rise due to increasing complexity and higher SerDes speeds. As a result, networking equipment now requires significantly larger power budgets than previous generations, creating additional demand for higher-capacity power supplies throughout the data center.

Exhibit 7: TDP roadmap for Nvidia racks

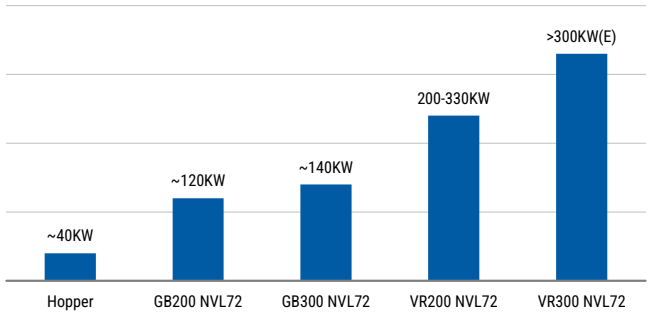
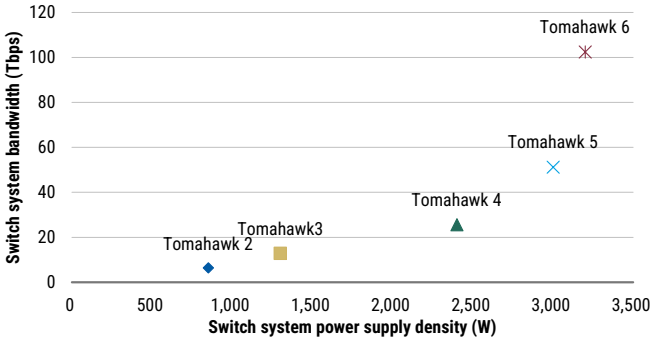


Exhibit 8: Power supply density for Tomahawks

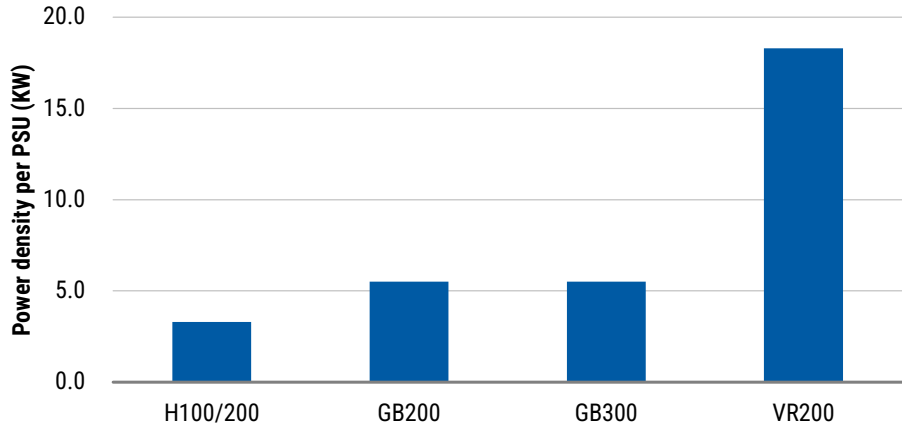


The growth in power consumption is driving a corresponding upgrade cycle for power-supply units (PSUs). Earlier AI server generations typically relied on multiple 3-3.3KW power supplies installed directly within individual server chassis. As the industry transitions toward rack-scale architectures, power conversion is increasingly centralized into dedicated power shelves capable of supporting substantially larger power loads. In Blackwell systems, power shelves utilize 5.5KW power supplies, while future Rubin-generation platforms are expected to increase PSU ratings to approximately 18KW, representing a significant increase in both PSU content and power density.

Delivering higher power within a similar footprint creates significant engineering challenges. Increasing PSU power density requires advances across multiple areas, including higher-performance power semiconductors such as silicon carbide (SiC) and gallium nitride (GaN), more efficient power-conversion topologies, smaller and more capable magnetic components, improved thermal-management solutions, and greater system integration. As output power increases, thermal constraints become increasingly important, making heat dissipation, airflow optimization, and mechanical design critical determinants of PSU performance.

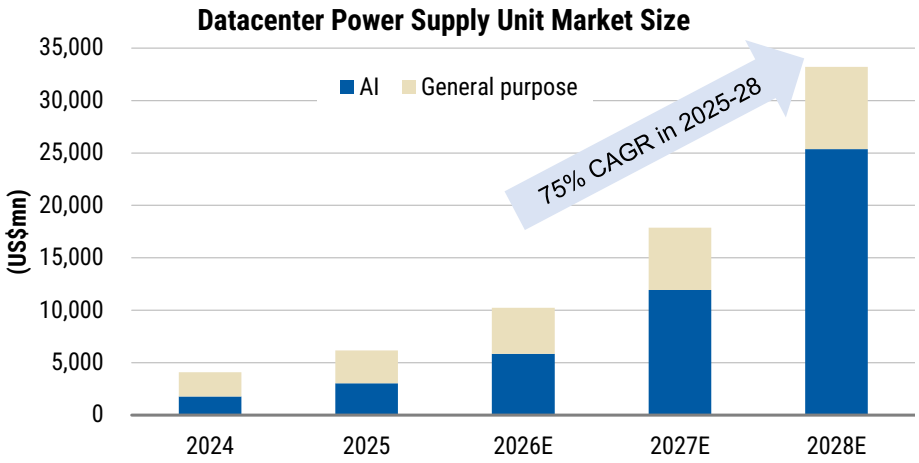
Consequently, on top of the volume driven growth, the power-supply industry is benefiting from the technology driven dynamics, where efficiency, power density, and thermal performance are becoming key competitive differentiators. As AI rack power requirements continue to increase, PSU innovation will become an increasingly important component of the broader AI infrastructure investment cycle.

Exhibit 9: Power density per PSU trend in Nvidia server systems



Overall, we estimate the global datacenter power supply market will grow at a three-year CAGR of 75% to reach US$33bn in 2028 from US$6bn in 2025, driven mainly by server shipment growth, increasing power consumption per rack, and the adoption of HVDC power racks from 2027.

Exhibit 10: Global datacenter power supply market forecast



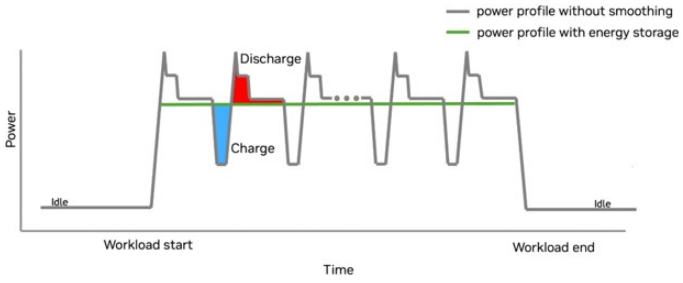
Complementary hardware for more resilient power delivery

As AI clusters become more power intensive, maintaining stable and resilient power delivery is emerging as a critical design priority. Traditional centralized power architectures are increasingly being complemented by localized energy-storage and power-conditioning technologies that can respond rapidly to dynamic workload demands.

Battery Backup Units (BBUs) represent one of the most important additions to modern AI power architectures. Positioned close to the computing workload, BBUs provide localized energy storage capable of responding significantly faster than facility-level backup systems. Their role extends beyond conventional backup power functionality. BBUs can

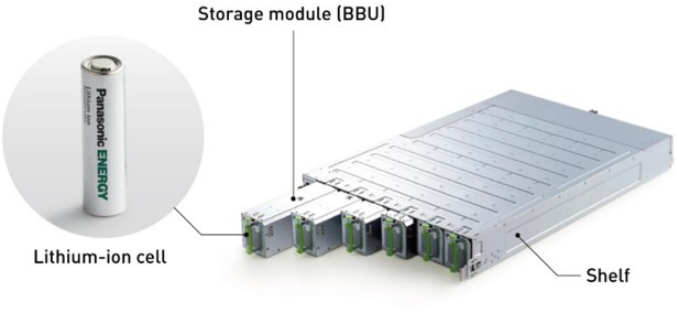
absorb transient fluctuations, smooth power ramps, and improve overall rack-level power stability. By providing local buffering capacity, BBUs help reduce the burden on upstream power distribution infrastructure while improving overall system resilience.

Exhibit 11: Concept of peak shaving for AI server power loads



Source: Nvidia

Exhibit 12: BBU



Source: Panasonic

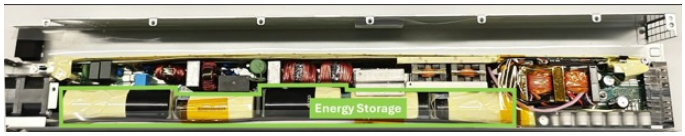
Capacitor Backup Units (CBUs) address an even shorter-duration segment of the power stability challenge. Utilizing supercapacitors and other high-power-density storage technologies, CBUs can deliver and absorb power within milliseconds, making them particularly effective for managing rapid GPU power spikes. These systems act as shock absorbers within the electrical architecture, preventing sudden load changes from propagating throughout the broader power network. As AI rack power densities continue increasing, capacitor-based energy storage is becoming an attractive solution for improving power quality without requiring substantial oversizing of facility infrastructure.

Exhibit 13: CBU



Source: Panasonic

Exhibit 14: GB300 power shelf with integrated capacitors



Source: Nvidia

Together, BBUs and CBUs provide complementary benefits. BBUs are designed to handle longer-duration events, while CBUs are optimized for extremely fast transient response. The integration of these technologies enables operators to improve power quality, reduce stranded infrastructure capacity, enhance system reliability, and increase effective utilization of expensive electrical assets. As rack densities continue to rise, localized energy-storage solutions are likely to become a standard component of AI data center design rather than an optional enhancement.

The adoption of BBU and CBU architectures may also enable future increases in rack density. As operators attempt to push toward several hundred kilowatts per rack, localized power-buffering solutions can help mitigate transient peaks and reduce stress on grid-facing infrastructure. This allows operators to deploy higher-density computing resources without proportionately increasing upstream electrical capacity.

Exhibit 15: Comparison of CBU and BBU

Dimension	CBU	BBU
Primary storage medium	EDLC supercapacitors, LICs or high-capacitance	Typically lithium-ion battery cells managed by a BMS
Primary strength	High power, fast charge/discharge and frequent	Higher stored energy and longer sustained
Main function	Transient smoothing, peak shaving, voltage	Outage backup, source-transfer bridging and orderly workload
Response profile	Well suited to millisecond-class load support and	Fast switchover, but generally optimized for sustained support rather
Typical duration	Milliseconds to seconds; selected systems can extend toward tens of seconds or more	Commonly tens of seconds to several minutes, depending on battery and rack design
Power versus energy density	Higher power capability but lower energy storage per unit volume than a battery	Higher energy density, making it more economical for longer events
Cycle life	Eaton specifies more than one million cycles and up to 20 years for certain supercapacitor modules	Battery life is more sensitive to depth of discharge, temperature and repeated cycling; BMS supervision is therefore essential.
State management	Requires cell balancing, voltage/temperature monitoring and state-of-health tracking, but does not use a conventional battery fuel	Requires BMS functions including charge control, state of charge, state of health and cell protection
Best-fit use case	Repetitive AI power spikes, inrush events, PSU load leveling and bus stabilization	Power continuity through outages, brownouts or generator/source-transfer intervals

HVDC emerging as a necessity, with SST in the pipeline

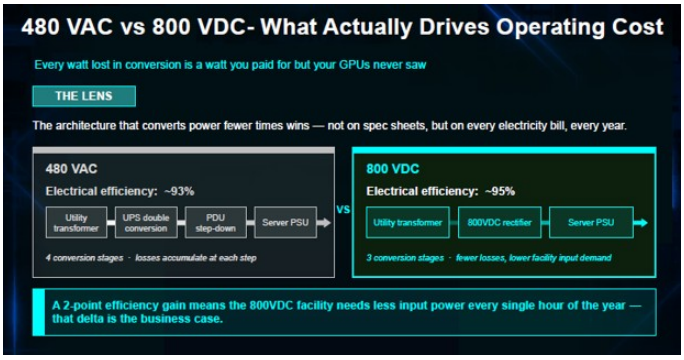
800VDC and ±400VDC progressing simultaneously

The rapid increase in AI rack power density is creating growing challenges for conventional power-distribution architectures. As rack-level power consumption rises from tens of kilowatts to several hundreds of KWs, electrical losses, cable size, thermal management requirements, and overall system complexity increase significantly. Under traditional low-voltage power architectures, delivering larger amounts of power requires proportionally higher current, which increases resistive losses and places greater stress on cables, connectors, busbars, and power-conversion equipment.

At the same time, every power-conversion stage introduces efficiency losses and generates additional heat. As AI deployments scale into multi-megawatt and eventually gigawatt-level campuses, even small efficiency improvements can translate into meaningful reductions in operating costs, cooling requirements, and infrastructure investment. Consequently, the industry is increasingly focused on power architectures that can reduce conversion stages, lower current levels, and improve overall power-delivery efficiency.

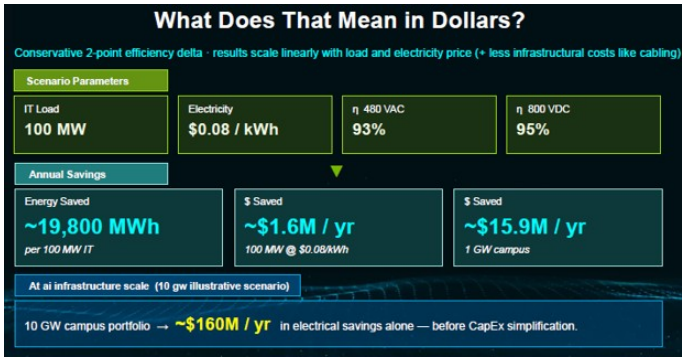
Against this backdrop, high-voltage direct current (HVDC) is emerging as one of the most important architectural upgrades for next-generation AI infrastructure. By increasing distribution voltage, HVDC allows the same amount of power to be delivered at a lower current, reducing conduction losses and improving system efficiency. Higher-voltage architectures also improve scalability for ultra-high-density rack deployments by alleviating some of the physical and thermal constraints associated with conventional power-distribution systems.

Exhibit 16: Better power efficiency for 800VDC architecture



Source: Delta

Exhibit 17: Better efficiency leads to cost savings



Source: Delta

Nvidia's roadmap illustrates the industry's gradual migration toward higher-voltage power architectures. Earlier AI server generations primarily utilized conventional power-delivery schemes centered around rack-level 48V distribution. As rack power requirements increased with Blackwell-generation systems, industry participants began evaluating higher-voltage architectures capable of supporting substantially greater power density. Looking ahead, the industry is increasingly converging around 800VDC as a potential next-generation standard for large-scale AI deployments, particularly as rack-level power consumption approaches several hundred kilowatts. The transition reflects the need to improve power-distribution efficiency while supporting the continued growth of AI computing infrastructure.

We believe that the adoption of the 800VDC power rack could start from late 4Q26, though it should be small scope deployment through 2027, since the in-rack power shelf is the reference design in the Vera Rubin NVL72 rack. Broader adoption should happen more likely from late 2027 and 2028, as the Vera Rubin Ultra rack further pushes up the power consumption.

Exhibit 18: HVDC deployment roadmap



Focus Area	Status	Impact
Grounding & Protection	Architecture converging with UL and OEMs	Foundation for safe and reliable deployment
Equipment Certification	Certification pathways established and progressing	Accelerates deployment readiness
Protection Devices	MCCB available; SSCB ecosystem maturing	Enables fast fault isolation and higher availability
Busway, Connector & Whips	Standardized 125A architecture under development	Supports scalable rack deployment
High-Density Power Conversion	Rectifiers available; SST roadmap advancing	Enables future 800 VDC Compute Product

Source: Nvidia

The 800VDC power rack BOM is primarily driven by the PSU, which accounts for 50–55% of total cost, followed by the BBU at 25–30%, with other components comprising the remaining 15–20%. This cost structure reflects the broader functionality integrated into an HVDC power rack, combining power conversion, backup power and supporting components within a centralized architecture.

For power supply players, one of the most direct implications is that an HVDC power rack carries around 4–5x higher dollar content per watt compared with a conventional in-rack PSU solution, since the former integrates more components into a robust system and power supply players usually ship the entire rack instead of just the components.

Exhibit 19: BOM cost breakdown for an 800VDC power rack

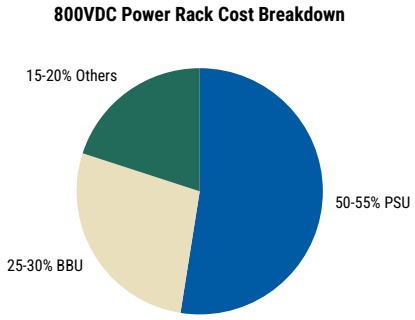
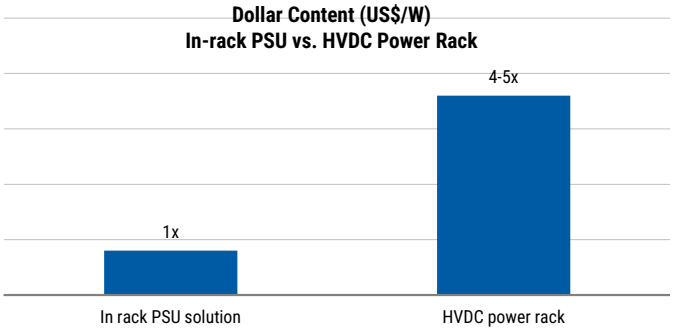


Exhibit 20: Higher dollar content for HVDC power racks



However, the industry has not yet fully converged on a single architecture. While Nvidia's roadmap increasingly points toward 800VDC, several hyperscalers and members of the Open Compute Project (OCP) ecosystem have expressed interest in ±400VDC architectures. Both approaches deliver similar total voltage levels and pursue many of the same objectives, including lower current, higher efficiency, and improved scalability. The debate largely centers on implementation trade-offs rather than end goals.

Supporters of 800VDC generally highlight the advantages of a simpler single-bus architecture, potentially lower system complexity, and closer alignment with existing industrial high-voltage DC equipment. The architecture avoids the need to maintain balanced positive and negative voltage rails and may simplify certain aspects of power conversion and system integration. These characteristics make 800VDC particularly attractive in environments where operational simplicity and interoperability with existing electrical infrastructure are key priorities.

Advocates of ±400VDC, meanwhile, emphasize the advantages of a bipolar architecture. By creating positive and negative voltage rails around a neutral reference point, ±400VDC can potentially reduce insulation requirements, simplify certain fault-management scenarios, and provide improved flexibility for load balancing and future system expansion. Many hyperscalers also view the bipolar approach as being more compatible with long-term visions for highly distributed power systems incorporating energy storage, renewable energy, and microgrid functionality.

Importantly, the ultimate beneficiaries may be largely independent of which standard gains wider adoption. With the exception of certain power-semiconductor and protection-system requirements, many component suppliers across the power ecosystem remain

broadly agnostic between 800VDC and ±400VDC implementations. Higher-voltage architectures in either form require similar upgrades in power conversion, power management, cooling, connectors, magnetics, and infrastructure equipment. As a result, the industry's migration toward HVDC matters significantly more than the eventual selection of any single voltage standard.

Exhibit 21: Comparison of 800VDC and ±400VDC

Metric	±400VDC	800VDC
Voltage to ground	400V relative to ground	800V relative to ground
Voltage differential	800V	800V
Power distribution efficiency	High	Slightly higher
Scalability	Good	Excellent
Insulation requirements	Lower	Higher
Maintenance safety	Better	More challenging
System architecture simplicity	Dual-rail system	Single-bus system
Grounding Complexity	More complicated	Simpler
Electrical footprint	Small	Smaller
Ability to reduce conversion stages	Moderate	Highest
Long-distance distribution capability	Excellent	Excellent
Maturity of ecosystem	Better today	Early but rapidly developing
Industry players	OCP, hyperscalers	Nvidia
Long term industry direction	Transitional architecture	Likely end-state architecture

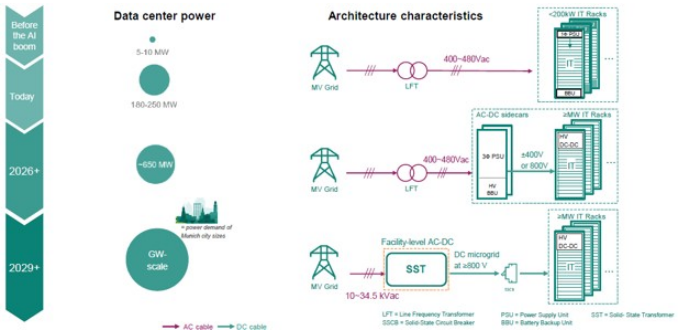
Solid-state transformers to follow after 800VDC/±400VDC matures

Looking further ahead, solid-state transformers (SSTs) represent a potential next-generation evolution of power infrastructure. Unlike traditional transformers that operate primarily as passive devices, SSTs leverage power electronics to provide voltage conversion, power conditioning, and bidirectional energy management within a single platform. This enables greater controllability and flexibility while supporting increasingly digitalized power networks.

Although SSTs remain at an early stage of commercialization and face challenges related to economics, reliability, and large-scale deployment, the technology aligns well with long-term data center trends. As AI campuses evolve into utility-scale power consumers incorporating energy storage, distributed generation, and advanced power management, SSTs could eventually become a key enabling technology.

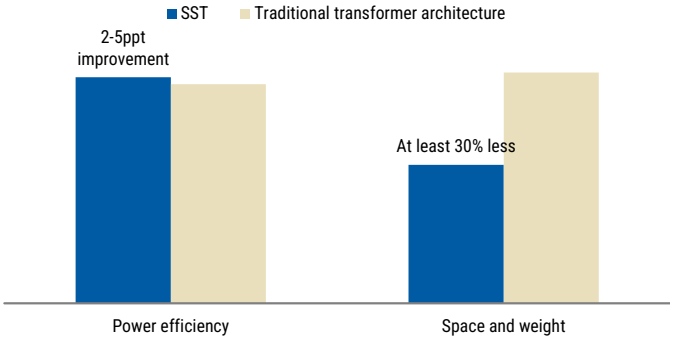
Therefore, while HVDC is expected to drive the next phase of power architecture upgrades, SSTs represent a longer-term technology roadmap that could further transform how future data centers distribute and manage electricity.

Exhibit 22: SST enabling more efficient power delivery architecture



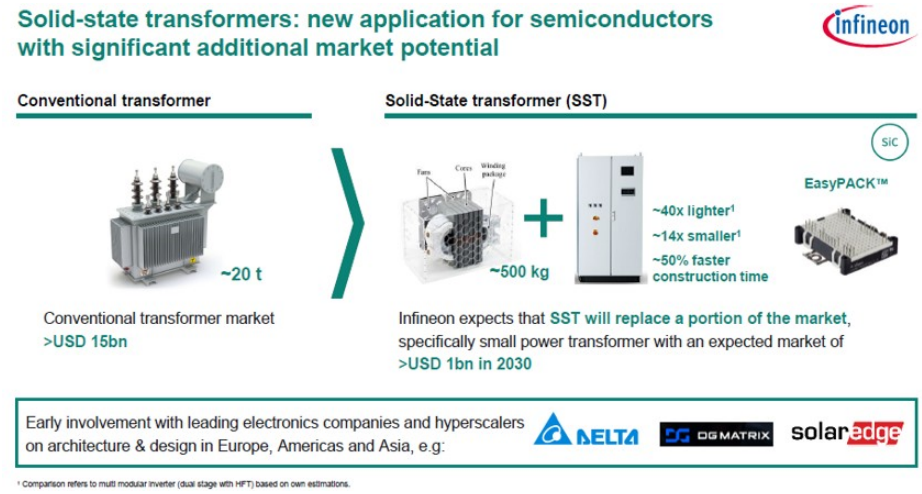
Source: Infineon

Exhibit 23: SST offers better power efficiency with at least 30% less space/weight



According to Infineon, SST could represent an emerging >US$1bn market opportunity by 2030, as the technology begins to replace a portion of the >US$15bn conventional transformer market, particularly in smaller-power applications. While SST penetration is expected to remain relatively modest versus the overall transformer market, the opportunity is significant for relevant players in the longer term.

Exhibit 24: SST to see initial adoption in 2030

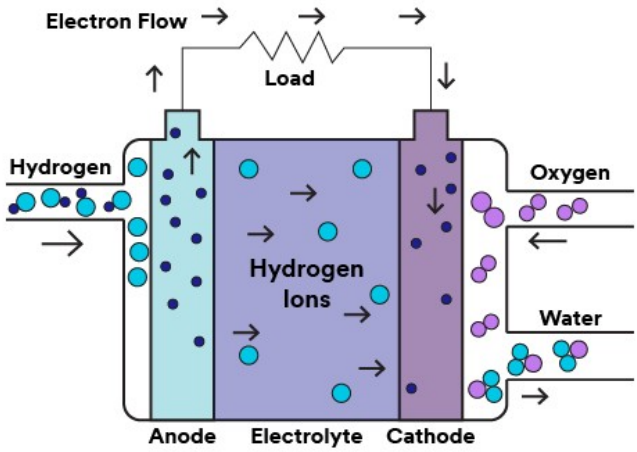


Source: Infineon

Fuel cells for better flexibility and resilience

Fuel cells generate electricity through an electrochemical reaction rather than combustion. Hydrogen, natural gas, biogas, or other fuels are supplied to the fuel cell stack, where chemical energy is converted directly into electricity, heat, and in some cases water. Unlike batteries, which store energy, fuel cells continuously generate power as long as fuel is supplied.

Exhibit 25: Fuel cell diagram



Source: 3M

Fuel cells can offer several benefits, including:

- **High energy efficiency:** Direct electrochemical conversion enables higher efficiency than conventional combustion-based power generation, reducing fuel consumption and operating costs.
- **Low carbon emissions**: Fuel cells produce minimal direct emissions and can support long-term decarbonization goals when powered by low-carbon fuels such as hydrogen.
- **Grid independence and energy security:** Provides reliable onsite power generation, reducing dependence on grid availability and improving energy resilience.
- **Long-duration energy storage enabler:** Hydrogen-based fuel-cell systems can complement renewable energy by storing energy over extended periods, supporting long-duration and seasonal energy storage.
- **Fuel flexibility:** Certain fuel-cell technologies can operate on hydrogen, natural gas, biogas, ammonia-derived hydrogen, and other hydrogen-rich fuels, allowing deployment using existing fuel infrastructure.
- **Low maintenance and quiet operation:** Limited moving parts reduce maintenance requirements while providing quieter operation than conventional combustion-based generators.
- **Scalable deployment:** Can be deployed across a wide range of power levels, from distributed backup systems to utility-scale installations, making fuel cells suitable for large AI data center campuses

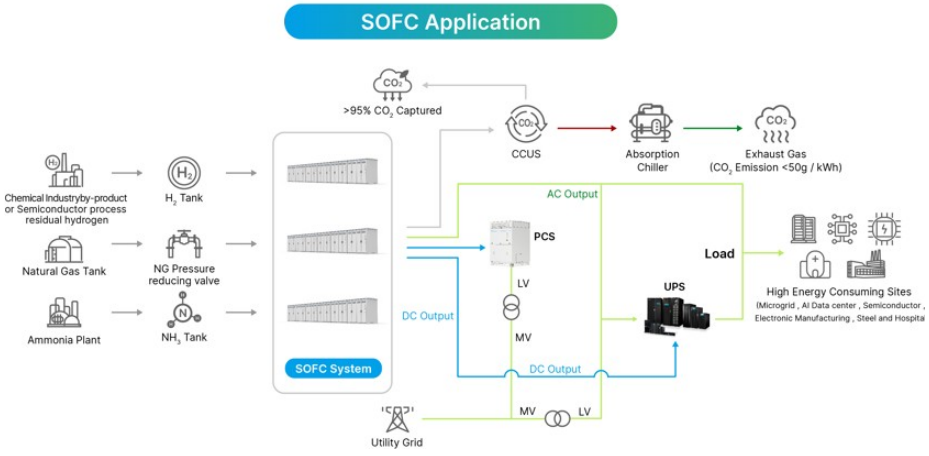
However, the following challenging are hurdles for wider deployment:

- **High system cost:** Fuel cells remain more expensive than conventional power-generation technologies, limiting large-scale adoption.
- **Limited hydrogen infrastructure:** Hydrogen production, transportation, storage, and refueling networks remain underdeveloped in most regions.
- **Durability requirements:** Long-term reliability and stack lifetime remain key challenges, particularly for large-scale and mission-critical deployments.
- **Regulatory and safety considerations:** Hydrogen handling requires specialized infrastructure and compliance with stringent safety regulations, increasing deployment complexity.
- **Lack of standardization:** Multiple fuel-cell technologies compete across different applications, creating uncertainty around long-term technology standards and adoption paths.

Fuel cells have the potential to fundamentally reshape future data center power architectures. Traditionally, data centers rely on grid electricity as the primary power source and diesel generators for backup power. However, as AI-driven workloads significantly increase power demand and power availability becomes a key deployment constraint, hyperscalers are increasingly evaluating fuel cells not only as a backup-power solution but also as a potential primary power source. This could accelerate the transition from a grid-centric architecture toward one incorporating distributed onsite generation, enabling data center deployment in regions where utility power capacity is limited.

Over the longer term, fuel cells could become a core component of AI data center microgrids that integrate renewable generation, battery storage, hydrogen storage, and intelligent energy-management systems. Such an architecture would improve scalability, resilience, and energy independence while reducing reliance on centralized utility infrastructure. In addition, because fuel cells inherently generate DC power, they align well with the industry's migration toward HVDC-based power distribution architectures. This could simplify power delivery by reducing conversion stages between generation and IT equipment, improving overall system efficiency and supporting next-generation high-density AI deployment

Exhibit 26: Fuel cell in datacenter power delivery architecture



Source: Delta

Stock picks

Delta (2308.TW, OW, PT NT$2,500)

Dominant supplier in the AI power supply offerings: Delta has been one of the most competitive players in power related components and systems, as its offerings spanning from the core power supply units, the complementary BBU/CBU for AI servers, the highly anticipated HVDC power rack, grey space power facilities to air/cooling solutions. We estimate that it has ~60% market share in the datacenter power supply market and is among the leading suppliers of the 800VDC/±400VDC power rack to Nvidia/hyperscalers in 2H26 (wider adoption likely from late 2027). Longer term, when the HVDC becomes more mature, the centralized power delivery architecture could migrate into SSTs, where Delta already has small volume shipment now. With the solid technology, close customer relationships and strong execution, we believe Delta should stay competitive even amid the fusion of the white space and grey space.

Extending from power delivery to power source: In addition to the offerings in the power delivery, Delta signed a licensing agreement with Ceres Power (CRW.GB) for its solid oxide fuel cell (SOFC) technology in early 2024. This could extend its business scope into the power source area. which we believe could add more value to its customers as it could help mitigate the power constraints for datacenter build and enhance the operational efficiency through active micro grid management.

Solid near term momentum: Despite the slower revenue growth owing to the component supply tightness in 3Q26, we expects its 4Q26 revenue to grow by double digit % QoQ on alleviated component supply and solid demand. 3Q26 GM should be flat to slightly up QoQ and 4Q26 might see further expansion on the tailwinds price hikes.

OW with PT of NT$2,500: We see power/thermal as one of the major secular AI related themes in the downstream tech space and Delta is well positioned to benefit, potentially driving a 63% earnings CAGR in 2025-28, in our estimate. We thus assume the coverage of Delta with an OW rating and the PT of NT$2,500, based on 30x 2027/2028 P/E. Our target multiple implies a PEG of 0.5 vs. the average PEG of 1.2 in 2020-25, when the stock was trading at an average 12M forward P/E of 22x with the earnings CAGR of 19%.

Lite-On (2301.TW, EW, PT NT$315)

Rising tide lifting all boats: Lite-On is the second largest power supply vendor globally and has been engaging with hyperscalers regarding the ASIC and Nvidia GPU servers. Looking ahead, it expects to expand its share in a couple of hyperscalers from late 2026. It should have over 50% of the total revenue generated from the cloud segment for PSUs, BBUs, and to a lesser extent, other mechanical components like liquid cooling, chassis/rack, etc. Moreover, it also has the HVDC power rack solutions ready and is qualifying with potential customers. The cloud shall remain as the major growth driver in the coming years.

Aspiration in the optical communication space: Lite-On has always had an opto business (10-20% of total revenue in 2026) in both visible and invisible light applications. Post the strategic investment at Dense Light (21% stake), an InP epi wafer and chip process manufacturer, it plans to leverage its expertise in packaging and provides the optical

modules and engine for the data interconnects. While the end applications looks promising, we think it might take some time to see more meaningful revenue contribution in Lite-On.

Cloud momentum extending into the coming quarters: We expect its 4Q26 revenue to grow by double digit QoQ driven mainly by the could segment and GM should normalize at ~25% in 2H26 vs. 27.2% in 2Q26 on some inventory write-off reversal and high margin product shipment delayed from the previous quarter.

EW with PT of NT$315: We think Lite-On will grow along with the AI capex cycle, driven by its power-related offerings (47% earnings CAGR in 2025-28, in our estimate), though execution and progress with new customers is worth monitoring. We assume coverage of Lite-On with an EW rating and a PT of NT$315, based on 17x 2027/2028 P/E. Our target multiple implies a PEG of 0.4 vs. the average PEG of 1.6 in 2020-25, when the stock was trading at an average 12M forward P/E of 14x on an earnings CAGR of 9%.

Panasonic Holdings (EW, PT ¥4,700, covered by Kazuo Yoshikawa)

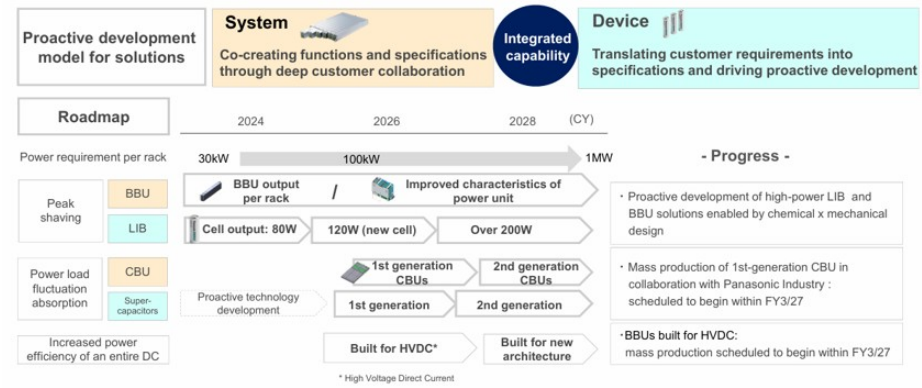
Panasonic has an approximately 80% share of the BBU market and is well positioned to benefit from rising AI data center power demand: The company expects data center energy storage revenue to grow from ¥322bn in FY3/26 to approximately ¥550bn in FY3/27 and around ¥1trn in FY3/29. Importantly, all of the business underpinning its FY3/29 revenue target has already been awarded, meaning agreements have been reached with customers on product development and orders.

BBU: Panasonic has begun limited mass production of new high-capacity BBU cells above 200W, with products adopted by a major GPU vendor and ramping from June-July. Discussions for next-generation server platforms are in the final stages, and the company expects to maintain its approximately 80% market share. It also sees scope for higher pricing as rack-level power consumption rises.

The company is also close to securing awards for BBUs used in dedicated power racks. As its medium-term plan includes only already-awarded business, additional design wins could provide upside.

CBU: CBUs are emerging as a new growth opportunity as GPU and ASIC power consumption increases. Panasonic is working with power-supply vendors to propose integrated BBU+CBU solutions and is receiving inquiries from US and Chinese hyperscalers, including customers that were not previously BBU users.

Exhibit 27: Panasonic Holdings: Strengthen proposal and development capabilities for the next generation



Source: Company data

Risks to our positive view

Risks to our positive view on power-related demand include:

- Component supply tightness leading to lower shipment of products
- Shortfall in hardware shipments, yield issues, electricity availability, etc. slowing down datacenter builds
- Lower-than-expected penetration of new power delivery architectures like 800VDC/±400VDC, SST, etc.
- More intense competition leading to margin erosion among most players
- Hyperscalers prove unable to finance aggressive capex through internal and external funding.

AI Cluster Performance Evolving to Fulfill Demand

AI adoption driving exponential growth in compute demand

AI is becoming a mainstream computing workload

Over the past three years, generative AI has evolved from a research application into a core technology platform. Consumer applications such as ChatGPT, Claude, Gemini, and Copilot now serve hundreds of millions of users globally, while enterprises are increasingly embedding AI into software development, customer service, document analysis, enterprise search, healthcare, financial research, and industrial automation.

Historically, software demand scaled roughly with the number of users. AI demand is fundamentally different because every user interaction consumes computational resources in real time. Unlike traditional cloud applications, where the bulk of computation occurs during software development, AI applications continuously consume compute whenever a model generates a response. Consequently, rising AI adoption directly translates into higher demand for data center compute resources.

Model complexity continues to increase

The AI industry has historically improved model performance through scaling laws: larger models trained on larger datasets using more compute generally produce better outcomes. While efficiency improvements continue, leading AI developers remain committed to scaling model capabilities.

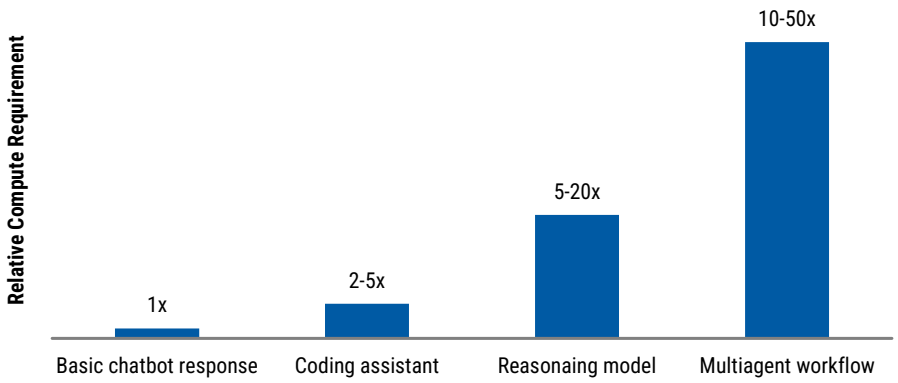
More importantly, the industry is moving beyond conventional language models toward significantly more compute-intensive architectures.

Reasoning models

Traditional LLMs generate answers directly from a prompt. New reasoning models perform multiple internal computational steps before producing a response. Instead of generating a few hundred output tokens, advanced reasoning systems may generate thousands or even tens of thousands of internal reasoning tokens while evaluating different solution paths. This dramatically increases inference workloads.

The industry is increasingly shifting from "parameter scaling" toward "test-time compute scaling," where additional compute is consumed during inference to improve answer quality. From an infrastructure perspective, this is highly significant because inference workloads occur continuously rather than once during model training.

Exhibit 28: Increasing demand for compute power from reasoning models



Agentic AI

The next wave of AI applications involves autonomous agents capable of planning, executing, verifying, and refining tasks.

Instead of generating a single answer, an AI agent may search multiple information sources, run code, validate results, generate follow-up tasks, and produce a final response. Each step requires additional inference cycles.

A workflow that previously required one model invocation may now require dozens of model calls, substantially increasing compute consumption per user.

Multimodal AI

AI models are increasingly processing images, video, audio, sensor data, and text simultaneously.

Video understanding is particularly compute intensive because each second of video contains vastly more information than a text prompt.

As multimodal AI adoption grows, compute demand per query rises significantly, creating another structural driver of accelerator demand.

Training demand remains robust

Although inference is becoming the largest source of AI compute demand, training workloads continue to expand. Frontier model development increasingly requires clusters containing tens of thousands of accelerators operating simultaneously.

Training runs are becoming larger because developers are pursuing larger parameter counts, more training data, longer training durations, multimodal datasets and reinforcement learning and reasoning optimization

Consequently, each new generation of foundation models typically consumes substantially more compute than the previous generation.

Importantly, training infrastructure sets the baseline for future inference deployment. Every major model developer must maintain access to large-scale training clusters to remain competitive, supporting continued demand for high-end AI accelerators.

More compute requests to be fulfilled by better chip performance

Rising compute demand cannot be addressed simply by adding more servers. Modern AI workloads increasingly encounter bottlenecks related to memory bandwidth, interconnect performance, power consumption, and cooling. As a result, each new generation of AI hardware delivers improvements across multiple dimensions.

Compute performance

The latest AI accelerators provide several-fold performance improvements versus prior generations. For example, NVIDIA's Blackwell architecture was specifically designed to support both large-scale training and high-volume inference, with significant gains in AI throughput compared with previous Hopper-based systems. Performance improvements are necessary because reasoning and agentic workloads require substantially more operations per query.

Memory capacity and bandwidth

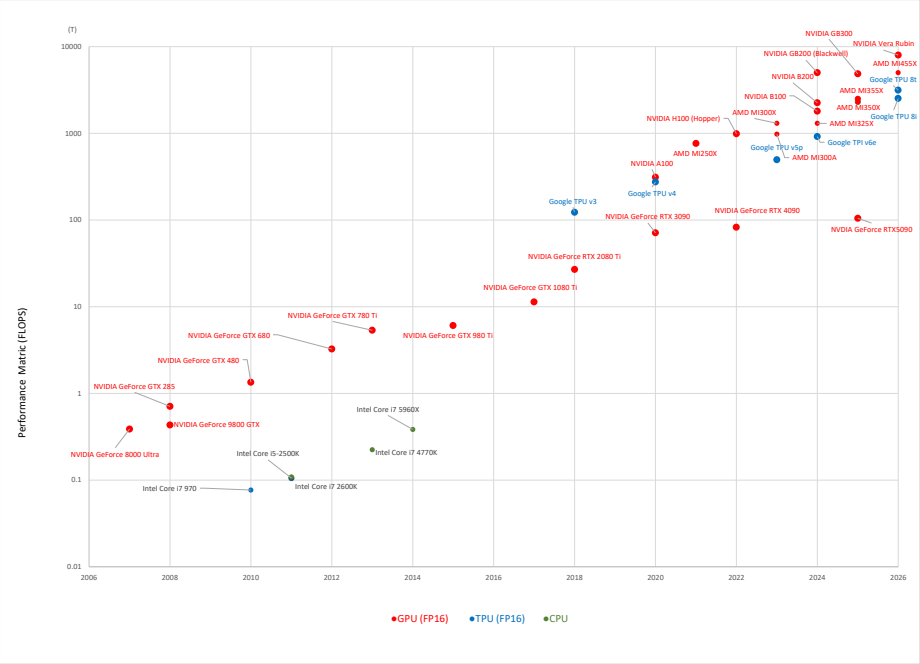
AI models are increasingly constrained by memory rather than raw compute, because longer models, longer context windows, reasoning models and agentic workflows all require more memory. As model complexity increases, modern high-end accelerators usually incorporate hundreds of gigabytes of High Bandwidth Memory (HBM), providing memory bandwidth measured in multiple terabytes per second.

Interconnect performance

As AI clusters scale from hundreds to thousands of accelerators, communication between processors becomes increasingly important. A large AI training cluster effectively behaves like a single massive computer. Consequently, networking technologies such as NVLink, InfiniBand, Ethernet, etc., have become critical enablers of AI scaling.

Many AI workloads can no longer be accelerated simply by increasing GPU performance; cluster interconnect bandwidth must also improve.

Exhibit 29: Compute power for HPCs continuing to evolve



AI Cluster Performance Is Fundamentally a Power Problem

As AI clusters become larger and more powerful, every component within the cluster must deliver higher throughput, lower latency, and greater bandwidth. Achieving these improvements inevitably requires more electrical power.

Therefore, the power challenge is no longer limited to GPUs. Modern AI clusters are highly integrated systems comprising compute servers, networking fabrics, memory subsystems, storage infrastructure, and cooling equipment. As cluster performance increases, power consumption rises across the entire infrastructure stack.

In our view, the next phase of AI infrastructure investment should be viewed not simply as a semiconductor upgrade cycle, but as a broad-based power consumption cycle driven by the need for ever-higher AI cluster performance.

Server systems seeing increasing TDP

Nvidia's AI-server power roadmap is best understood as a transition from chip-defined, air-cooled nodes to power-, cooling- and fabric-constrained rack-scale computers: the Ampere-era A100 SXM carried a disclosed 400W standard TDP—up to 500W in custom thermal designs—with 40/80GB HBM2/HBM2e and 600GB/s NVLink, while the eight-GPU, dual-AMD-EPYC DGX A100 was a conventional 6U server with a 6.5KW maximum input rating, meaning GPU silicon represented only roughly half of the node envelope once CPUs, NVSwitches, memory, NICs, drives, fans and conversion losses were included.

Hopper stepped the H100 SXM to a disclosed 700W TDP, 80GB HBM3, roughly 3TB/s memory bandwidth and 900GB/s fourth-generation NVLink, while H200 retained essentially the same Hopper compute architecture and power profile but substituted 141GB of faster HBM3e at 4.8TB/s—making H200 primarily a memory-capacity/bandwidth upgrade rather than a compute-TDP step-up—and Nvidia rates the eight-GPU, dual-Xeon DGX H100/H200 at 10.2KW maximum input, versus 6.5KW for DGX A100;

The architectural discontinuity arrives with GB200. Rather than treating an eight-GPU server as the scale-up unit and using InfiniBand/Ethernet primarily between nodes, Grace Blackwell NVL72 makes the entire rack the scale-up domain, integrating 18 one-U compute trays—each with two 72-core Grace Arm CPUs and four Blackwell GPUs—plus nine NVLink switch trays, management switches, , a 50–51V DC busbar, copper NVLink cartridges and liquid manifolds, so that 36 CPUs and 72 GPUs operate as one low-latency NVLink-connected system. Based on the same compute die, Nvidia upgraded the HBM content to further improve the performance of GB300 chip. Overall, the rack level power consumption is estimated at ~120KW for the GB200 NVL72 and ~140KW for the GB300 NVL72.

Vera Rubin—often called VR200/R200 in industry roadmaps—pushes the same rack-level philosophy further. Nvidia's disclosed architecture combines 72 dual-die Rubin GPUs with 36 new Vera CPUs, replacing Grace's 72 Arm Neoverse cores with 88 custom Olympus cores, raising CPU memory capacity and bandwidth to as much as 1.5TB LPDDR5X and 1.2TB/s, and doubling coherent CPU–GPU NVLink-C2C bandwidth to 1.8TB/s; Rubin itself moves from HBM3e to 288GB HBM4 at up to 22TB/s, introduces a third-

generation Transformer Engine and raises sixth-generation NVLink scale-up bandwidth to 3.6TB/s per GPU, while ConnectX-9 provides 1.6Tbps/GPU scale-out connectivity and BlueField-4 increasingly offloads storage, security and infrastructure processing.

Nvidia has not published a simple Rubin chip TDP in its primary datasheet, but our supply chain checks indicate roughly 1.8kW at an efficiency-oriented operating point and up to approximately 2.3kW at the performance mode, which is consistently corroborated by system vendors, with ASUS specifying 187kW Max-Q and 227kW Max-P for its fully liquid-cooled Vera Rubin NVL72.

Looking ahead, we expect this trend to extend into the future along with the technology advancement and more critical workloads.

Exhibit 30: Compute power vs. TDP for HPCs

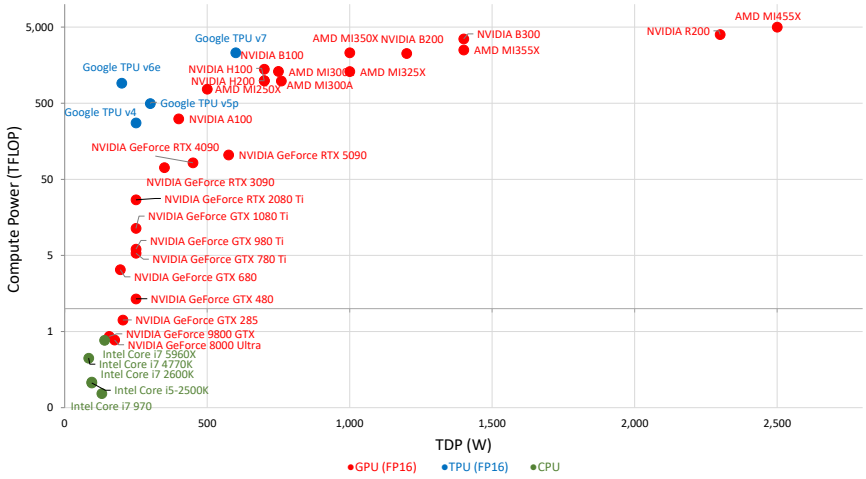
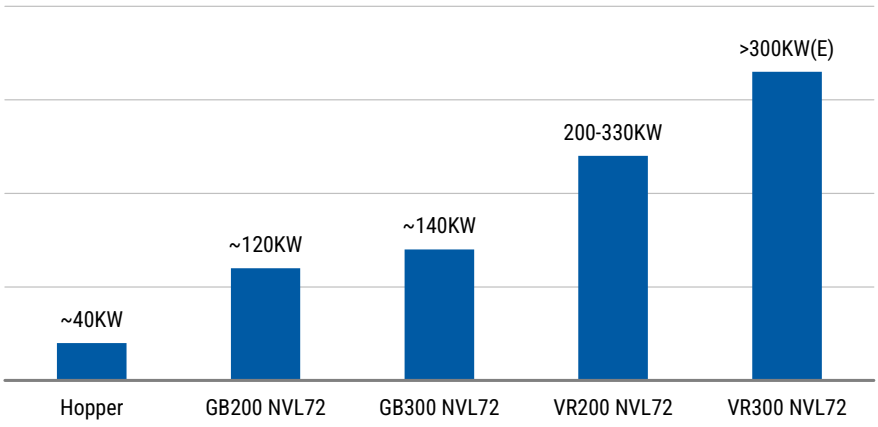


Exhibit 31: TDP roadmap for Nvidia racks



Networking consuming more power as well

The network switch system is another key devices within the datacenter and has been evolving over the past years, driven primarily by advances in the switch ASIC. From Tomahawk 1 (2014) to Tomahawk 6 (2025), the process node shrank from 28nm to 3nm, while aggregate switching bandwidth increased from 3.2 Tbps to 102.4 Tbps, representing a 32× increase. Over the same period, SerDes speed improved from 25 Gbps NRZ to 200 Gbps PAM4, enabling maximum port speeds to scale from 100G to 1.6T.

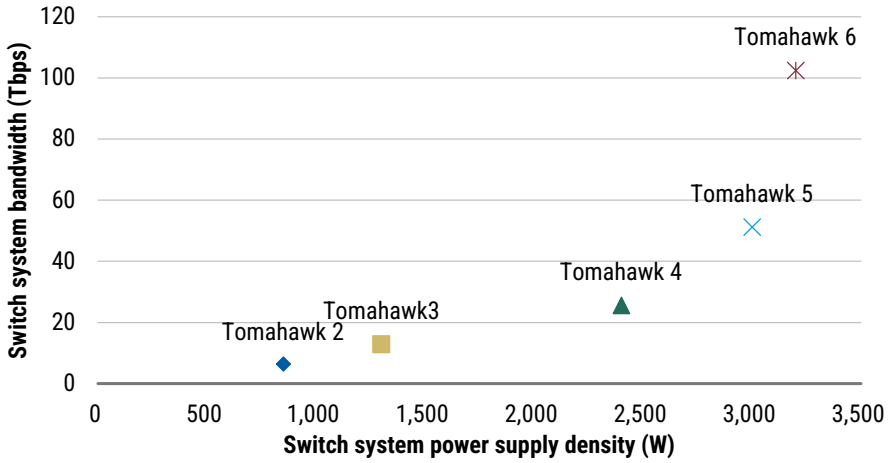
While process scaling significantly improved switching efficiency, it did not prevent ASIC power from rising. ASIC TDP increased from 150 W to 800 W, reflecting the growing complexity of higher radix switches and faster SerDes interfaces.

The impact of ASIC scaling is evident at the system level. As switch bandwidth increased from 100G-class systems to 800G and 1.6T platforms, power delivery requirements grew dramatically. Representative PSU ratings increased from 850W in Tomahawk 2 systems to 3,200W in Tomahawk 6 systems, nearly a 4× increase.

The rise in PSU power is driven not only by the higher ASIC TDP, but also by the growing power consumption of high-speed optics, increased cooling requirements, and larger overall system power budgets. As systems evolved from 100G to 1.6T ports, switch vendors were forced to add more fan redundancy, larger heatsinks, and in some cases larger chassis footprints. For example, Tomahawk 6 systems have expanded to 4U form factors with 7+1 fan redundancy to accommodate the thermal load associated with 102.4 Tbps switching capacity.

Alongside the increasing power consumption of the server racks, network switches are also witnessing higher power budgets as well.

Exhibit 32: Power supply density for different network switch generations



Power supply growing on both volume and specs

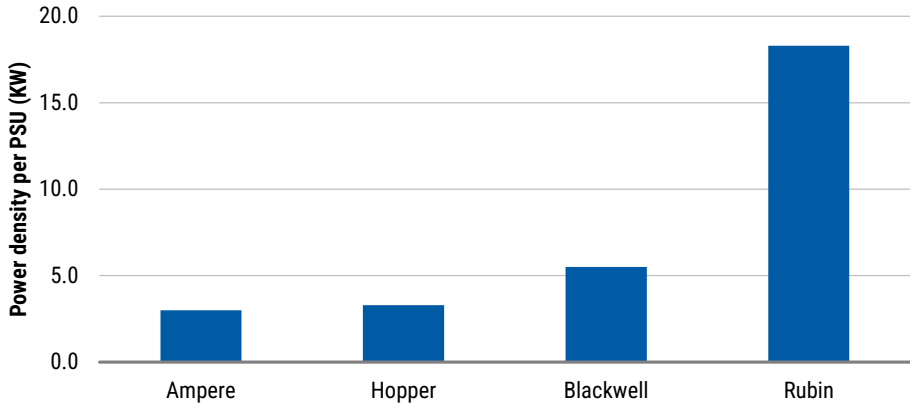
PSU spec upgrade trend

In response to the increasing power consumption of the AI servers, the power supply architecture has been migrating as well. in the A100 and H100/200 single system servers, 6 power supply units are placed directly inside the chassis, with each capable to deliver 3.0-3.3KW of power.

When Nvidia make a broad based move to the rack level server solution from the Blackwell generation, the power supply units were no longer placed directly within the chassis housing the HPCs that consume power; instead, power supply units are separated and grouped into power shelves. In the GB200/300 systems, there are 8 power shelves sitting on the same rack as the compute trays and switch trays, each enclosing six 5.5KW power supply units.

For the upcoming Vera Rubin rack, the power supply unit will be upgraded to 18.3KW of output to comprise a power shelf of 110KW, and each power shelf will take 3U of space on the rack vs. 1U in the Blackwell generation.

Exhibit 33: Power density for PSUs in Nvidia server systems



Challenges increasing the power density of PSUs

To come up with power supply unit with higher power density, manufacturers might need to solve the following potential hurdles:

Better power semiconductors (SiC and GaN)

The biggest improvement in PSU power density comes from replacing traditional silicon power chips with newer materials such as silicon carbide (SiC) and gallium nitride (GaN). These devices waste less energy during power conversion and can operate at much higher speeds. As a result, many supporting components inside the PSU—such as transformers and inductors—can be made significantly smaller. This allows PSU manufacturers to fit more output power into the same physical size while also improving efficiency. In many ways, SiC and GaN have become the foundation for the latest generation of high-density AI server power supplies.

More efficient power conversion designs

Modern PSUs use more sophisticated circuit designs, i.e., bridgeless PFC vs. conventional one, that reduce energy losses during power conversion. A conventional boost PFC includes a diode bridge and therefore incurs multiple semiconductor conduction drops. A bridgeless totem-pole PFC reduces the number of devices in the active current path and is now a preferred architecture for dense server PSUs. Instead of simply converting electricity from one voltage to another, newer architectures are designed to minimize wasted energy throughout the process. Higher efficiency means less heat is generated, which directly enables higher power density. This is particularly important in AI servers, where even small efficiency improvements can translate into meaningful reductions in heat generation and cooling requirements.

Smaller and smarter magnetic components

Transformers and inductors are often among the largest components inside a PSU. A major area of innovation is reducing their size without sacrificing performance. Manufacturers are increasingly using planar magnetics, where copper traces on circuit boards replace traditional wire windings. These designs are thinner, easier to cool, and more space-efficient. Some suppliers are also combining multiple magnetic functions into a single component, further reducing overall volume. As PSU power levels increase from 3 KW to 5.5 KW and beyond, shrinking magnetic components becomes one of the most important contributors to power-density improvement.

Improved thermal management

As power output increases, heat becomes the primary obstacle to higher density. Even highly efficient PSUs still generate a significant amount of heat. Therefore, manufacturers are investing heavily in better cooling solutions, including optimized airflow paths, advanced heatsinks, vapor chambers, and liquid-cooling technologies. The goal is to remove heat more effectively so that components can operate safely even when packed more tightly together. In many modern PSU designs, thermal engineering is just as important as electrical engineering in determining achievable power density.

Higher component integration

Another way to increase power density is to reduce the amount of unused space inside the PSU. Manufacturers are integrating more functions into fewer components and using three-dimensional layouts that make better use of available volume. Components are stacked, arranged vertically, or combined into integrated modules to maximize space utilization. This approach is similar to how smartphones have become more powerful without becoming larger—the same principle is increasingly being applied to server power supplies.

Higher efficiency creates more thermal headroom

Efficiency improvements provide a double benefit. First, less electricity is wasted. Second, less heat is generated. For example, a 5.5 kW PSU operating at 97.5% efficiency generates roughly 30% less heat than one operating at 96.5% efficiency. While a one-percentage-point improvement may appear small, it can remove tens of watts of heat from the system. This allows engineers to either increase output power within the same form factor or reduce cooling requirements while maintaining existing power levels.

Reducing the size of energy storage components

Server PSUs contain capacitors that store energy and maintain power delivery during short interruptions. Historically, these capacitors occupied a significant amount of space. Newer designs use smarter control techniques that allow manufacturers to achieve the same performance with smaller capacitor banks. Reducing capacitor volume frees up valuable internal space that can be used for higher power output or additional functionality.

For example, in Infineon's 3.3 KW example, the calculated capacitance for a conventional 10ms hold-up implementation was approximately 2.5 mF; allowing deeper discharge through a baby-boost stage reduced the requirement to about 900 µF. The 12 KW design similarly uses a bidirectional buck-boost buffer to achieve 20ms hold-up at full load without the capacitor volume that a passive solution would require. The penalty is additional semiconductors, an inductor, control software and failure modes, but the net volumetric benefit can be substantial.

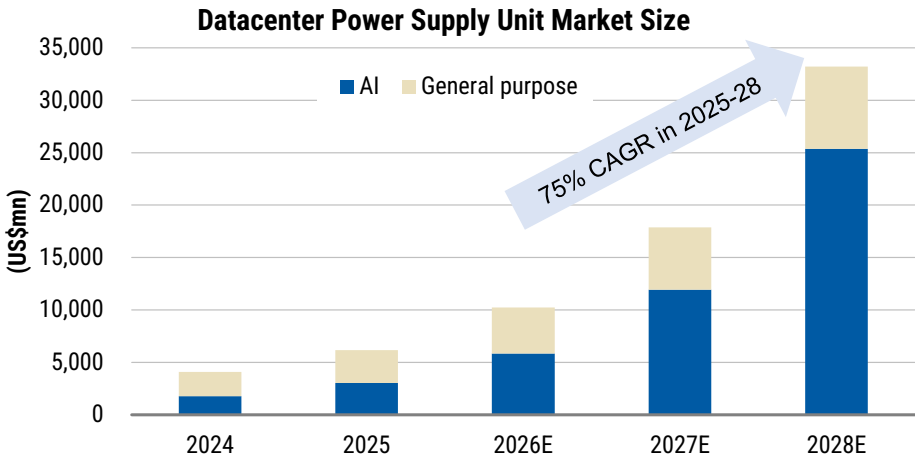
Advanced mechanical design

Power density is ultimately measured in watts per unit volume, making physical packaging extremely important. Manufacturers continuously optimize component placement, airflow channels, connector design, and enclosure structures to eliminate unused space. The best PSU designs are often not defined by a single breakthrough component but rather by hundreds of small mechanical optimizations that collectively improve volumetric efficiency.

Datacenter PSU market size to grow at a 75% CAGR in 2025-28

Overall, we estimate the global datacenter power supply market will grow at a three-year CAGR of 75% to reach US$33bn in 2028 from US$6bn in 2025, driven mainly by server shipment growth, increasing power consumption per rack, and the adoption of HVDC power racks from 2027.

Exhibit 34: Global datacenter power supply market forecast



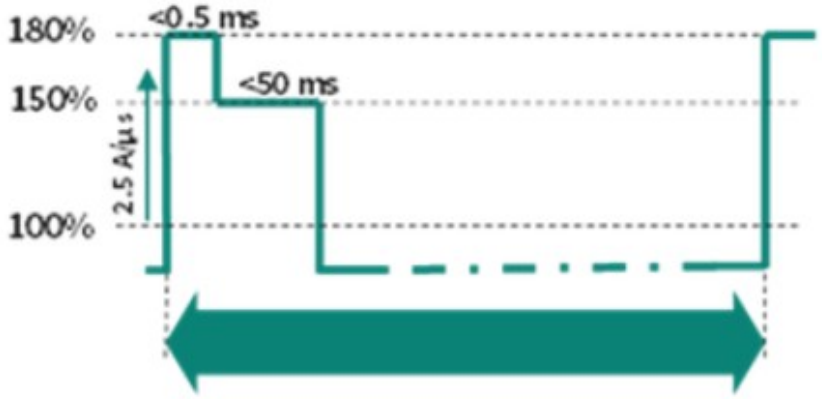
Complementary Hardware for More Resilient Power Delivery

Transient power spikes at AI servers

AI servers have a much more dynamic power profile than conventional enterprise servers because GPU utilization changes sharply throughout a workload. During large-model training, thousands of GPUs typically operate in a bulk-synchronous manner, and power consumption approaches the accelerator's design limit during tensor-compute phases, then falls during collective communications, check pointing, data movement, or periods in which GPUs wait for slower devices to synchronize. Inference can create a similar pattern as request arrivals, batching, context length and token-generation activity change rapidly.

Because the GPUs perform these transitions together, the power variation of one accelerator is multiplied across the server, rack and cluster rather than being averaged out by unrelated workloads, so the magnitude of the challenge is also rising with system density. Power spikes are therefore not simply the result of higher average consumption; they reflect the combination of higher accelerator power, more GPUs per system and increasingly synchronized workload behavior.

Exhibit 35: Transient power spikes at AI servers



Source: Infineon

System performance and reliability potentially compromised

The key issue is that power delivery systems are designed to supply a relatively stable amount of power, while AI workloads can change power demand almost instantaneously. When the load changes faster than the power-delivery system can react, a temporary mismatch occurs between power demanded by the GPUs and power supplied by the PSU, creating electrical instability throughout the server.

Sudden increases in GPU demand lead to voltage drops

When GPUs transition from a low-utilization state (communication, synchronization, checkpointing, etc.) to a full-compute state, their current demand can increase within microseconds. However, the PSU cannot instantly increase its output because its control loop requires time to detect the load change and adjust power conversion.

During this response period, the additional energy must come from capacitors located throughout the power chain. As these capacitors discharge, the voltage on the DC bus begins to fall, creating voltage drop.

If the drop becomes sufficiently large, several operational issues can occur:

- GPU throttling: The GPU reduces clock speed or performance to remain within safe operating limits.
- Computation errors: Voltage instability can affect timing margins in high-speed processors and memory.
- Unexpected resets or shutdowns: Protection circuits may trigger if voltage falls below specified thresholds.
- Training interruption: Large AI training jobs running across thousands of GPUs may fail if even a small number of nodes experience instability.

In a large AI cluster, a single server failure can force the entire training job to restart from the latest checkpoint, resulting in significant productivity loss.

Sudden decreases in GPU demand results in voltage spike

The opposite problem occurs when GPUs rapidly reduce power consumption.

For example, after a compute-intensive phase, GPUs may enter a communication phase where power consumption drops sharply. The PSU, however, is still delivering power based on the previous load condition. The excess energy has nowhere to go immediately and accumulates in the DC bus, causing a temporary voltage spike.

This can create:

- Overvoltage stress on components
- Protection-circuit activation
- Accelerated degradation of capacitors and voltage regulators
- Higher risk of component failure over time

Although individual spikes may be small, AI servers can experience these transitions continuously during training, causing repeated stress throughout the power-delivery chain.

Excessive current surges throughout the system

Power spikes are not only a voltage problem—they also create large current excursions.

A 1KW increase in power demand at 12V requires roughly 83A of additional current. In modern AI servers, where multiple GPUs can change utilization simultaneously, current swings can reach hundreds of amperes within a very short period.

These current surges lead to:

- Higher conductor losses
- Increased connector heating
- Greater stress on voltage-regulator modules (VRMs)
- Larger thermal fluctuations in power components
- As GPU power ratings rise to over 1KW, the magnitude of these transient currents becomes increasingly difficult to manage.

Thermal cycling and reliability degradation

Perhaps the most overlooked issue is the impact on long-term reliability.

Every power excursion creates a corresponding temperature excursion. When power rises, semiconductor junction temperatures increase; when power falls, temperatures decrease. This repeated heating and cooling cycle occurs thousands or even millions of times over a system's lifetime.

The result is accelerated wear on:

- Electrolytic capacitors
- Power MOSFETs
- Solder joints
- PCB interconnects
- Connectors and busbars

Consequently, even if average power remains within specification, transient power spikes can shorten component lifetime and increase failure rates.

Infrastructure overdesign and lower efficiency

To prevent instability, operators often provision the power infrastructure for the worst-case peak rather than average consumption.

For example:

- A rack averaging 80KW may need infrastructure sized for 120-140KW transient peaks.
- PSUs must be oversized to accommodate short-duration excursions.
- UPS systems must reserve capacity for sudden load increases.
- Cooling systems must handle peak thermal output.

This creates a significant amount of stranded capacity, where expensive electrical and cooling infrastructure remains underutilized most of the time. Furthermore, PSUs and UPS systems typically operate most efficiently near their optimal loading point; oversized systems therefore tend to operate at lower efficiency, increasing operating costs.

Conventional power delivery systems not designed to deal with power volatility

In the conventional power delivery architecture, the UPS protects the data hall against utility outages and power-quality disturbances, while the PSU converts upstream power into the tightly regulated DC output required by the server. Both remain essential in the modern AI server clusters, but neither was originally optimized to absorb large, repetitive and synchronized AI load steps economically.

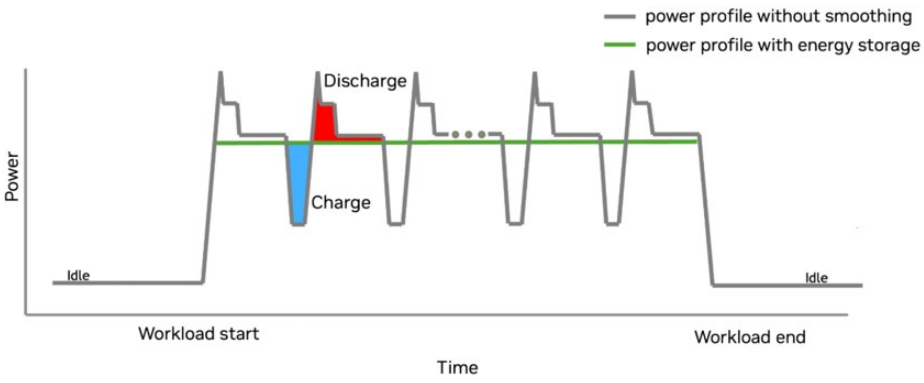
The PSU can use its internal capacitors to cover the initial transient, but capacitor volume is limited; once that stored energy is depleted, the PSU control loop must raise input current, potentially producing current overshoot and transferring the disturbance upstream.

The UPS is located farther from the accelerator and sees the aggregated load of many PSUs. A transient originating near the GPU can therefore occur on microsecond-to-millisecond timescales—including permitted GPU electrical-design-power excursions around 50 ms—before the complete upstream conversion chain can rebalance power.

This is a response time and location mismatch, rather than evidence that UPSs or PSUs are fundamentally incapable of supporting AI. New AI-tolerant UPS platforms are being designed for faster and larger load steps, but many legacy systems may still experience voltage excursions, nuisance alarms, lower partial-load efficiency or unwanted bypass transfers under steep synchronized loads.

To complement the conventional power delivery architecture, more AI server clusters are adopting BBUs or CBUs to shave peak loads and mitigate the above-mentioned issues.

Exhibit 36: Concept of peak shaving for AI server power loads



Source: Nvidia

Battery backup units seeing increasing importance

What's a BBU?

A BBU is an integrated battery, power-conversion and control system that supplies DC power to server loads when the normal power source is interrupted or temporarily insufficient. In an OCP-style rack, the BBU modules are housed in a dedicated shelf located near the PSU power shelf and connect directly to the common DC busbar. Under normal conditions, the PSUs power the servers and maintain the BBU's charge; when the bus

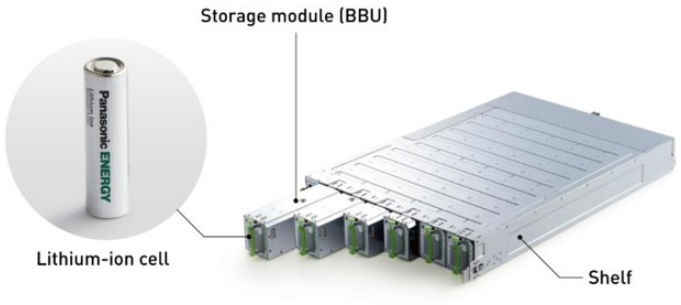
voltage falls below a programmed threshold, the BBU's bidirectional converter changes from charging or standby into discharge mode and supports the rack. Analog Devices' ORV3 reference design detects a bus voltage below 48.5V for more than 2ms and is required to ramp to full support within another 2ms, while maintaining the bus above the specified minimum voltage.

Key components within a BBU

- **Battery-cell pack:** The cells store the energy used during discharge. Lithium-ion cylindrical cells are commonly arranged in series and parallel to obtain the required voltage, power and runtime; the ORV3 reference module uses an interchangeable 11-series, 6-parallel battery pack. The cell must provide not only sufficient energy but also high discharge power, low internal resistance and predictable performance over its qualified life.
- **Bidirectional DC/DC converter:** The converter controls energy flow between the battery pack and the rack DC bus. It steps voltage in the appropriate direction during charging and discharging and regulates the output so that the battery's changing voltage does not destabilize the server bus. Multiphase designs divide the large current across several parallel power stages; TI's 12 KW design uses six 2 KW phases and reports peak efficiency above 99%, including above 99.5% when battery and bus voltages are close.
- **Battery management system:** The BMS measures individual cell voltages, current and temperature, balances cells and estimates state of charge and state of health. It also protects against overcharge, excessive discharge, overheating and short circuit. Accurate state-of-health information is important because the operator must know whether an apparently charged BBU can still deliver its required rack power and runtime after years of standby operation and periodic discharge events.
- **Power-path and protection circuitry:** Fuses, current sensors, disconnect switches, protection MOSFETs and hot-swap circuits isolate faults and prevent uncontrolled current flow. Typical protections include overvoltage, undervoltage, overcurrent and overtemperature shutdown. Hot-swap capability allows a failed or end-of-life module to be replaced without intentionally shutting down the entire rack, while active current sharing ensures that parallel BBU modules carry an appropriate portion of the load.
- **Control, monitoring and communications:** A module controller manages operating modes, converter control, fault handling, diagnostics and communication with the BMS and rack controller. Interfaces such as PMBus, Modbus, CAN, RS-485 or Ethernet can expose battery voltage, current, temperature, state of charge, state of health, fan speed and fault status to data-center management systems. Analog Devices' reference platform uses separate controllers for battery monitoring and overall BBU logic, communications and housekeeping.
- **Thermal-management system:** Fans, temperature sensors, airflow paths and mechanical spacing around cells remove heat during high-rate charging and discharging. Thermal management is both a performance and a safety requirement: excessive temperature accelerates battery ageing and can restrict available discharge power. Analog Devices' ORV3 design therefore incorporates fan control, airflow simulation and cell-to-cell spacing, while Murata's commercial 3 KW BBU includes internal forced-air cooling and overtemperature protection.

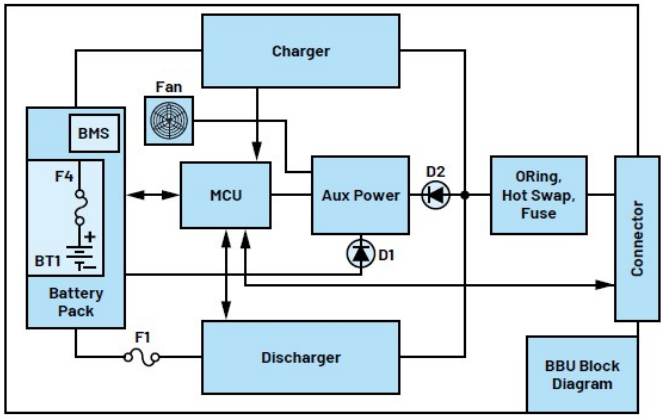
- **Shelf, backplane and mechanical integration:** The shelf distributes DC power and communications across several modules and provides retention, service access and cooling. Modularity allows power and runtime to be scaled by adding units and enables redundancy such as the ORV3 5+1 configuration. Murata specifies a 2U, hot-swappable 3 KW module, while Eaton's 19-inch BBU shelf accepts six 2.75 kW batteries and scales runtime according to the application requirement

Exhibit 37: BBU



Source: Panasonic

Exhibit 38: BBU diagram



Source: ADI

Key functions of BBUs

Backup power and outage ride-through

The BBU's fundamental function is to maintain rack operation during a brief loss of upstream power. The ride-through interval gives the data center time to restore the primary source, transfer to another source, preserve critical data, migrate workloads or perform an orderly shutdown. Rather than attempting long-duration operation, rack BBUs are generally optimized for high power over a relatively short period; the ORV3 example specifies 3KW per module for four minutes, and Murata guarantees a 240-second hold-up period at full-rated load under stated age and temperature conditions.

For AI systems, this protection is particularly valuable because a rack interruption can affect a synchronized workload extending well beyond the failed server. Local ride-through prevents a momentary upstream disturbance from immediately removing GPU, CPU, networking and storage power from the rack. Infineon consequently describes the server BBU as an integrated backup supply intended to prevent data loss and maintain AI workloads during outages or grid failures.

Peak shaving and load smoothing

A BBU can be operated as an active energy buffer. as it charges during lower-demand periods and discharges when server consumption temporarily exceeds a defined rack or PSU limit. This allows a relatively stable upstream power source to support a more volatile downstream AI workload. The economic benefit is that PSU, UPS and electrical capacity can be designed closer to sustained demand rather than every short theoretical peak, although the achievable benefit depends on the frequency, duration and predictability of those peaks.

The important limitation is battery cycling. A BBU designed primarily for emergency standby may not be optimized for continuous, repetitive peak shaving, so frequent cycling affects heat generation, state of charge, usable life and replacement schedules. Very short and highly repetitive transients may therefore be better supported by a CBU or hybrid battery-capacitor system, leaving the battery to cover longer events. Renesas correspondingly distinguishes distributed BBUs from hybrid systems that combine batteries with capacitors for fast transient response and longer-duration backup.

Power stabilization, redundancy and fault isolation

A rack BBU adds another independent source to the local power path. If a power shelf, upstream feed or AC source is interrupted, the BBU can sustain the rack while the fault is isolated or another source assumes the load. A modular shelf further improves resilience because one BBU failure need not remove the entire backup function; the remaining modules can continue operating within the available redundancy and power limits.

Distributed deployment also narrows the fault domain. A defective rack BBU affects a much smaller set of compute equipment than a failure within a centralized backup system serving a large room or data hall. The tradeoff is a much larger population of batteries, converters, controllers and fans that must be monitored, tested and eventually replaced. Digital state-of-health monitoring, fault reporting and hot-swappable construction are therefore central to the operating model rather than optional features.

BBUs vs. UPSs

While some of the functions that BBUs offer look similar to what traditional UPSs can do, we think BBUs and UPSs should be viewed as complementary rather than substitutive technologies, as they address different challenges within the datacenter power architecture.

A UPS is designed to provide facility- or row-level protection against utility outages, voltage sags, surges, and other power-quality disturbances, ensuring continuous AC power delivery to a broad range of IT equipment. In contrast, a BBU operates much closer to the AI workload, supplying DC power directly at the rack level and responding to localized power events that occur faster than the upstream power infrastructure can efficiently manage.

As AI racks become denser and more power intensive, the UPS continues to serve as the first line of defense against external power disruptions, while the BBU acts as a second layer of protection by stabilizing the rack DC bus, supporting transient load spikes, and providing short-duration ride-through capability. Together, they create a hierarchical power protection architecture in which the UPS ensures overall system availability and power quality, while the BBU enhances rack-level resilience and power-delivery performance. Rather than replacing the UPS, BBUs help reduce the burden on centralized backup systems, improve power utilization efficiency, and enable the reliable operation of next generation AI infrastructure with increasingly dynamic power profiles

Exhibit 39: Comparison of BBUs and UPSs

Dimension	BBU	UPS
Location	Inside the server rack, in a BBU shelf, or in an adjacent rack/sidecar close to the compute	Typically deployed at equipment room, row or facility level and protects multiple racks or other critical loads
Power path	Usually DC-coupled to the rack bus through a bidirectional DC/DC converter	Takes AC input and delivers AC output; the server PSU subsequently converts AC again for the server
Primary objective	Short rack-local ride-through, DC-bus support, modular redundancy and—when designed for it—transient or peak-load assistance	Broad power conditioning and uninterrupted AC supply during utility sags, surges, frequency disturbances and outages
Response	Monitors the local bus and can inject power within a few milliseconds; the cited ORV3 reference requires full support within 2ms after a qualifying 2ms event	An online double-conversion UPS already powers the load through its inverter and therefore has zero transfer time between utility and battery operation
Runtime	Generally optimized for short-duration, high-power support; OCP's reference example is four minutes at full load	Runtime varies widely with battery-bank sizing and system design; its broader function is to maintain protected AC loads until another source or operating response becomes available
Scalability	Added rack-by-rack or module-by-module; power and runtime can be matched to each compute deployment	Scaled centrally in larger KW blocks and requires capacity planning for the aggregate protected load
Fault domain	A module or shelf fault is normally localized to one rack or a limited group of racks	A centralized system protects a much broader load domain, so its redundancy and maintenance design are correspondingly critical

Capacitor backup unit (CBU) another complementary power device

What's a CBU?

A CBU is an integrated energy storage subsystem positioned close to the AI load, typically in a dedicated rack shelf, capacitor tray or adjacent power sidecar, and connected to the server rack's DC power path. Its principal storage medium may be an electric double-layer capacitor (EDLC or supercapacitor), lithium-ion capacitor (LIC), or a bank of high-capacitance electrolytic capacitors, depending on the required response time, energy duration, footprint and cost. Dedicated CBU products increasingly use EDLCs or LICs because their low internal resistance supports rapid, repetitive charge and discharge.

The CBU sits electrically alongside the PSU. Under steady-state conditions, the PSU supplies the AI rack and maintains the CBU at its target state of charge. When GPU demand temporarily exceeds the desired PSU output, the CBU discharges through a controlled power converter into the DC bus; when GPU demand falls, the converter reverses direction and recharges the capacitors at a controlled rate. This prevents the full GPU load step from passing through the PSU and upstream equipment. Infineon notes that a separate CBU must be compatible with, and potentially communicate with, the PSU shelves to coordinate rack power.

An alternative to deploying a standalone CBU shelf is to integrate the energy-storage capacitors directly into the power shelf, effectively combining power conversion and transient-energy buffering within the same subsystem. NVIDIA's GB300 NVL72 architecture is a good example: its enhanced power shelves incorporate a substantial capacitor bank alongside the PSU modules, with NVIDIA allocating approximately 65 joules of stored energy per GPU.

NVIDIA reported that this approach reduced peak upstream power demand by roughly 30% in a measured Megatron LLM workload. Compared with a separate CBU, integrating capacitors into the power shelf can save valuable rack space, shorten the electrical path between the energy buffer and the load, and simplify coordination between the PSU and

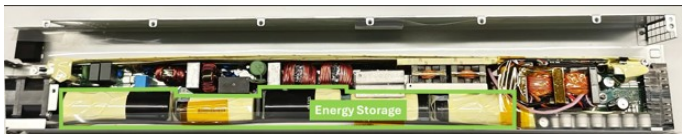
energy-storage system, since charging, discharging, and PSU output can be managed as one integrated power architecture. The trade-off is that capacitor capacity becomes more tightly coupled to the power-shelf design, potentially reducing flexibility in independently scaling or servicing the energy-storage system. Nevertheless, as AI racks become increasingly power dense, integrating transient-energy storage directly into the power shelf could become an attractive alternative to standalone CBUs, particularly where rack space and fast transient response are key design priorities.

Exhibit 40: CBU



Source: Panasonic

Exhibit 41: GB300 power shelf with integrated capacitors



Source: Nvidia

Key components within a CBU

- **Capacitor bank:** The storage section contains many capacitor cells arranged in series to reach the required voltage and in parallel to provide the required power and energy. Individual EDLC cells commonly operate at only a few volts, so high-voltage systems require long series strings; Abracon cites approximately 2.7 V as a typical supercapacitor cell rating, while TI's NVIDIA-oriented reference architecture demonstrates an 800 V CBU using EDLC cells. LICs can provide more stored energy than conventional EDLCs, while EDLCs generally emphasize lower resistance, high peak power and very high cycling endurance. Delta, for example, offers a power-capacitance shelf using fast-charge/discharge LICs, rated to provide 15 seconds of hold-up at 20 KW.
- **Cell-balancing system:** Because cell capacitance, leakage current and internal resistance vary across a series string, the voltage does not divide perfectly equally. A balancing circuit prevents an individual cell from exceeding its rated voltage during charging or being reverse-biased during deep discharge. Passive balancing uses resistors to bleed excess charge, while active approaches use MOSFETs, amplifiers or switched converters to redistribute energy more efficiently; the design tradeoff is between cost, balancing speed, component count and standby loss.
- **Bidirectional DC/DC converter:** This is the active interface between the capacitor bank and the rack DC bus. During a GPU power spike, it converts the capacitor's changing voltage into a regulated bus contribution; during lower demand, it reverses power flow to recharge the bank. The converter also allows more of the capacitors' stored energy to be used than would be possible with a direct connection, because capacitor voltage declines materially during discharge.

Infineon specifically identifies a capacitor bank plus bidirectional converter as the main CBU structure, while Skeleton Technologies notes that converter-based integration provides greater energy utilization and more precise power-flow control than a direct bus connection.

- **Controller and management software:** The controller determines when the CBU should charge, remain idle or discharge, based on bus voltage, PSU output, rack power limits and capacitor state. It may also calculate state of health, manage balancing, regulate recharge current and coordinate multiple CBU modules.
- **Protection and switching:** A high-power capacitor bank can release very large fault current, so fuses, contactors, pre-charge circuits and overvoltage, overcurrent and overtemperature protection remain essential. Pre-charge limits inrush current when the CBU is connected to the DC bus, while contactors or solid-state switches isolate the bank during a fault or service event. Skeleton's reference system integrates switchgear, fuse-fault detection, pre-charge/discharge contactors and protection against high-current and overheating events; industry discussions also highlight the need for faster DC fault isolation as CBU voltage rises.
- **Monitoring and communications:** Voltage, current, temperature, equivalent series resistance and state-of-health measurements allow operators to identify degradation before the CBU can no longer meet its specified peak-power requirement. Panasonic emphasizes voltage, temperature and ESR monitoring for predictive maintenance, while Skeleton provides single-cell voltage sensing, multiple temperature sensors and communication bus to its master controller.
- **Thermal and mechanical system:** Although supercapacitors are efficient, high repetitive current produces resistive heating, particularly in cells, busbars, converters and connectors. The shelf therefore requires designed airflow, temperature sensing and, at higher power levels, potentially liquid cooling. Skeleton's cabinet uses controlled fan cooling and also offers liquid-cooled configurations, while the Open Compute Project presentation from Panasonic identifies lower-resistance capacitors and liquid cooling as development priorities for future higher-output CBUs

Main functions of CBU

Transient power stabilization

One of the most fundamental roles of a CBU is to stabilize the rack DC power bus when GPU power consumption changes very rapidly. AI workloads frequently alternate between compute-intensive phases and lower-power communication, synchronization, or idle phases. Because many GPUs operate in parallel, these transitions can create very large changes in rack power within milliseconds.

When GPU demand suddenly increases, the PSU cannot instantaneously increase its output. The CBU therefore discharges stored energy into the DC bus during this short response gap, helping prevent voltage drop and excessive current excursions. Conversely, when GPU demand suddenly falls, the CBU can absorb excess energy while the PSU reduces its output, helping limit voltage overshoot.

The key objective here is power quality and system stability, rather than reducing the overall power capacity required by the rack. In other words, the CBU acts as a fast

electrical shock absorber between the highly dynamic GPU load and the slower PSU. This becomes increasingly important as AI racks move above 100 KW and synchronized GPU workloads create larger absolute power swings.

Peak shaving

Another major function is peak shaving, whose objective is to limit the maximum power that the PSU and upstream infrastructure need to supply.

The benefit is potentially significant because the PSU, connectors, protection devices, and portions of the upstream power infrastructure no longer need to be oversized solely for short-duration GPU peaks. It can also allow PSUs to operate closer to their optimal loading range, improving utilization and potentially efficiency.

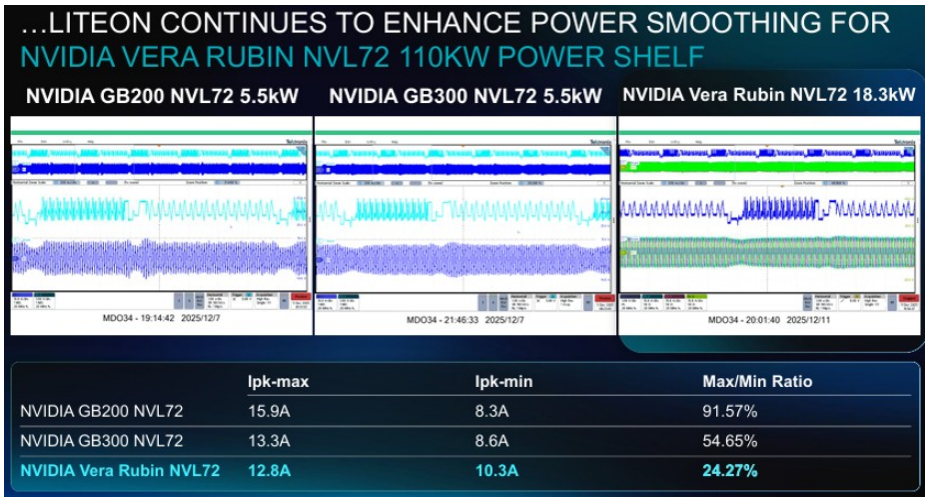
Short duration ride-through

The third function is ride-through, where the CBU temporarily powers the rack when the normal power source is interrupted altogether. This is fundamentally different from the first two functions because the issue is no longer a mismatch between GPU demand and PSU response—the upstream power source itself has temporarily disappeared.

A sufficiently sized CBU can maintain the DC bus during events such as a brief AC interruption, PSU switchover, source-transfer event, or temporary power-shelf failure. Depending on the amount and type of capacitance installed, this support can range from milliseconds to several seconds. For example, some LIC-based capacitance shelves are designed to provide multi-second hold-up at high rack power.

However, ride-through is generally a secondary function for a CBU. Capacitors have high power density but relatively low energy density, meaning they can deliver a large amount of power very quickly but cannot economically sustain that power for long periods. As the required backup duration increases from seconds toward minutes, a battery-based BBU generally becomes more suitable.

Exhibit 42: Capacitors help smooth power delivery



	Ipk-max	Ipk-min	Max/Min Ratio
NVIDIA GB200 NVL72	15.9A	8.3A	91.57%
NVIDIA GB300 NVL72	13.3A	8.6A	54.65%
NVIDIA Vera Rubin NVL72	**12.8A**	**10.3A**	**24.27%**

Source: Lite On

CBUs vs. BBUs

The two technologies should consequently be regarded as complementary rather than directly substitutive. A CBU handles the high-power, short-duration and high-cycle portion of the load profile, while a BBU supplies the larger amount of energy required when the upstream source remains unavailable. Some describe BBUs and CBUs as coexisting in the same power sidecar, and hybrid architectures can allocate fast transients to capacitors while preserving battery charge and cycle life for longer outages.

The CBU's major advantages are rapid bidirectional response, long cycle life, high efficiency and the ability to absorb repeated GPU peaks without the same electrochemical wear mechanism as a conventional battery. It can also recharge quickly between workload cycles and may require less routine replacement. However, the CBU has materially lower energy density, its voltage falls continuously during discharge, and large banks require balancing, high-current protection, cooling and valuable rack or sidecar space.

Exhibit 43: Comparison of CBUs and BBUs

Dimension	CBU	BBU
Primary storage medium	EDLC supercapacitors, LICs or high-capacitance	Typically lithium-ion battery cells managed by a BMS
Primary strength	High power, fast charge/discharge and frequent	Higher stored energy and longer sustained
Main function	Transient smoothing, peak shaving, voltage	Outage backup, source-transfer bridging and orderly workload
Response profile	Well suited to millisecond-class load support and	Fast switchover, but generally optimized for sustained support rather
Typical duration	Milliseconds to seconds; selected systems can extend toward tens of seconds or more	Commonly tens of seconds to several minutes, depending on battery and rack design
Power versus energy density	Higher power capability but lower energy storage per unit volume than a battery	Higher energy density, making it more economical for longer events
Cycle life	Eaton specifies more than one million cycles and up to 20 years for certain supercapacitor modules	Battery life is more sensitive to depth of discharge, temperature and repeated cycling; BMS supervision is therefore essential.
State management	Requires cell balancing, voltage/temperature monitoring and state-of-health tracking, but does not use a conventional battery fuel	Requires BMS functions including charge control, state of charge, state of health and cell protection
Best-fit use case	Repetitive AI power spikes, inrush events, PSU load leveling and bus stabilization	Power continuity through outages, brownouts or generator/source-transfer intervals

HVDC Emerging, with SST in the Pipeline

HVDC coming along with the increasing power consumption of AI clusters

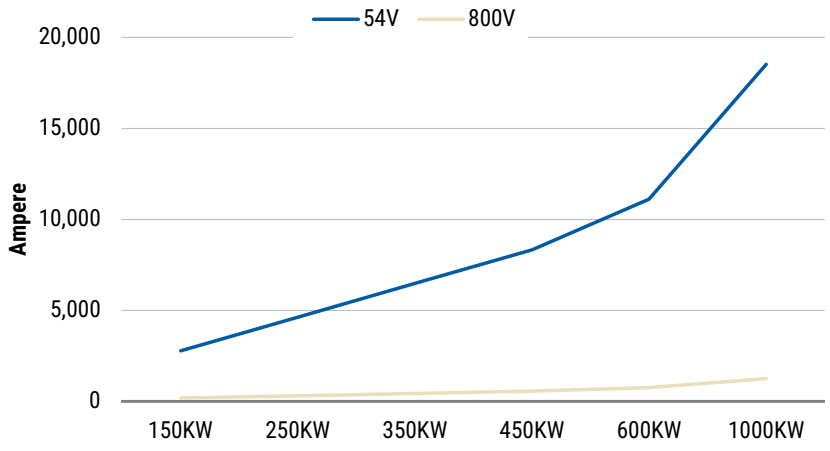
Why do we need HVDC?

Low voltage turns higher rack power into an extreme-current problem

The fundamental relationship is that at a given rack load, increasing the bulk-distribution voltage from 54 V to 800 V reduces current by approximately 14.8 times. A balanced ±400 V architecture also presents 800 V between its +400 V and –400 V conductors, so its bulk-transfer current is broadly equivalent to a 0–800 V architecture.

This current differential has an outsized mechanical consequence. For the same conductor resistance, I^2R loss at 54V would theoretically be approximately 219 times that at 800 V, though real systems avoid this level of loss by adding copper, parallel conductors and cooling. The result is not a 219-fold difference in actual operating loss, but a trade-off in which low-voltage designs consume increasing amounts of busbar cross-section, connector area, rack volume and cooling capacity to remain within voltage-drop and temperature limits.

Exhibit 44: Transmission current for different power scales: 54V vs 800V



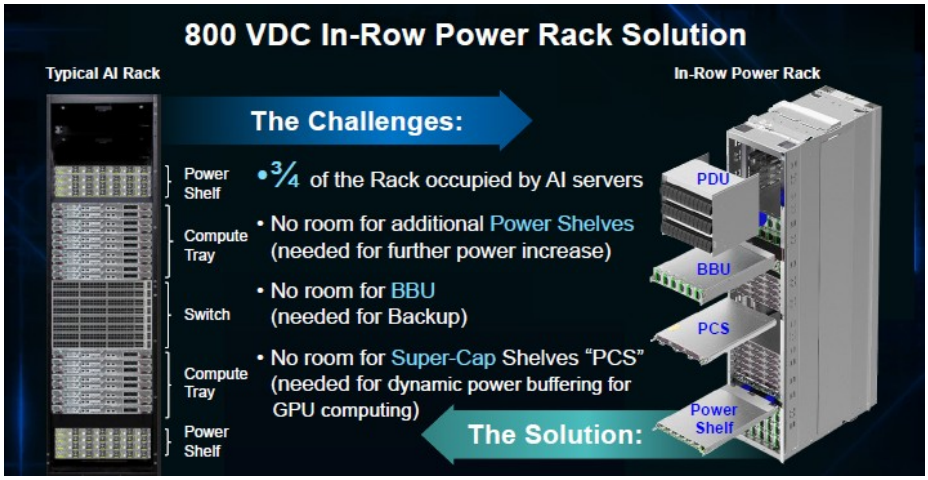
Source: Delta

Copper, connectors and thermal management become system constraints

At several thousand amperes, distribution cannot rely on a single ordinary conductor or connector. Designers must use broad laminated busbars, multiple parallel contacts, tightly controlled joint resistance and temperature monitoring. A small increase in contact resistance produces disproportionately high heat at these currents, raising the risk of hot spots, connector ageing and nuisance protection events. OCP's 100KW HVDC-to-LVDC shelf illustrates where the current problem moves: only two 62.5 A HVDC inputs are required on the ±400V side, while the 49–50V output approaches 2,000A and therefore uses a substantial bolted output connection with dedicated thermistors.

The same issue appears in power-shelf packaging. Delta notes that a typical AI rack may already allocate roughly three-quarters of its volume to servers, leaving insufficient room for additional power shelves, BBUs or supercapacitor shelves as power rises. Moving AC/DC conversion, protection and storage into an adjacent power rack therefore releases compute space and separates the thermal design of power electronics from that of GPUs. Microsoft's Mt Diablo rationale similarly states that disaggregated power can enable up to 35% more accelerators in the IT rack and allows power shelf count to be sized independently for each workload configuration

Exhibit 45: Not enough room for power devices on the traditional AI server rack



Source: Delta

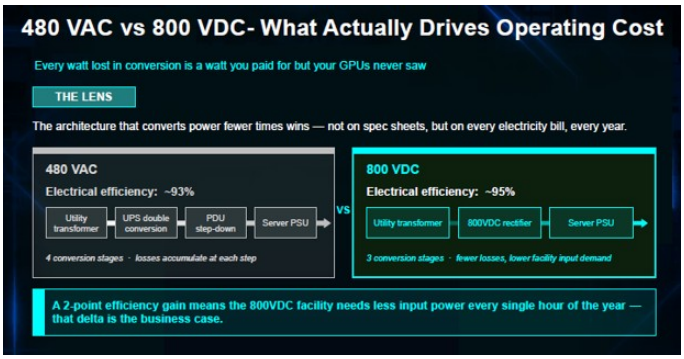
HVDC enables denser, fewer and more efficient power stages

As AI rack power scales from today's GB200/GB300 systems to Vera Rubin and future multi-hundred-kilowatt racks, power distribution efficiency becomes increasingly important. Our checks indicate that a conventional 480VAC architecture delivers an end-to-end electrical efficiency of roughly 93%, while an 800VDC architecture improves efficiency to about 95% by eliminating one power conversion stage and simplifying the power path from grid to rack.

The higher-voltage 800VDC distribution also reduces current levels for the same power load, lowering resistive (I^2R) losses in cables and distribution equipment and enabling more efficient delivery to increasingly power-dense AI racks.

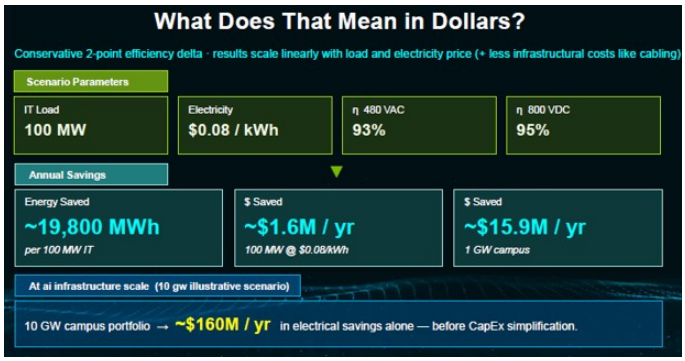
While a 2ppts of efficiency improvement may appear modest, the savings become substantial at hyperscale AI data center deployments. Delta's example shows that for a 100MW IT load, improving facility efficiency from 93% to 95% reduces annual energy consumption by approximately 19,800MWh, translating into roughly US$1.6 million of annual electricity cost savings assuming a power price of US$0.08/kWh. At a 1GW AI campus, the savings scale to nearly US$16 million per year, and at emerging multi-gigawatt AI factory deployments, the annual benefit can exceed US$100 million. Consequently, the transition from conventional AC distribution and 54VDC rack power architectures toward 800VDC distribution is not only a technical necessity for supporting future 500kW-to-1MW AI racks, but also a compelling economic proposition that lowers operating expenses and improves total cost of ownership through reduced energy losses and simplified electrical infrastructure.

Exhibit 46: Better power efficiency for 800VDC architecture



Source: Delta

Exhibit 47: Better efficiency leads to cost savings



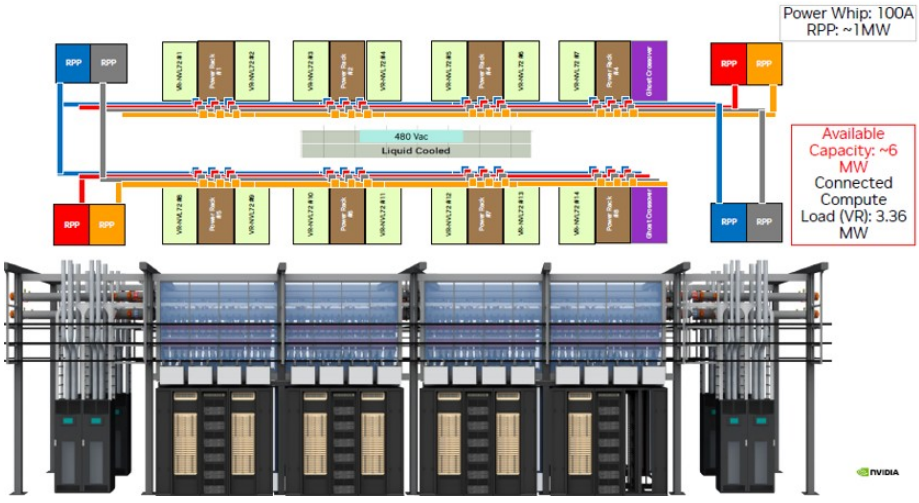
Source: Delta

What's the deployment roadmap for HVDC?

The deployment roadmap for 800VDC is best understood as a multi-stage evolution rather than a single-step transition. Both Nvidia and major hyperscalers recognize that existing 480VAC infrastructure cannot efficiently support the rapid increase in rack power. As rack densities increase, conventional power whips, busbars, connectors, and PSU architectures increasingly become bottlenecks, driving the need for higher voltage distribution

The first phase of deployment focuses on introducing 800VDC power racks within existing AC-based data centers. In this architecture, the facility remains largely unchanged, but a dedicated power rack converts 415-480VAC into 800VDC before distributing power locally to AI racks. This approach minimizes disruption to existing infrastructure while allowing operators to deploy Vera Rubin-class systems that exceed the practical limits of traditional power shelves. Nvidia refers to this as the row-level power rack architecture, which serves as the industry's bridge solution between legacy AC facilities and future native DC environments.

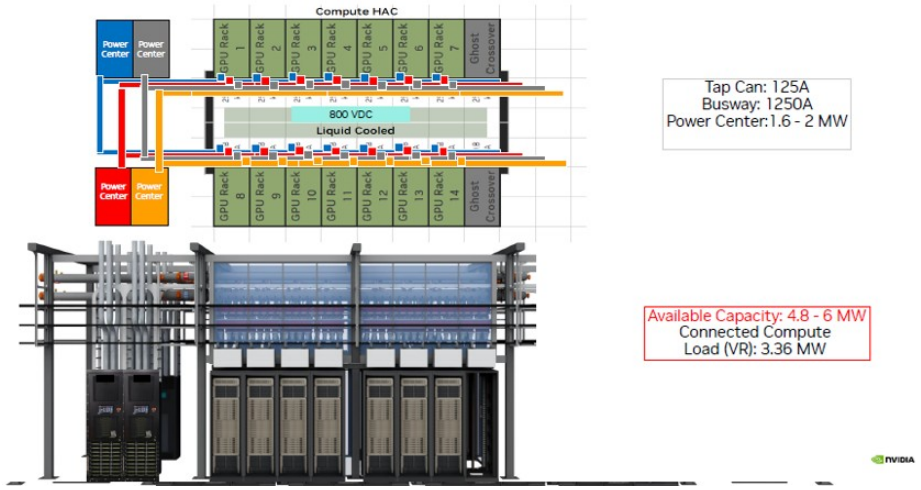
Exhibit 48: Phase I - 800VDC power rack



Source: Nvidia

The second phase expands 800VDC beyond individual rows into data hall-level power centers. Instead of deploying a dedicated power rack for every few compute cabinets, centralized rectifier or power-center systems distribute 800VDC through DC busways across an entire hall. Nvidia's roadmap identifies 1.6-2MW power centers and 1250A DC busway architectures as the next step, improving power density while reducing equipment duplication. This stage also aligns well with the deployment strategies being discussed by hyperscalers and the OCP community because it preserves existing medium-voltage and low-voltage infrastructure while moving a significant portion of distribution into DC

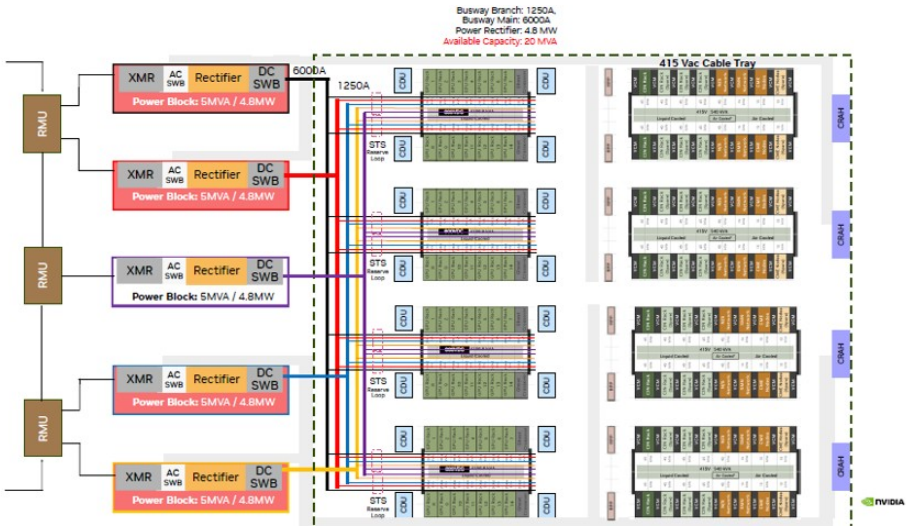
Exhibit 49: Phase II – 800VDC power center



Source: Nvidia

The third phase is the transition toward facility-wide 800VDC distribution, where AC is converted to DC near the utility entrance or medium-voltage substation and distributed throughout the data center via high-capacity DC switchboards and busways. Nvidia describes this as the DC power block architecture, utilizing approximately 4.8MW power blocks, 6000A DC busways, and integrated redundancy schemes. At this stage, the architecture begins to resemble an "AI factory," where power distribution is simplified, energy storage can be integrated directly onto the DC bus, and multiple megawatts of compute capacity can be deployed as modular building blocks

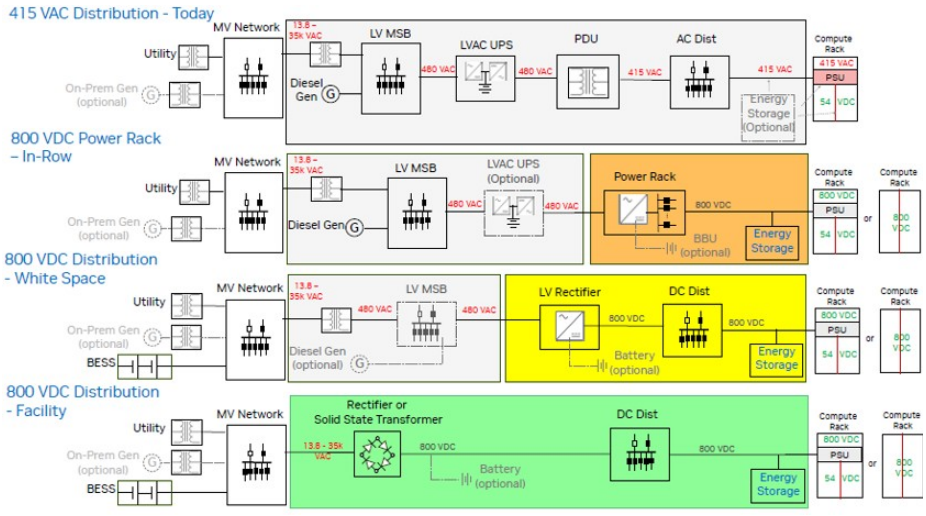
Exhibit 50: Phase III – power block



Source: Nvidia

The final stage is the emergence of native 800VDC compute racks, which Nvidia identifies as the long-term destination. Today, even in 800VDC deployments, power is generally stepped down to 54V before entering the servers. Future systems, however, are expected to accept 800VDC directly within the rack and potentially at the server level, eliminating another conversion stage and further reducing power losses. Nvidia's public roadmap shows a progression from current 54V architectures, to 800VDC-fed PSUs, and ultimately to 1MW-class native 800VDC racks capable of supporting next-generation AI accelerators

Exhibit 51: 800VDC native server racks



Source: Nvidia

In sum, deployment is likely to occur gradually because the ecosystem remains under development. NVIDIA notes that several enabling technologies—including DC protection devices, busways, connectors, grounding standards, arc-flash modeling, and solid-state transformers— are still undergoing certification and standardization. The roadmap presented by NVIDIA suggests a phased commercialization timeline, with 800VDC power racks reaching deployment readiness in 2026, hall-level power centers in 2027, and facility-scale DC power blocks becoming deployment-ready around 2028. Industry

collaboration among NVIDIA, Microsoft, Google, and OCP participants is focused on ensuring interoperability and regulatory acceptance before broader hyperscale adoption.

Exhibit 52: HVDC deployment roadmap



Focus Area	Status	Impact
Grounding & Protection	Architecture converging with UL and OEMs	Foundation for safe and reliable deployment
Equipment Certification	Certification pathways established and progressing	Accelerates deployment readiness
Protection Devices	MCCB available; SSCB ecosystem maturing	Enables fast fault isolation and higher availability
Busway, Connector & Whips	Standardized 125A architecture under development	Supports scalable rack deployment
High-Density Power Conversion	Rectifiers available; SST roadmap advancing	Enables future 800 VDC Compute Product

Source: Nvidia

800VDC represents 4-5x dollar content vs. in-rack power supply architecture

The 800VDC power rack BOM is primarily driven by the PSU, which accounts for 50–55% of the total cost, followed by the BBU at 25–30%, with other components comprising the remaining 15–20%. This cost structure reflects the broader functionality integrated into an HVDC power rack, combining power conversion, backup power and supporting components within a centralized architecture.

For power supply players, one of the most implications is that an HVDC power rack carries around 4–5x higher dollar content per watt compared with a conventional in-rack PSU solution, since the former integrates more components into a robust system and power supply players usually ship the entire rack instead of just the components.

Exhibit 53: BOM cost breakdown for an 800VDC power rack

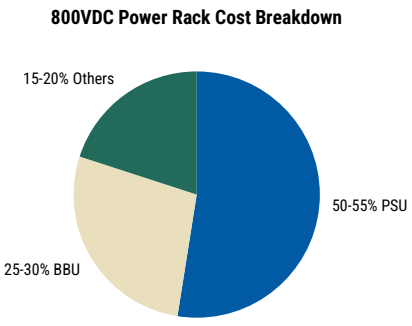
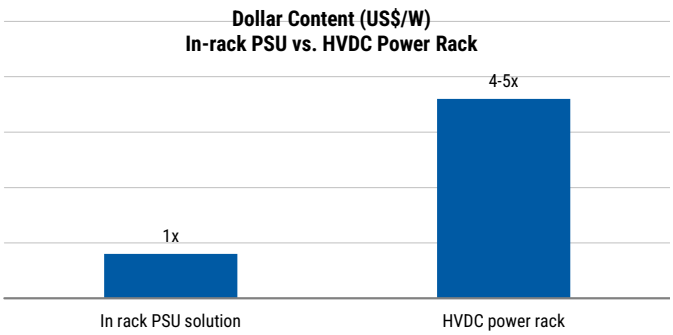


Exhibit 54: Higher dollar content for HVDC power racks



800VDC vs. ±400VDC?

The debate between ±400VDC and 800VDC centers on balancing near-term practicality against long-term scalability. While both architectures deliver an identical 800V differential voltage and therefore achieve similar reductions in current, copper loss, and cable size versus traditional 54V rack power systems, they differ significantly in implementation.

Under a ±400VDC architecture, each conductor operates at only 400V relative to ground, resulting in less stringent insulation, creepage, clearance, and protection requirements. This makes the architecture easier to deploy, maintain, and certify, which explains why many hyperscalers and OCP members view ±400VDC as an attractive migration path toward high-density AI racks.

By contrast, NVIDIA and several power infrastructure vendors are increasingly advocating for a single-bus 800VDC architecture because it simplifies overall power distribution, reduces conversion stages, and is better suited for future 500kW to 1MW rack deployments.

We believe that ±400VDC appears well positioned as a near-term bridge architecture due to its lower deployment risk and more mature ecosystem, whereas 800VDC is increasingly viewed as the long-term end-state architecture for AI factories as rack power scales toward the megawatt era.

Exhibit 55: Comparison of 800VDC and ±400VDC

Metric	±400VDC	800VDC
Voltage to ground	400V relative to ground	800V relative to ground
Voltage differential	800V	800V
Power distribution efficiency	High	Slightly higher
Scalability	Good	Excellent
Insulation requirements	Lower	Higher
Maintenance safety	Better	More challenging
System architecture simplicity	Dual-rail system	Single-bus system
Grounding Complexity	More complicated	Simpler
Electrical footprint	Small	Smaller
Ability to reduce conversion stages	Moderate	Highest
Long-distance distribution capability	Excellent	Excellent
Maturity of ecosystem	Better today	Early but rapidly developing
Industry players	OCP, hyperscalers	Nvidia
Long term industry direction	Transitional architecture	Likely end-state architecture

From a component supplier perspective, the transition toward HVDC is far more important than which architecture ultimately wins, as most hardware content increases under either scenario. PSU vendors, BBU suppliers, busbar manufacturers, connector providers, power-whip suppliers, magnetics vendors, and rack-level DC-DC converter manufacturers are largely architecture-agnostic, since both ±400VDC and 800VDC require a similar shift toward high-voltage power distribution, centralized power shelves, and rack-level conversion.

Our understanding is that the primary exception is the power semiconductor industry, where the choice of architecture could materially influence device content and technology adoption. A ±400VDC system can often be supported using existing 650V-900V MOSFET and SiC platforms, benefiting suppliers with exposure to lower-voltage power devices. By contrast, 800VDC architectures increasingly require 1200V-class SiC MOSFETs, higher-

voltage gate drivers, hot-swap controllers, isolation components, and advanced protection circuitry, driving greater semiconductor content per watt and raising technology barriers.

Solid state transformer potentially being the holy grail in AI datacenter power delivery

What is a solid-state transformer ?

A datacenter solid state transformer (SST) integrates power electronics to replace or augment traditional line-frequency transformers, and it's typically built around three major functional stages. The first is the medium-voltage AC-to-DC rectifier stage, which uses high-voltage semiconductor switches—often silicon carbide (SiC) or gallium nitride (GaN) devices due to their superior switching speed and thermal performance—to convert incoming grid-level medium voltage AC into a regulated DC bus. The second is the isolated DC-DC conversion stage, usually implemented as a dual active bridge (DAB) or LLC resonant converter, which provides galvanic isolation through a high-frequency transformer (operating in the kHz to tens-of-kHz range rather than 50/60 Hz), dramatically shrinking the magnetic core and transformer size compared to conventional units while stepping voltage down toward datacenter distribution levels. The third is the low-voltage DC-AC or DC-DC output stage, which conditions power into the form needed by the facility—either regulated DC for direct DC-bus distribution to IT racks or inverted AC for legacy PDU/UPS infrastructure. Supporting these power stages are the control and sensing subsystems (digital controllers, gate drivers, voltage/current sensors) that manage switching, fault protection, and grid synchronization; thermal management components (heat sinks, cold plates, or liquid cooling loops) to handle the concentrated losses in the semiconductor devices; and increasingly, integrated energy storage interfaces or DC microgrid ports that allow the SST to natively support battery buffering, renewable integration, and finer-grained power quality control than a passive transformer ever could

A Solid-State Transformer (SST) is a power-electronics-based transformer that replaces the traditional 50/60Hz iron-core transformer with a combination of power semiconductors, control electronics, and a medium-frequency transformer. Unlike conventional transformers that perform only voltage conversion, SSTs actively process power through multiple stages, enabling both AC and DC inputs/outputs, bidirectional power flow, voltage regulation, reactive power compensation, and intelligent grid control

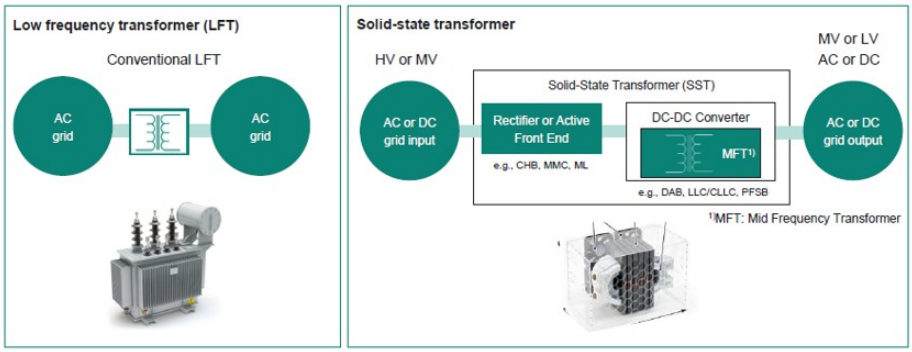
A SST is typically built around three major functional stages:

- The first is the medium-voltage AC-to-DC rectifier stage, which uses high-voltage semiconductor switches—often silicon carbide (SiC) or gallium nitride (GaN) devices due to their superior switching speed and thermal performance—to convert incoming grid-level medium voltage AC into a regulated DC bus.
- The second is the isolated DC-DC conversion stage, usually implemented as a dual active bridge (DAB) or LLC resonant converter, which provides galvanic isolation through a high-frequency transformer (operating in the KHz to tens-of-LHz range rather than 50/60 Hz), dramatically shrinking the magnetic core and transformer size compared to conventional units while stepping voltage down toward datacenter distribution levels.
- The third is the low-voltage DC-AC or DC-DC output stage, which conditions

power into the form needed by the facility—either regulated DC for direct DC-bus distribution to IT racks or inverted AC for legacy PDU/UPS infrastructure.

Supporting these power stages are the control and sensing subsystems (digital controllers, gate drivers, voltage/current sensors) that manage switching, fault protection, and grid synchronization; thermal management components (heat sinks, cold plates, or liquid cooling loops) to handle the concentrated losses in the semiconductor devices; and increasingly, integrated energy storage interfaces or DC microgrid ports that allow the SST to natively support battery buffering, renewable integration, and finer-grained power quality control than a passive transformer ever could.

Exhibit 56: Solid state vs. traditional transformers



Source: Infineon

Why do we need SSTs?

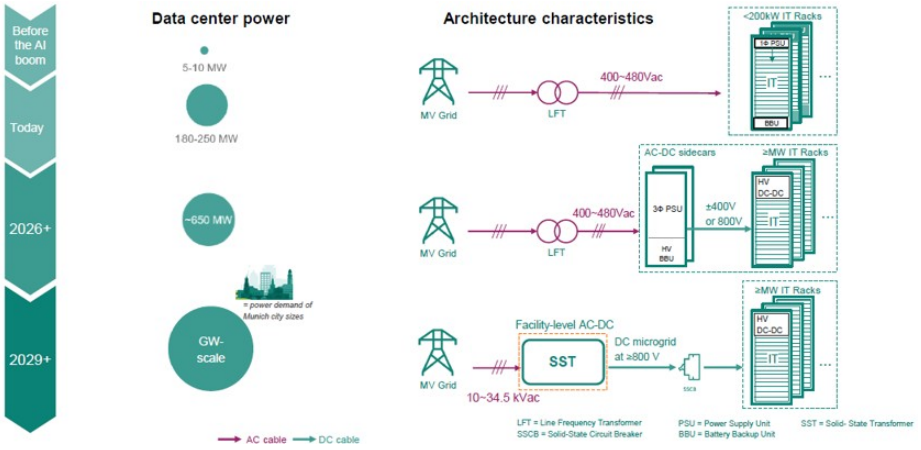
Improving efficiency in an era of gigawatt data centers

One of the strongest drivers for SST adoption is the rapid increase in AI data center power demand, which has progressed from 5-10MW facilities in the pre-AI era to campuses exceeding 650MW today, with gigawatt-scale facilities expected beyond 2029.

Traditional architectures use medium-voltage AC distribution, followed by low-frequency transformers and multiple power conversion stages before power reaches IT racks. SSTs allow facilities to convert 10-34.5KV medium-voltage power directly into ±400V or 800V DC distribution networks. This reduces conversion losses and improves overall grid-to-chip efficiency.

The efficiency advantage could be at the order of about 2-5ppt vs, the traditional transformer architecture, but at a 1GW AI campus, this magnitude of improvement could translate into roughly 20-50MW of continuous power savings. Over the lifetime of a facility, this can result in substantial reductions in electricity cost and carbon emissions. SSTs are therefore becoming increasingly attractive for hyperscaler operators seeking lower total cost of ownership (TCO).

Exhibit 57: SST enabling more efficient power delivery architecture



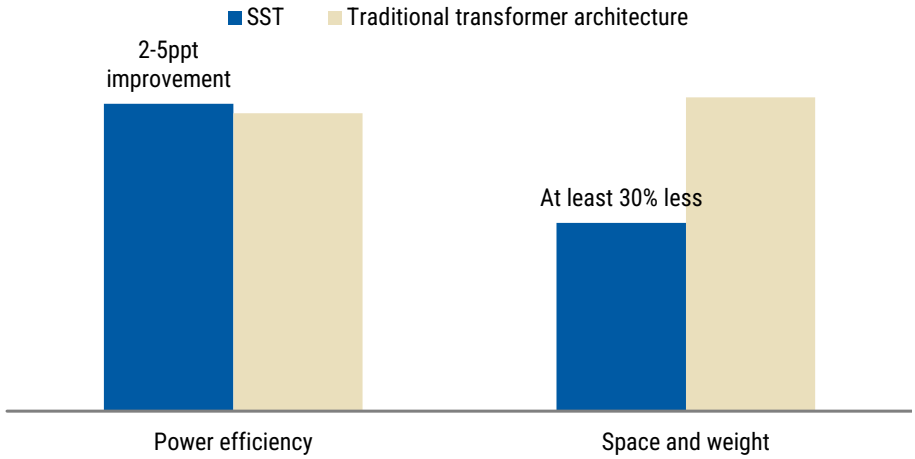
Source: Infineon

Higher power density and smaller footprint

Because SSTs use medium-frequency transformers instead of bulky 50/60Hz transformers, they can achieve significantly higher power density. Infineon indicates that SST systems can reduce both footprint and weight by at least 30% compared with conventional transformer-based installations. They also accelerate site energization and facilitate modular deployment

This becomes particularly valuable in urban data centers, utility substations, EV charging hubs, and energy storage sites where space constraints directly affect project economics.

Exhibit 58: SST offers better power efficiency with at least 30% less space/weight



Supporting distributed energy resources

As renewable energy penetration increases, future grids will require seamless integration of solar generation, battery storage, fuel cells, gas turbines, and other distributed energy resources.

SSTs provide intelligent power routing, bidirectional operation, and dynamic voltage control. This allows them to function as energy hubs that connect renewable generation,

storage, and loads through a common DC infrastructure. Such capabilities are difficult or impossible to achieve using passive transformers alone.

Improved grid intelligence

Unlike conventional transformers, SSTs can actively regulate voltage and frequency, provide reactive power compensation, support grid-forming functionality, and improve overall power quality. These capabilities make SSTs an important enabler for future smart grids, microgrids, and virtual power plants.

What's the deployment status now?

Despite decades of research, SSTs remain at an early commercial deployment stage. While we started to see the SST models being unveiled by relevant suppliers and there start to be some pilot adoption, we believe that the 800VDC power rack/center/block architecture shall be deployed first, and SSTs will take more time, as more fundamental design change in the datacenter and all hardware inside will involved. We believe wider adoption probably won't happen before 2030.

Exhibit 59: SST from Delta



Source: Delta

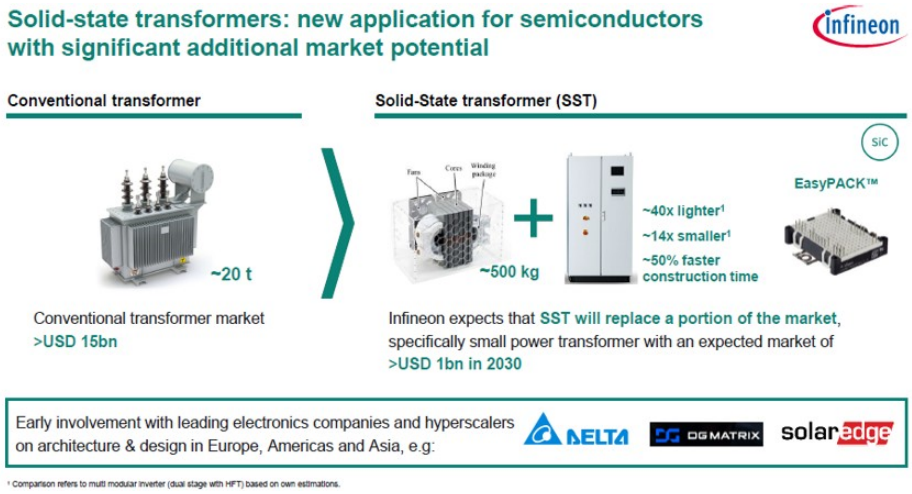
Exhibit 60: SST from Eaton



Source: Eaton

According to Infineon, SST could represent a >US$1bn market opportunity by 2030, as the technology begins to replace a portion of the >US$15bn conventional transformer market, particularly in smaller-power applications. While SST penetration is expected to remain relatively modest versus the overall transformer market, the opportunity is significant for relevant players in the longer term.

Exhibit 61: SST to see initial adoption in 2030



Source: Infineon

Key hurdles for the wider adoption of SSTs

- **High system cost remains the primary barrier:** Despite their technical advantages, SSTs remain significantly more expensive than conventional transformers. Traditional transformers consist primarily of steel cores, copper windings, insulation materials, and cooling systems, all of which benefit from mature manufacturing processes and economies of scale. In contrast, SSTs require large numbers of high-voltage power semiconductors, gate drivers, digital controllers, sensors, protection systems, and sophisticated thermal management solutions. The extensive use of silicon carbide (SiC) devices, which remain substantially more expensive than conventional silicon-based components, further increases the overall system cost. Although SSTs can deliver efficiency gains, improved controllability, and reduced footprint, the economic benefit is often insufficient to offset the substantial upfront capital expenditure in many utility applications. Consequently, SST deployment is currently limited to applications where efficiency, power density, or system flexibility carry exceptionally high economic value, such as AI data centers, megawatt EV charging stations, and grid-scale battery storage systems.
- **Long term reliability has yet to be proven:** Electric utilities and large infrastructure operators place an extremely high value on reliability because transformers are expected to operate continuously for several decades with minimal maintenance. Conventional transformers routinely achieve service lifetimes exceeding 30 to 40 years and have accumulated more than a century of operating history. SSTs, however, are fundamentally different devices that incorporate thousands of semiconductor junctions, complex control electronics, communication networks, cooling systems, and software algorithms. While each individual technology may be mature, the long-term reliability of the complete SST system remains largely unproven at scale. Operators therefore remain cautious about replacing a passive asset with an active power-electronic system whose failure modes are more complex and less understood. For critical infrastructure such as substations, industrial facilities, or hyperscale data centers, reliability concerns often outweigh potential efficiency improvements.

- **DC protection technology ss still evolving:** One of the most significant technical challenges associated with SST deployment is fault protection, particularly in medium-voltage and high-voltage DC networks. Unlike AC systems, where current naturally crosses zero multiple times each cycle and facilitates fault interruption, DC currents continue flowing continuously during fault conditions. As a result, fault currents can rise extremely quickly, making them more difficult to detect and interrupt safely. Future SST-based power architectures therefore require advanced protection devices such as solid-state circuit breakers (SSCBs), which can react far faster than conventional mechanical breakers. However, DC protection technologies are still evolving, and there is currently no universally accepted approach equivalent to the mature protection schemes available in AC grids. The absence of standardized protection architectures increases deployment risk and remains one of the key obstacles to the wider adoption of SST-based DC distribution systems.
- **System complexity creates engineering, qualification and maintenance challenges:** Compared with traditional transformers, SSTs are considerably more complex systems that require the integration of multiple high-performance subsystems. A complete SST includes medium-voltage power converters, isolated DC/DC stages, high-frequency transformers, semiconductor modules, digital controllers, communication interfaces, and advanced protection mechanisms. As voltage and power levels increase, the number of converter modules and semiconductor devices required can rise significantly. Although the introduction of higher-voltage SiC devices can reduce subsystem count and simplify design, SST architectures remain substantially more complicated than conventional transformers. This complexity increases development costs, certification requirements, maintenance needs, and operational risks, all of which slow adoption among traditionally conservative utility customers.
- **Industry standards and ecosystems are still developing:** The transformer industry benefits from decades of established standards, testing procedures, maintenance practices, and supply chain infrastructure. SSTs, by comparison, represent an emerging technology with limited deployment history and evolving standards. Many utilities remain uncertain regarding qualification procedures, maintenance requirements, spare-parts strategies, and long-term asset management. Furthermore, widespread SST adoption will require parallel development of supporting technologies such as DC protection systems, DC switchgear, energy management software, and medium-voltage DC distribution networks. Until the broader ecosystem matures, many customers will continue to view SSTs as a promising but relatively high-risk alternative to proven transformer technologies.

Fuel Cells for Better Flexibility and Resilience

What are fuel cells?

Fuel cells are electrochemical energy-conversion devices that generate electricity by combining hydrogen and oxygen without combustion. Unlike conventional power generation technologies, which convert fuel into heat and then into electricity through multiple intermediate steps, fuel cells produce electricity directly through an electrochemical reaction, resulting in higher efficiency and significantly lower emissions.

When powered by hydrogen, the primary byproducts are water and heat, making fuel cells one of the cleanest forms of power generation available today.

Fuel cells can be deployed across a wide range of applications, including passenger vehicles, heavy-duty trucks, buses, trains, ships, stationary power generation, backup power systems, and increasingly data centers. As governments and industries pursue decarbonization while seeking reliable and dispatchable power sources.

Major fuel-cell technologies include:

- SOFC (Solid Oxide Fuel Cell): Preferred for stationary power and data centers.
- PEMFC (Proton Exchange Membrane Fuel Cell): Dominant in transportation and backup power.
- MCFC (Molten Carbonate Fuel Cell): Large-scale distributed generation.
- PAFC (Phosphoric Acid Fuel Cell): Commercial combined heat and power.
- DMFC (Direct Methanol Fuel Cell): Portable applications

Exhibit 62: Fuel cell offering - Delta



Source: Delta

Exhibit 63: Fuel cell offering - Bloom Energy



Source: Bloom Energy

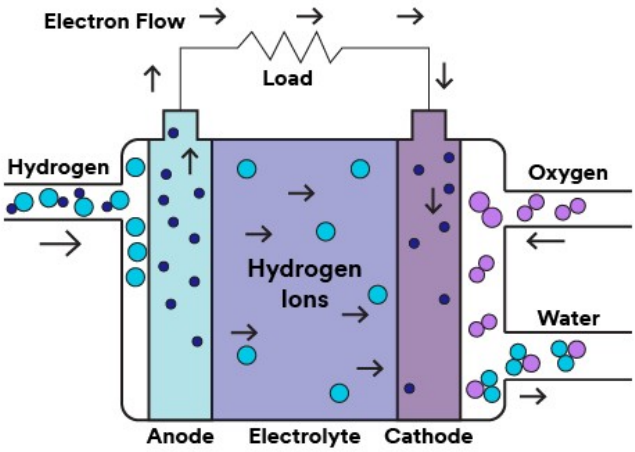
How do fuel cells work?

While there are several fuel-cell architectures, the fundamental operating principle is largely the same: fuel is supplied to the anode, oxygen is supplied to the cathode, and an electrolyte in between allows charged ions to move while forcing electrons to flow through an external circuit. This flow of electrons generates usable electricity, while the

reaction also produces heat and water as byproducts. Unlike conventional generators, which rely on burning fuel to drive a turbine or engine, fuel cells produce electricity directly through electrochemical reactions, allowing them to achieve significantly higher efficiency and lower emissions.

Using solid oxide fuel cell (SOFC) technology as an example, air is supplied to the cathode side of the cell, where oxygen molecules gain electrons and form oxygen ions (O^{2-}). These oxygen ions migrate through a ceramic electrolyte toward the anode side of the cell. At the same time, hydrogen or another hydrogen-rich fuel is supplied to the anode. When the oxygen ions reach the anode, they react with the fuel to produce water, heat, and electrons. The generated electrons cannot pass through the electrolyte and therefore travel through an external electrical circuit before returning to the cathode, creating a continuous flow of electricity. As long as fuel and air continue to be supplied, the fuel cell can generate power continuously.

Exhibit 64: Fuel cell diagram



Source: 3M

What benefits can fuel cells offer?

- **High energy efficiency:** Fuel cells convert the chemical energy of hydrogen directly into electricity through an electrochemical reaction, avoiding the multiple energy-conversion steps associated with combustion-based power generation. This allows fuel cells to achieve electrical efficiencies of over 60%, with some systems exceeding 80% total efficiency when waste heat is recovered through combined heat-and-power configurations, compared with roughly 35-45% for conventional gas-fired power generation. The efficiency advantage translates directly into lower fuel consumption, reduced operating costs, and improved energy utilization over the lifetime of the system.
- **Low carbon emissions:** When powered by hydrogen, fuel cells produce only water and heat as byproducts, generating no direct carbon emissions at the point of use. Even when operating on hydrogen derived from natural gas, fuel cells typically produce significantly lower greenhouse gas emissions than conventional combustion-based power systems. As green hydrogen production scales, fuel cells offer a potential pathway toward near-zero-carbon power generation for transportation, industrial facilities, and distributed energy systems.

- **Grid independence and energy security:** Fuel cells can provide reliable on-site power generation independent of the electrical grid, improving energy security and resilience. This makes them attractive for mission-critical applications such as data centers, hospitals, telecommunications infrastructure, defense installations, and industrial facilities that require uninterrupted power. In regions facing grid congestion or reliability concerns, fuel cells can also serve as a distributed power source located directly at the point of consumption.
- **Long duration energy storage enabler:** Fuel cells play an important role in the emerging hydrogen economy by enabling excess renewable energy from solar and wind farms to be stored as hydrogen and later converted back into electricity. Unlike batteries, which become increasingly expensive for longer storage durations, hydrogen storage can economically provide energy over days, weeks, or even seasonal periods. This capability could become increasingly valuable as renewable energy penetration rises worldwide.
- **Fuel flexibility:** Certain fuel-cell technologies, particularly Solid Oxide Fuel Cells (SOFCs), can operate on a variety of fuels including hydrogen, natural gas, biogas, ammonia-derived hydrogen, and other hydrogen-rich fuels. This flexibility allows users to deploy fuel-cell systems today using existing fuel infrastructure while maintaining a pathway toward pure hydrogen operation as hydrogen supply chains mature.
- **Quiet operation and reduced maintenance:** Because fuel cells contain relatively few moving parts compared with combustion engines and turbine-based generators, they operate quietly and often require less routine maintenance. The absence of combustion also reduces mechanical wear and vibration, which can improve system reliability and lower lifecycle operating costs, particularly in stationary power applications.
- **Scalability across multiple applications:** Fuel cells can be deployed across a wide range of power levels, from portable power systems measured in watts to utility-scale installations generating multiple megawatts. This scalability allows the technology to serve residential buildings, commercial facilities, industrial plants, transportation platforms, microgrids, and large-scale data centers, making fuel cells one of the most versatile clean-energy technologies available today.

Limitations and constraints of fuel cells

- **High system cost:** Despite significant technological progress, fuel cells remain more expensive than many conventional power-generation technologies. The cost challenge stems from complex stack manufacturing, specialized materials, balance-of-plant components, and relatively low production volumes compared with mature technologies such as internal combustion engines, gas turbines, and batteries. While costs have declined over the past decade, achieving widespread commercial adoption will likely require further manufacturing scale, supply-chain development, and technological improvements.
- **Limited hydrogen infrastructure:** One of the biggest challenges facing fuel-cell deployment is the lack of hydrogen production, transportation, storage, and refueling infrastructure. Unlike electricity, which is already widely distributed through established grids, hydrogen supply networks remain in the early stages of

development in most regions. This infrastructure gap creates a chicken-and-egg challenge, where end users are reluctant to adopt fuel cell technologies without reliable hydrogen availability, while infrastructure developers are hesitant to invest until demand reaches sufficient scale.

- **Durability and lifetime requirements:** Although fuel-cell durability has improved substantially, long-term reliability remains a key focus area, particularly for commercial transportation and utility-scale installations. Continuous exposure to heat, humidity, contaminants, and repeated operating cycles can gradually degrade fuel cell performance over time. Customers in sectors such as transportation, power generation, and industrial infrastructure often require operating lifetimes measured in tens of thousands of hours, creating significant engineering challenges for system developers.
- **Regulatory and safety considerations:** Hydrogen is a highly flammable gas and requires specialized handling, storage, and transportation systems. Although hydrogen technologies have well-established safety standards, new fueling stations, pipelines, and storage facilities must comply with stringent regulations. Developing these standards and obtaining approvals can extend deployment timelines and increase project costs.
- **Technology fragmentation and lack of standardization:** The fuel cell industry encompasses multiple technology platforms, including PEM, SOFC, MCFC, and other architectures, each optimized for different applications. While this diversity encourages innovation, it can also slow industry standardization and create uncertainty for customers evaluating long-term technology choices. Unlike the battery industry, which is increasingly converging around a few dominant chemistries, the fuel cell sector remains relatively fragmented.

Implications to datacenter power architecture

Transition from backup power to primary power source

Traditionally, data centers rely on grid electricity for primary power and diesel generators for backup power during outages. Fuel cells are increasingly being evaluated not only as a replacement for diesel backup generators but also as a potential primary power source. As AI-driven workloads substantially increase power density and overall electricity demand, hyperscalers are exploring fuel cells as a means of deploying computing capacity in locations where grid power is constrained or unavailable. This could fundamentally shift data-center power architecture from a grid-centric design to one that incorporates distributed on-site generation.

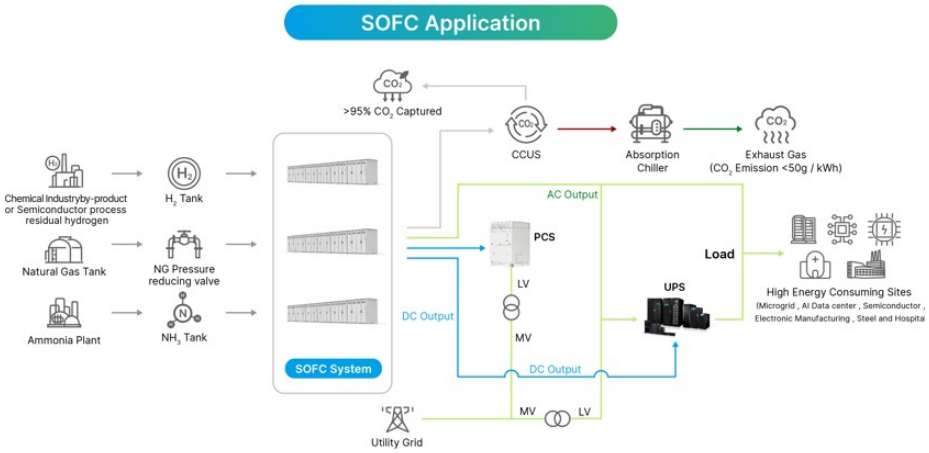
Potential long term evolution toward microgrid-based datacenters

In the longer term, fuel cells could become a core component of data-center microgrids combining renewable generation, battery storage, hydrogen storage, and intelligent energy management systems. Rather than functioning as passive consumers of grid electricity, future AI-focused data centers may increasingly operate as self-contained energy ecosystems capable of managing their own generation, storage, and load balancing. Such an architecture could provide greater scalability, resilience, and energy independence than traditional grid-connected designs

Potential simplification of power distribution architecture

Fuel cells inherently generate DC power, which aligns more closely with the power requirements of modern servers and IT equipment. Current data centers typically involve multiple AC/DC and DC/DC conversion stages between the utility grid and the processor, each introducing efficiency losses. Over time, fuel cell adoption could encourage a greater migration toward direct DC power distribution architectures, reducing conversion losses and improving overall facility efficiency. Moreover, its DC output fits well with the shift into HVDC power delivery architecture for datacenters.

Exhibit 65: Fuel cells in datacenter power delivery architecture



Source: Delta

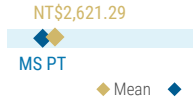
Risk Reward – Delta Electronics Inc. (2308.TW)

Strong combination: Power + cooling for AI infrastructure

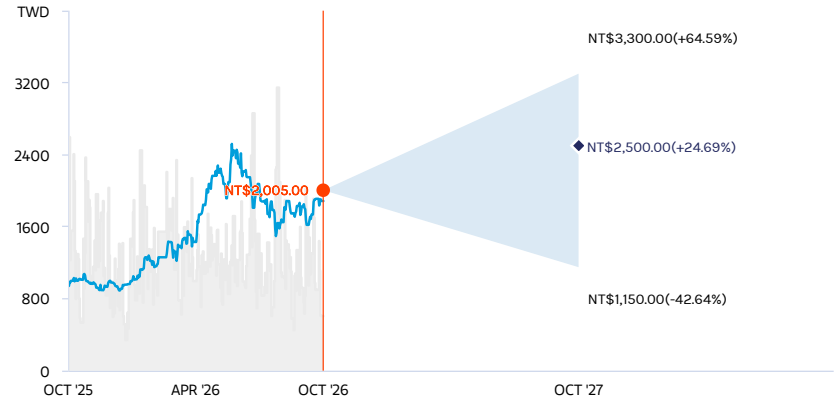
PRICE TARGET NT$2,500.00

Base case: P/E multiple of 30x based on average 2027-28 EPS. We believe our target multiple is justified by the 63% earnings CAGR we estimate for 2025-28, implying a PEG of 0.5, versus the stock's ~23x P/E and 23% earnings CAGR during 2022-25, which implied a PEG of 1.0.

Consensus Price Target Distribution



RISK REWARD CHART

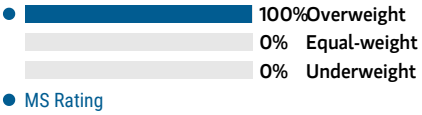


Key: Historical Stock Performance ● Current Stock Price ◆ Price Target

OVERWEIGHT THESIS

- We view AI server power supplies and power management/cooling solutions for data centers as key growth drivers for the coming decade.
- Our expectation for a re-rating reflects expanding gross/operating margins and increasing exposure to AI-related businesses.
- Our 2026-28 earnings estimates reflect strength in AI server power supplies and liquid cooling.
- Our price target implies a P/E of 30x based on average 2027-28 EPS.

Consensus Rating Distribution



Risk Reward Themes

Secular Growth:	*Positive*
Self-help:	*Positive*
Technology Diffusion:	*Positive*

View descriptions of Risk Rewards Themes here

BULL CASE NT$3,300.00

40x average 2027-28e P/E

Strong sales growth with significant margin expansion: We assume a greater-than-expected contribution from higher-margin businesses such as power electronics and liquid cooling as well as greater synergies from acquired companies.

BASE CASE NT$2,500.00

30x average 2027-28e P/E

Strong YoY revenue growth in 2026-28, coupled with operating leverage: We expect steady growth in AI server power supply and data center infrastructure to be the key drivers. We see favorable mix and scale merits driving profit to grow more than double in 2027 vs. 2025.

BEAR CASE NT$1,150.00

14x average 2027-28e P/E

Moderate sales growth with slight margin improvement: We assume further profit drag from legacy products, smaller profit contribution from higher-margin businesses, higher opex and/or less meaningful contribution from acquired companies.

Risk Reward – Delta Electronics Inc. (2308.TW)

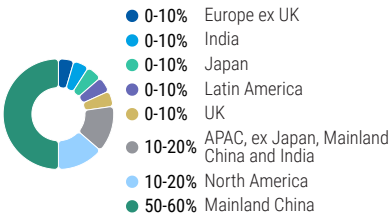
KEY EARNINGS INPUTS

Drivers	Dec 2025	Dec 2026e	Dec 2027e	Dec 2028e
Power Electronics YoY (%)	25.0	48.9	52.6	51.8
Automation YoY (%)	4.8	10.9	5.2	5.2
Infrastructure YoY (%)	81.6	44.6	51.0	48.2

INVESTMENT DRIVERS

- Transition to EV-related business and data infrastructure
- Operating margin recovery
- FINI holdings

GLOBAL REVENUE EXPOSURE



View explanation of regional hierarchies here

MS ALPHA MODELS

2/5 **MOST** 3 Month Horizon

the highest favored Quintile and 5 is the least favored Quintile

RISKS TO PT/RATING

RISKS TO UPSIDE

- Faster-than-expected global economic growth
- Falling raw material prices
- Margin improvement

RISKS TO DOWNSIDE

- Global economic slowdown
- Ongoing component supply tightness
- Rising raw material prices
- Margin contraction as a result of higher operating expenses or faster ASP erosion

OWNERSHIP POSITIONING

Inst. Owners, % Active 64.7%

MS ESTIMATES VS. CONSENSUS

FY Dec 2026e

Sales / Revenue (NT$, mn)	746,914	775,166 / 774,709	807,877
EBITDA (NT$, mn)	163,606	169,213 / 175,311	191,720
Net income (NT$, mn)	96,767	103,307 / 105,731	111,870
EPS (NT$)	37.25	39.77 / 40.50	43.07

◆ Mean ◆

Financial Summary – Delta

Exhibit 66: Delta: Financial Summary

Consolidated Income Statements

NT$mn (Year End Dec)	2024	2025	2026E	2027E	2028E
Net sales	**421,148**	**554,885**	**775,166**	**1,131,877**	**1,659,520**
COGS	(284,567)	(364,729)	(494,374)	(707,257)	(1,028,632)
Gross profit	136,580	190,157	280,792	424,620	630,887
Operating expenses	(88,928)	(106,224)	(142,491)	(191,135)	(263,905)
Operating income	**47,652**	**83,932**	**138,300**	**233,484**	**366,983**
Non-operating income	**3,664**	**3,934**	**10,365**	**7,100**	**7,300**
Interest income	1,859	1,501	1,687	2,100	2,300
Investment income	(162)	(200)	372	-	-
Exchange gain	1,041	644	2,091	-	-
Other	926	1,989	6,215	5,000	5,000
Pre-tax income	**51,316**	**87,866**	**148,666**	**240,584**	**374,283**
Income tax	(10,925)	(19,930)	(32,925)	(52,189)	(81,172)
Minority interests	(5,163)	(7,828)	(12,433)	(20,158)	(31,363)
Net income	**35,229**	**60,108**	**103,307**	**168,237**	**261,747**
Modelware EPS (NT$)	**13.56**	**23.14**	**39.77**	**64.77**	**100.77**

Consolidated Cash Flow Statements

NT$mn (Year End Dec)	2024	2025	2026E	2027E	2028E
Operating CF	**72,895**	**98,474**	**116,226**	**177,240**	**263,652**
Net Profits	35,229	60,108	103,307	168,237	261,747
Depreciation & Amort.	25,016	27,839	30,912	33,985	37,058
Change in WC	(2,486)	(24,190)	(55,485)	(86,295)	(125,847)
Other adjustments	15,137	34,717	37,491	61,313	90,694
Investing CF	**(40,360)**	**(53,214)**	**(70,000)**	**(70,000)**	**(60,000)**
Change of FA (Capex)	(33,435)	(45,950)	(70,000)	(70,000)	(60,000)
Change of L/T invest.	(4,739)	-	-	-	-
Other adjustments	(1,644)	(3,517)	-	-	-
Financing CF	**(11,118)**	**(9,401)**	**36,107**	**(12,414)**	**(54,050)**
Increase in L/T debt	8,552	(11,308)	12,530	2,978	3,829
Increase in S/T debt	(1,417)	446	29	212	272
Issuance of stock	-	-	-	-	-
Cash dividends	(18,696)	(20,313)	(30,131)	(51,787)	(84,335)
Other adjustments	444	21,774	53,680	36,183	26,183
FX adjust.	8,586	(2,146)	-	-	-
Net change in cash	**30,003**	**33,713**	**82,333**	**94,826**	**149,602**

Consolidated Balance Sheets

NT$mn (Year End Dec)	2024	2025	2026E	2027E	2028E
Cash	**117,459**	**151,172**	**233,505**	**328,332**	**477,934**
Mkt securities	4,333	3,941	3,941	3,941	3,941
A/R	89,921	121,717	170,036	248,282	364,023
Inventory	83,856	101,478	142,217	205,395	298,726
Others	15,355	17,473	24,410	35,642	52,257
Current Assets	**310,925**	**395,781**	**574,109**	**821,592**	**1,196,881**
Long-term investments	4,740	4,876	4,876	4,876	4,876
Fixed assets	115,710	142,040	127,820	127,652	124,410
Other assets	100,523	96,922	96,922	96,922	96,922
Total Assets	**531,898**	**639,619**	**803,727**	**1,051,042**	**1,423,089**
S/T borrowings	1,918	2,364	2,393	2,605	2,877
AP/NP	69,223	94,452	128,025	183,154	266,379
Other ST liabilities	81,142	112,849	157,649	230,195	337,504
Total current liabilities	**152,283**	**209,665**	**288,067**	**415,954**	**606,759**
L/T debt	32,209	21,143	33,673	36,650	40,480
Other LT liabilities	67,840	83,422	83,422	83,422	83,422
Total Liabilities	**252,332**	**314,230**	**405,162**	**536,026**	**730,661**
Common shares	25,975	25,975	25,975	25,975	25,975
Retained Earnings	102,429	140,925	214,101	330,552	507,964
Other SH equity	151,161	158,489	158,489	158,489	158,489
Shareholders' equity	**279,565**	**325,389**	**398,565**	**515,016**	**692,428**
Total Liab./SE	**531,898**	**639,619**	**803,727**	**1,051,042**	**1,423,089**

Consolidated Financial Ratios

	2024	2025	2026E	2027E	2028E
Margins (%)					
Gross margin	32.4	34.3	36.2	37.5	38.0
Operating margin	11.3	15.1	17.8	20.6	22.1
Pretax margin	12.2	15.8	19.2	21.3	22.6
Net margin	8.4	10.8	13.3	14.9	15.8
YoY growth (%)					
Sales	5	32	40	46	47
Operating profits	16	76	65	69	57
Pretax profits	5	71	69	62	56
Net profits	5	71	72	63	56
Others					
Cash div payout (%)	52%	50%	50%	50%	50%
Adjusted cash div (NT$)	7.0	11.6	19.9	32.5	50.5
Yield (%)	0%	1%	1%	2%	3%
Net Debt/Equity (%)	-31%	-40%	-51%	-57%	-63%
Liabilities/Equity (%)	90%	97%	102%	104%	106%
Liabilities/Assets (%)	47%	49%	50%	51%	51%
ROAE (%)	13%	20%	29%	37%	43%
ROAA (%)	7%	10%	14%	18%	21%
AR/NR days	73	70	69	67	67
Inventory days	103	93	90	90	89
AP/NP days	79	82	82	80	80
Cash conversion	97	80	77	77	77

Exhibit 67: Delta: Quarterly Financials

Delta (2308.TW) (NT$ mn)	1Q26	2Q26	3Q26E	4Q26E	1Q27E	2Q27E	3Q27E	4Q27E	2025	2026E	2027E
Net sales	**159,353**	**183,256**	**201,379**	**231,177**	**231,791**	**265,526**	**302,930**	**331,630**	**554,885**	**775,166**	**1,131,877**
COGS	(100,392)	(117,948)	(129,404)	(146,630)	(146,821)	(166,230)	(188,191)	(206,015)	(364,729)	(494,374)	(707,257)
Gross profit	**58,961**	**65,308**	**71,976**	**84,547**	**84,970**	**99,295**	**114,739**	**125,615**	**190,157**	**280,792**	**424,620**
Operating expenses	(30,544)	(34,738)	(36,983)	(40,226)	(40,422)	(45,405)	(50,589)	(54,719)	(106,224)	(142,491)	(191,135)
SG&A	(16,886)	(18,956)	(20,208)	(21,963)	(22,067)	(24,959)	(27,870)	(30,178)	(57,482)	(78,014)	(105,074)
R&D exp.	(13,657)	(15,782)	(16,775)	(18,263)	(18,356)	(20,445)	(22,720)	(24,541)	(48,743)	(64,477)	(86,061)
Operating income	**28,417**	**30,570**	**34,992**	**44,321**	**44,548**	**53,891**	**64,149**	**70,896**	**83,932**	**138,300**	**233,484**
Non-operating income	**2,676**	**4,430**	**1,630**	**1,630**	**1,775**	**1,775**	**1,775**	**1,775**	**3,934**	**10,365**	**7,100**
Interest income	428	398	430	430	525	525	525	525	1,501	1,687	2,100
Investment income	487	(115)	-	-	-	-	-	-	(200)	372	-
Disposal of investment	-	-	-	-	-	-	-	-	-	-	-
Exchange gain	1,688	403	-	-	-	-	-	-	644	2,091	-
Other	72	3,743	1,200	1,200	1,250	1,250	1,250	1,250	1,989	6,215	5,000
Pre-tax income	**31,093**	**35,000**	**36,622**	**45,951**	**46,323**	**55,666**	**65,924**	**72,671**	**87,866**	**148,666**	**240,584**
Income tax	(7,258)	(7,831)	(7,910)	(9,925)	(10,423)	(11,968)	(14,174)	(15,624)	(19,930)	(32,925)	(52,189)
Minority interest	(3,279)	(2,033)	(3,158)	(3,963)	(3,841)	(4,676)	(5,537)	(6,104)	(7,828)	(12,433)	(20,158)
Net income	**20,556**	**25,136**	**25,554**	**32,063**	**32,059**	**39,022**	**46,213**	**50,943**	**60,108**	**103,307**	**168,237**
EPS (NT$)	**7.91**	**9.68**	**9.84**	**12.34**	**12.34**	**15.02**	**17.79**	**19.61**	**23.14**	**39.77**	**64.77**
Margins (%)											
Gross Margin	37.0	35.6	35.7	36.6	36.7	37.4	37.9	37.9	34.3	36.2	37.5
Operating Margin	17.8	16.7	17.4	19.2	19.2	20.3	21.2	21.4	15.1	17.8	20.6
Pretax Margin	19.5	19.1	18.2	19.9	20.0	21.0	21.8	21.9	15.8	19.2	21.3
Net Margin	12.9	13.7	12.7	13.9	13.8	14.7	15.3	15.4	10.8	13.3	14.9
QoQ Growth (%)											
Sales	(1)	15	10	15	0	15	14	9			
Gross Profit	5	11	10	17	1	17	16	9			
Operating Profit	8	8	14	27	1	21	19	11			
Pretax Profit	21	13	5	25	1	20	18	10			
Net Profit	19	22	2	25	(0)	22	18	10			
YoY Growth (%)											
Sales	34	48	34	43	45	45	50	43	32	40	46
Gross Profit	56	48	37	51	44	52	59	49	39	48	51
Operating Profit	102	64	41	68	57	76	83	60	76	65	69
Pretax Profit	99	79	36	79	49	59	80	58	71	69	62
Net Profit	101	80	37	85	56	55	81	59	71	72	63

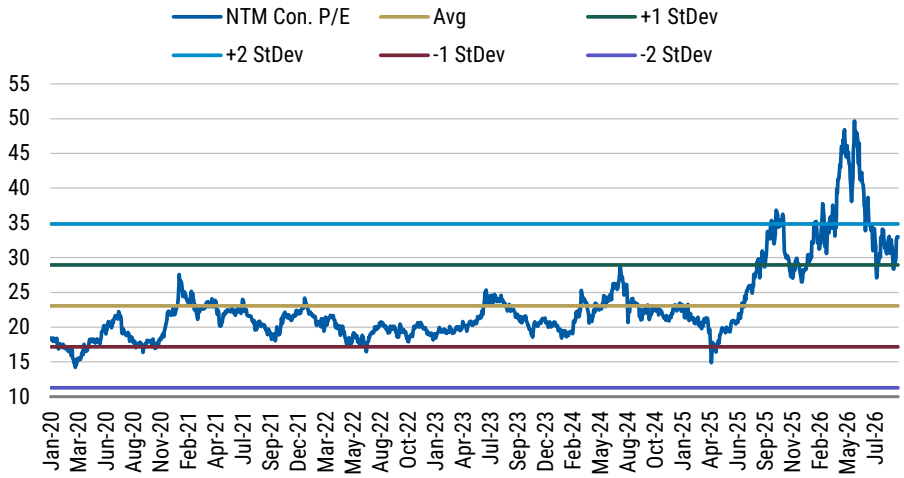
Valuation Methodology – Delta

Our price target (base case scenario value) is NT$2,500, derived from 30x 2027-28e P/E: We use price/earnings as our valuation methodology because we believe it appropriately captures Delta's earnings growth potential. We believe our target P/E multiple of 30x is justified by an earnings CAGR of 63% over 2025-28, which implies a PEG of 0.5. This is lower than the average of 1.0 in 2022-25, when the stock was trading at an average P/E of 23x for a 23% earnings CAGR.

Bull case scenario value – NT$3,300: We apply a 40x P/E multiple to our average base case estimate for 2027-2028. This scenario assumes a greater-than-expected contribution from higher-margin businesses such as power electronics and liquid cooling as well as greater synergies from acquired companies.

Bear case scenario value – NT$1,150: We apply a 14x P/E multiple to our average base case estimate for 2027-2028. This scenario assumes further profit drag from legacy products, smaller profit contribution from higher-margin businesses, higher opex and/or less meaningful contribution from acquired companies.

Exhibit 68: Delta: NTM P/E chart



Why Does Delta Trade at a Significant Valuation Discount to Delta Thailand?

Discussion of the valuation gap between Delta Electronics (2308.TW, currently 33x 12M forward consensus P/E) and Delta Electronics Thailand (DELTA.TH, 71x currently) has been ongoing for several years. If we look at their valuations before 2020, Delta was actually trading at a premium, with an average 17x 12M forward P/E in 2010-2020 vs. Delta Thailand's 13x. This made sense, as Delta is much more competitive in the global arena and Delta Thailand was set up as a subsidiary focused on manufacturing.

The situation flipped after 2020. Though Delta has been re-rated to trade at an average 23x 12M forward P/E since 2020, Delta Thailand has seen its valuation surge to an average of 60x over the same period. This cannot be explained by moves in their respective markets, either, as TWSE went from 14x forward P/E in 2010-20 to 16x after 2020, while the SET went from 13x to 16x.

We think this could be attributable to a few factors:

- **Different positioning within the Delta group:** Delta group has increased its stake in Delta Thailand to over 60% from about 20% in 2019 through a tender offer. Since then, there has been much more collaboration between Delta group and Delta Thailand. As Thailand emerged as one of the major manufacturing hubs within the Delta group, its earnings growth (28% in 2020-25) outpaced that of Delta Taiwan (19%).
- **Scarcity in the Thai stock market:** Delta Thailand is one of the relatively few large-cap Thai-listed companies that offer direct exposure to AI/data-center power and EV-related themes, whereas there are many such alternatives in Taiwan.
- **Limited free float:** As the Delta group owns a ~62% stake in Delta Thailand and there have been a few quite stable long-term shareholders, we think the effective free float is probably 20% or less.

Exhibit 69: Delta NTM P/E

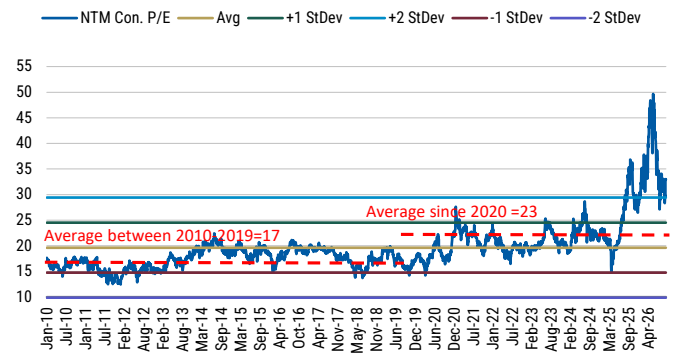


Exhibit 70: DET NTM P/E

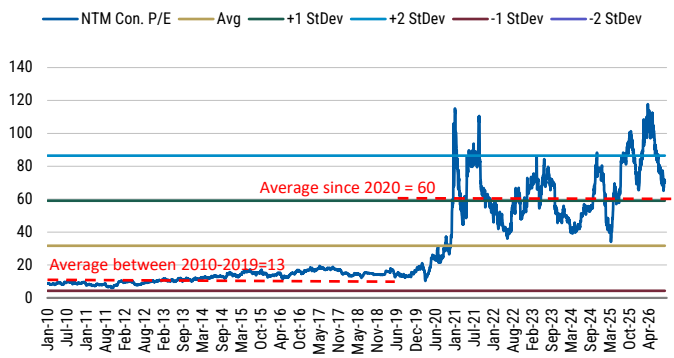


Exhibit 71: TWSE NTM P/E

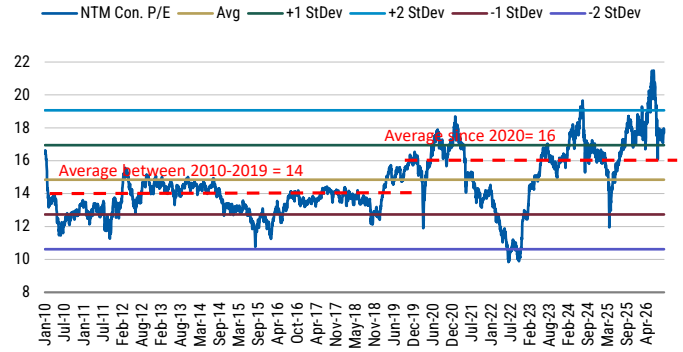


Exhibit 72: SET Index NTM P/E

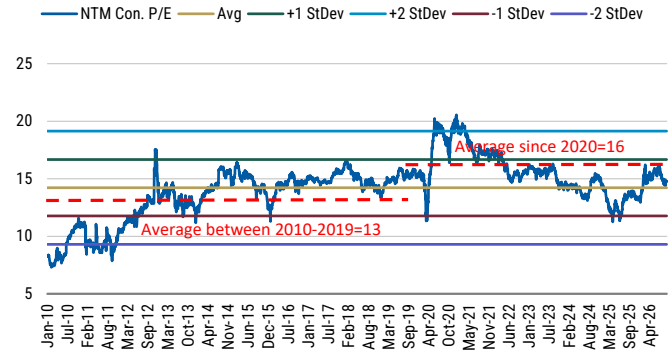
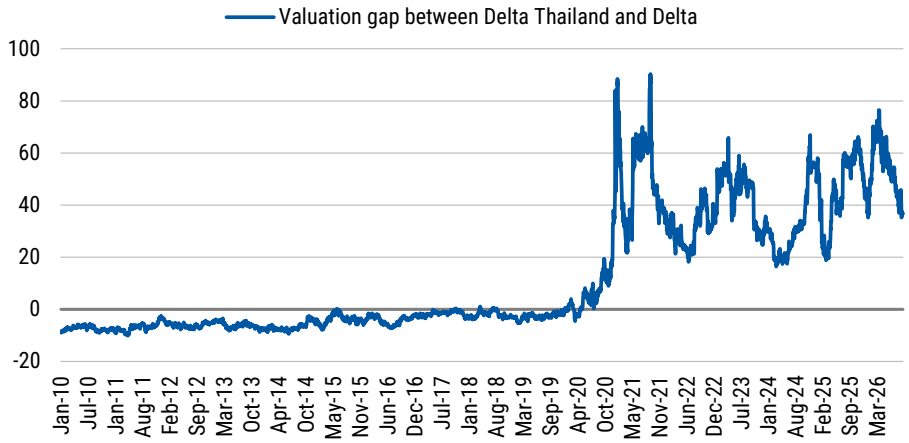


Exhibit 73:

Valuation gap between Delta Thailand and Delta based on 12M forward P/E



Global Footprint – Delta

Comprehensive Global Operations

Aggressive capacity expansion plans

The company has been continuously diversifying its production base away from China owing to the trade turmoil back in 2018-2019, and then by supply chain disruptions during Covid. Management explicitly set a target of lowering China production to 60% of total capacity (from >70% previously) for the first time back in 2020. New capacity was added in Thailand, India, and Taiwan during 2020-2021. It then saw demand surge from various sectors, including EVs, automation and telecom, during 2022-23, followed by AI-related power suppliers and cooling solutions emerging in 2024. In order to grasp the emerging opportunities, relevant plant build-outs were planned accordingly, on top of the ongoing geographical diversification. It expanded its presence—production sites, R&D, and sales teams—in Europe and the US to better cope with operational disruptions and capture domestic demand, in addition to its existing setups in Taiwan, Thailand, and India. There was then a substantial increase in capex of NT$12.2bn in 2025 (vs. NT$5.6-6bn in 2023-24) as the company further accelerated its expansion plan, mainly on AI-related power/cooling demand and increasing capacity tightness, specifically in Thailand and the US. Delta continued to expand in Thailand by adding three new plants, with another three under construction for power and liquid cooling system offerings. In view of the capacity tightness reserved for AI-related business, it also started to prioritize margin-accretive business—i.e., de-emphasize business with net margins below 10%—which further optimized product mix.

Looking ahead into 2026 and beyond, expansion is likely to remain aggressive in order to fulfill substantial power and liquid cooling turnkey solution demand for large-scale data centers, while management also indicated that the new capacity added at the end of 2025 has already been fully taken. With capacity staying tight for AI power and cooling offerings, the company will continue to expand actively in China, Thailand, the US, and Tainan, Taiwan. It also plans to add mass production lines in Taoyuan, scheduled to be ready by 2028, in addition to the pilot run lines in Tainan. Against this backdrop, the capex budget has been revised up several times and is now roughly NT$50bn for 2026.



Exhibit 74: Delta's global footprint



Financial Analysis – Delta

44% Revenue CAGR in 2025-28e

Delta's revenue growth was moderate in 2020 at merely 5%, making it the only year of the Covid disruption (2019-22) that did not see double-digit growth. Growth picked up to 11% in 2021, underpinned by surging demand for cloud, networking, and work-from-home infrastructure amid the pandemic, while the next uplift to 22% growth in 2022 was mainly fueled by increasing demand for EV- and data center-related components (server power supplies, networking switches, energy storage systems, etc). Revenue momentum subsequently moderated in 2023, growing just 4%, as the EV component business was impacted by the UAW strike in the US and automation demand in China stayed subdued, despite sustained demand for AI server power supplies and modules. Growth in 2024 was at a similar level, as strength in data center offerings were not enough to offset weakness in its legacy businesses.

Revenue surged 32% in 2025 to a record NT$554,885mn, driven by AI server rack power and cooling solutions, along with further spec upgrades and value growth.

We expect revenue to expand at a 44% CAGR over 2025-28, supported by ongoing AI server shipments, spec upgrades, 800 VDC power solutions, and turnkey liquid cooling offerings. As the upgrade to 800V HVDC power architecture is likely to take years given the complex design and high safety requirements, we expect AI server power solution upgrades to be separated into a few stages – from power shelves in AI server racks to standalone side-car power racks, from 50V DC to 800V DC output, from 72kW per power shelf to >250kW, etc, as well as power wattage increases for next-generation GPU rack designs.

By segment, Power Electronics has been the primary revenue driver, contributing at least 50% over the past years despite the company separating the Mobility segment from Power Electronics since 2024 so as to further distinguish revenue generated from EVs/ autos. We expect the segment to remain the key revenue contributor, supported by strong demand from AI server-related power solutions, and estimate it will reach 58% of revenue by 2028 vs. 50% in 2025. Infrastructure, the second-largest revenue contributor historically, is expected to grow along with the Power Electronics segment, fueled by demand for data center power-related solutions (i.e. UPSs, CDUs, sidecars, etc) as well as energy infrastructure-related offerings (i.e. SSTs (solid-state transformers), fuel cells, modular power solutions, micro grid management, etc.). We model the segment's revenue contribution to reach 35% by 2028 from 24% in 2024 and 33% in 2025. For Automation and Mobility, we expect their revenue contributions to decline sequentially throughout the years, with Automation dropping from 10% in 2025 to 4% in 2028 and Mobility falling from 7% in 2025 to 2% in 2028.

Exhibit 75: Delta: Revenue CAGR, 2025-28e

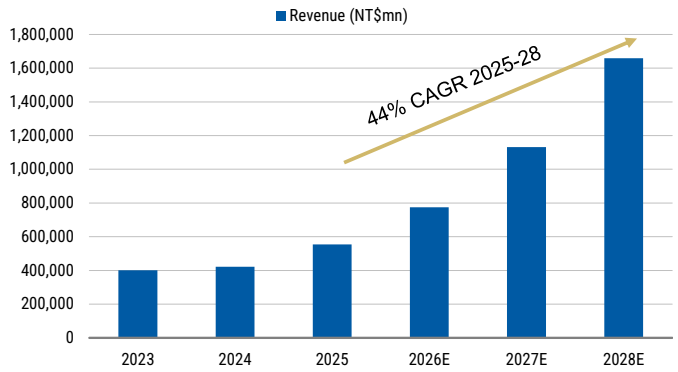


Exhibit 76: Delta: Annual revenue growth

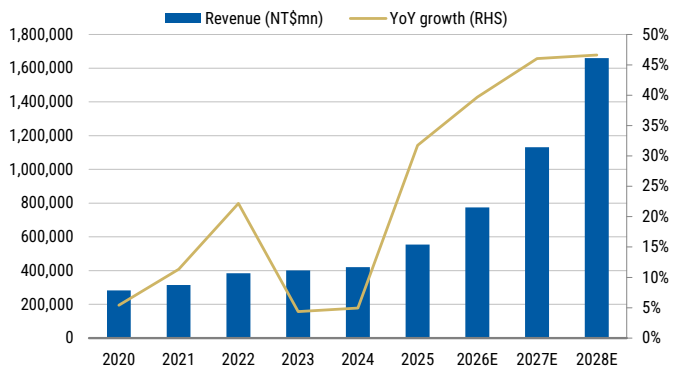


Exhibit 77: Delta: Power electronics revenue

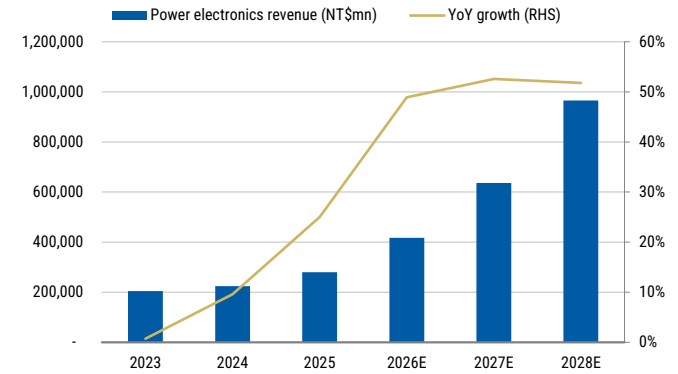


Exhibit 78: Delta: Infrastructure revenue

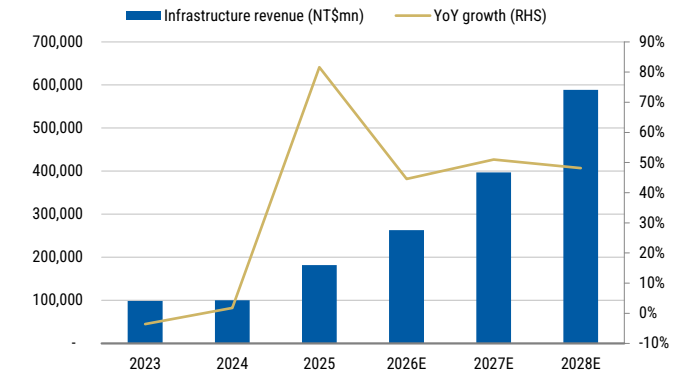


Exhibit 79: Delta: Automation revenue

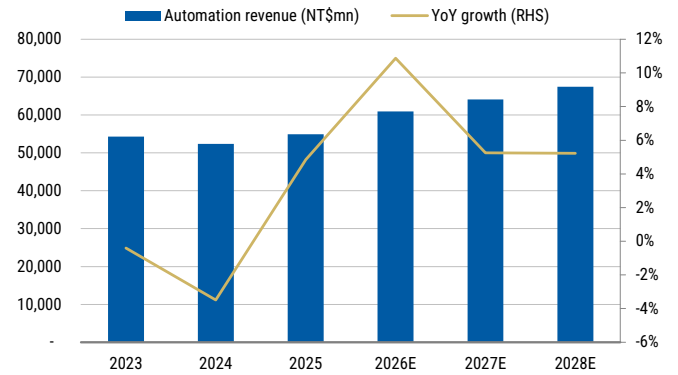


Exhibit 80: Delta: Mobility revenue

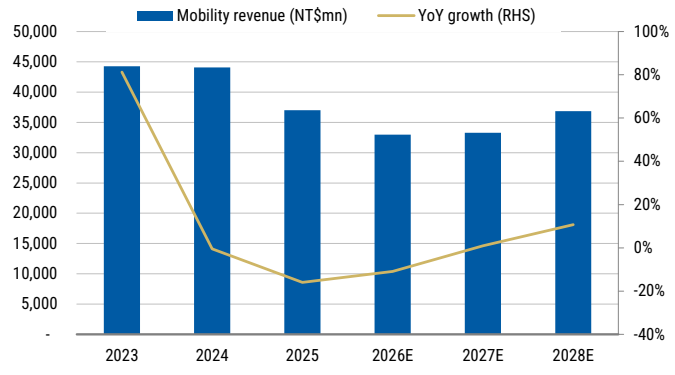
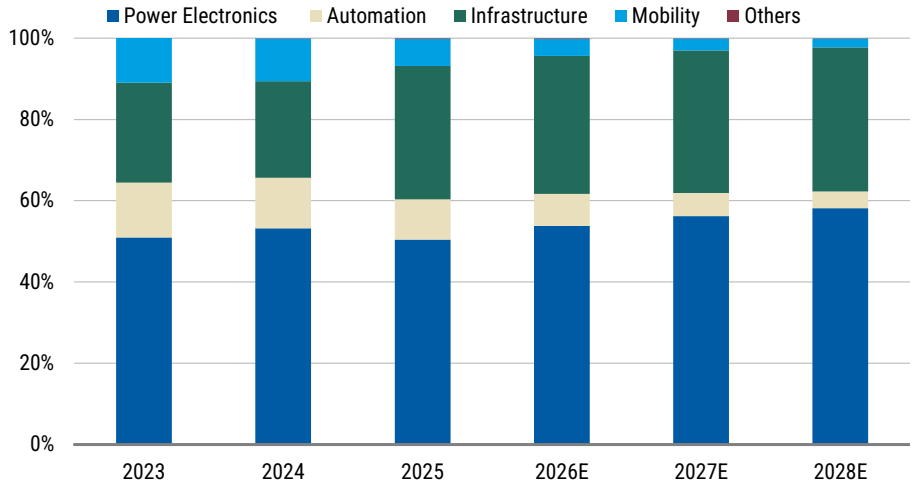


Exhibit 81:
Delta: Revenue breakdown by segment



Continuous margin expansion on favorable product mix

Delta's gross margin stayed largely between 28-29% during 2021-23, as revenue was mainly driven by relatively lower-margin products, including EV- and consumer/IT-related products, while server-related offerings had yet to scale up. 2024 served as an inflection point, with the product mix changing structurally, supported by robust demand for AI data center-related offerings, while its legacy businesses stayed relatively weak. Gross margin thus grew by 3.2ppt YoY during the year, from 29.2% in 2023 to 32.4% in 2024. Increasing AI computing needs further increased power consumption and cooling requirements per rack, boosting demand for high-density AI power and cooling products along with related system-level solutions in 2025, which lifted gross margin by another 1.8ppt YoY to 34.3% for the year.

The opex ratio was largely stable, fluctuating around 18-19% between 2020 and 2023. It then increased to 19% in 2023 and 21.1% in 2024, mainly due to operating deleverage, higher R&D costs, and ongoing capacity expansion. The ratio declined back to 19.1% in 2025 on the back of an enlarged revenue base, despite R&D-related investments remaining elevated.

Looking ahead, we expect gross margin to continue expanding on the back of increasing revenue contribution from AI-related power supplies. We thus pencil in a 2.0ppt gross margin expansion to 36.2% in 2026, followed by further increases of 1.3ppt and 0.5ppt YoY in 2027 and 2028, respectively, reaching 38.0% by 2028. On the opex ratio, we expect Delta to continue benefiting from operating leverage in the coming years, and thus estimate operating margin to improve more substantially than gross margin, expanding by 2.7ppt YoY in 2026 to 17.8%, followed by further expansions of 2.8ppt and 1.5ppt in 2027 and 2028, reaching 20.6% and 22.1%, respectively.

Exhibit 82: Delta: Margin trend

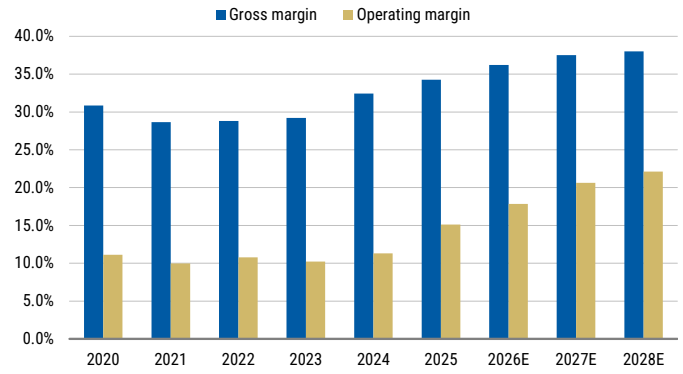
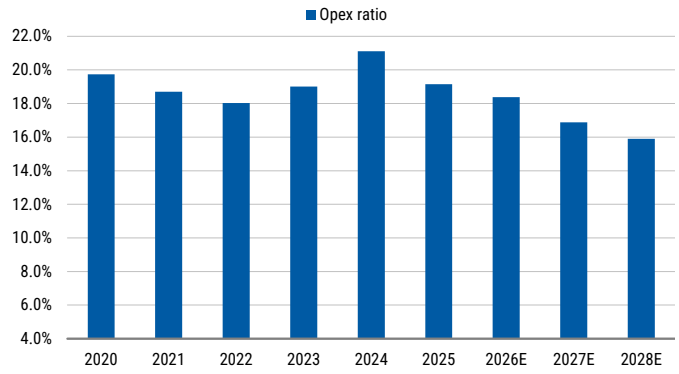


Exhibit 83: Delta: Opex ratio



Net profit to grow at a 63% CAGR in 2025-28

Delta's operating profit grew sequentially throughout 2020-22, supported by gross margin expansion and a lower opex ratio, from NT$31,426mn in 2020 to NT$41,439mn in 2022. It then saw a slight 1% YoY decline in 2023 due to higher operating expenses as the company further stepped up its R&D investments, despite a 0.4ppt gross margin expansion in 2023. Operating margin also edged down accordingly by 0.6ppt to 10.2% during the year.

Operating profit returned to a growth trajectory from 2024 onward, growing 16% YoY in 2024 and 76% YoY in 2025, as AI-related power offerings started to scale up. Operating margin expanded faster than gross margin, supported by operating leverage and disciplined cost controls, increasing by 1.1ppt in 2024 to 11.3% and another 3.8ppt to 15.1% in 2025.

With the tax rate staying largely stable at ~20% during 2020-23 before starting to edge up in 2024 and reaching 23% in 2025, Delta's net profit growth was largely aligned with its operating profit growth trajectory. 2023 was an exception, as net profit still managed to grow 2% YoY despite a 1% YoY decline in operating profit, mainly due to higher FX gains.

The net margin trend has been largely aligned with that of operating margin in recent years. After dropping from 9.0% in 2020 to 8.5% in 2021, as lower gross margin was offset by a lower opex ratio, net margin hovered around the 8.3-8.5% level during 2021-24. Net margin was largely flat YoY in 2024 despite a 1.1ppt expansion in operating margin, mainly due to lower non-operating income and a slightly higher tax rate. We then saw a surge in 2025, with net margin expanding 2.5ppt YoY to 10.8%, benefiting from a larger contribution from higher-margin AI power offerings.

Looking ahead, we forecast net profit to grow 72% YoY in 2026e, followed by 63% YoY growth in 2027e and 56% YoY growth in 2028e, implying a 63% CAGR in 2025-28e, assuming a stable tax rate of 22%. Net margin is likely to rise accordingly, from 10.8% in 2025 to 13.3% in 2026, followed by further expansion to 14.9% in 2027 and 15.8% in 2028.

Exhibit 84: Delta: Operating income trend

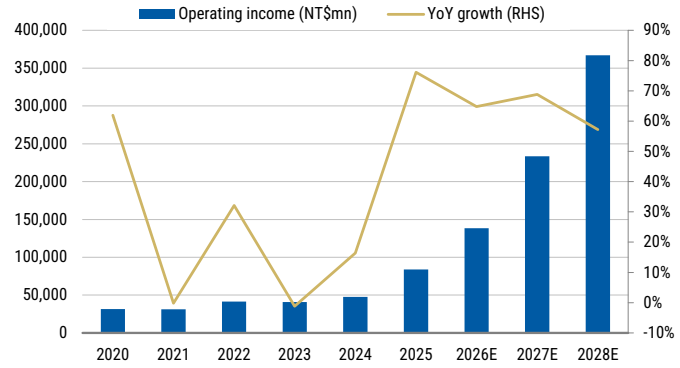


Exhibit 85: Delta: Operating margin

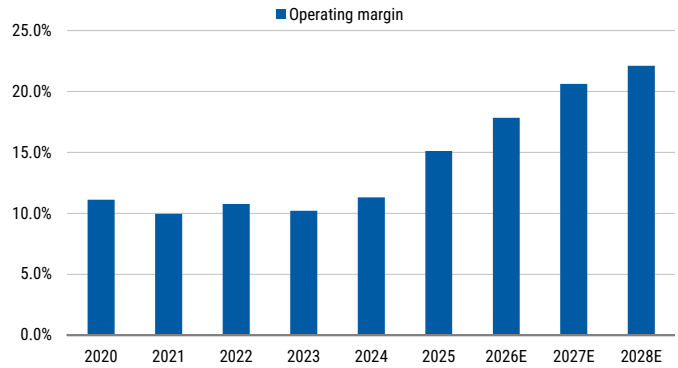


Exhibit 86: Delta: Net profit trend

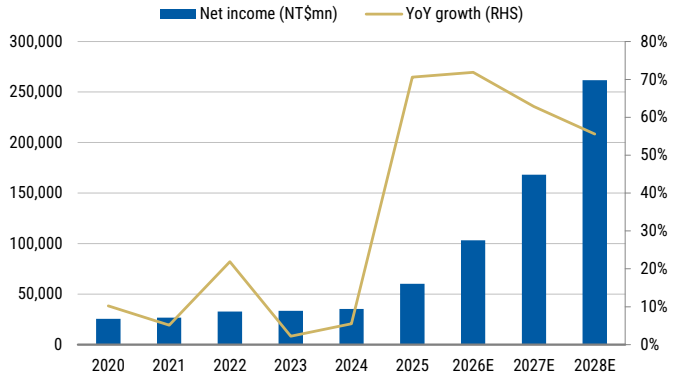
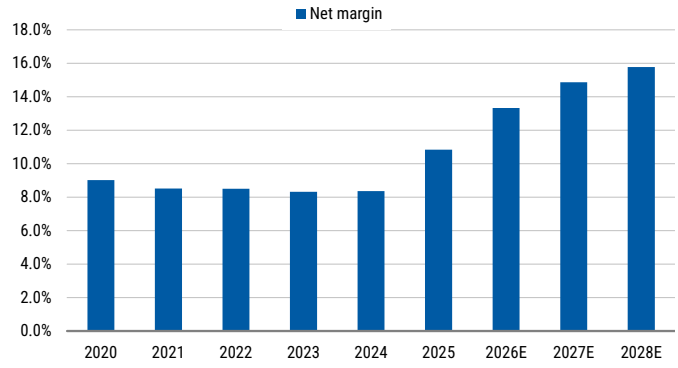


Exhibit 87: Delta: Net margin



ROE to rise alongside margin expansion

ROE rose gradually but steadily during 2020-22, from 14.3% to 14.5%, mainly driven by improving asset turnover, offset by lower net margin. ROE edged down to 13.8% in 2023 and fell further to 12.6% in 2024 on the back of lower asset turnover, while net margin and financial leverage remained at similar levels. It then rebounded to 18.5% in 2025, supported by both improving asset turnover and net margin expansion, which was fueled by an increasingly favorable product mix and operating leverage.

Looking ahead, we expect ROE to continue rising, reaching 25.9% in 2026, 32.7% in 2027, and 37.8% in 2028, primarily driven by continued margin expansion and improving asset turnover. Net margin is projected to expand from 10.8% in 2025 to 15.8% in 2028, supported by gross margin improvement—from 34.3% to 38.0%—and a declining operating expense ratio from 19.1% in 2025 to 15.9% in 2028.

Exhibit 88: Delta: ROE trend

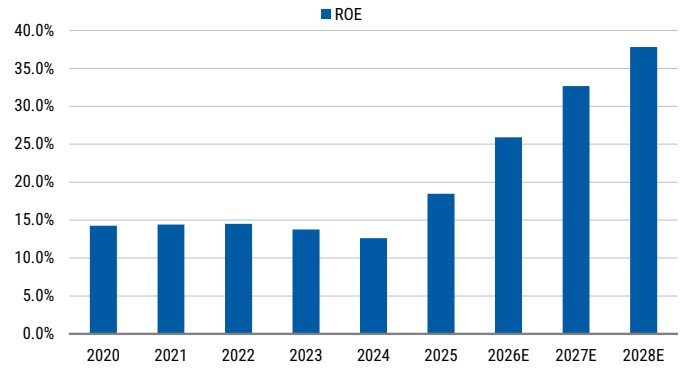
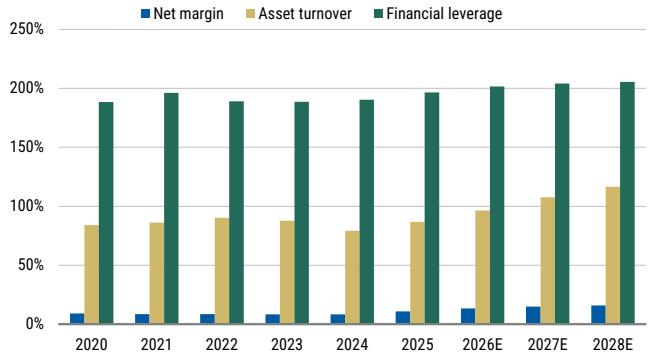


Exhibit 89: Delta: DuPont analysis



Improving balance sheet and continuous free cash flow generation

Delta has generated positive operating cash flow in recent years. Operating cash flow declined in 2021, mainly due to higher inventory buildup amid supply chain disruptions during the year. Free cash flow generation has also remained positive, while fluctuating with capital expenditures over the years. Free cash flow was exceptionally low at NT $5,292mn in 2021, as the company started to step up capex mainly for geographic capacity diversification and greater production flexibility—adding new capacity in Thailand, India, Tainan/Zhongli in Taiwan, and Wuhu, Wujiang, and Chongqing in Mainland China. Both absolute capex and the capex-to-sales ratio have trended upward since 2022, as Delta continued to add capacity for EV-related offerings and power and cooling solutions during 2023-24, with further acceleration in 2025 due to capacity tightness for AI power and cooling offerings, especially in Thailand and the US.

Looking ahead, the company aims to continue adding capacity worldwide to fulfill substantial demand for power and liquid-cooling turnkey solutions for large-scale data centers operated by major hyperscalers beyond 2027. Capacity remains tight for AI power and cooling offerings, so it continues to expand actively in China, Thailand, the US, and Tainan. Management also plans to add mass-production lines in Taoyuan, Taiwan, scheduled to be ready by 2028, in addition to the pilot-run lines in Tainan. Accordingly, we now estimate capex to rise to NT$70bn in 2026 from ~NT$46bn in 2025 and stay at a similar level in 2027 before edging down slightly to NT$60bn in 2028.

Delta's net cash position has been rising since 2021, with its net cash-to-equity ratio improving from just 2% in 2021 to 40% in 2025. We expect this ratio to continue rising, reaching 51% in 2026, 57% in 2027, and 63% in 2028, on strong free cash flow generation despite a higher total debt level amid ongoing capacity expansion.

Exhibit 90: Delta: Positive operating cash flow

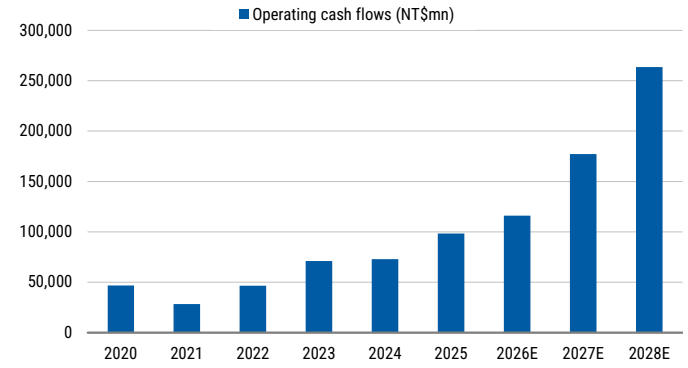


Exhibit 91: Delta: Improving free cash flow generation

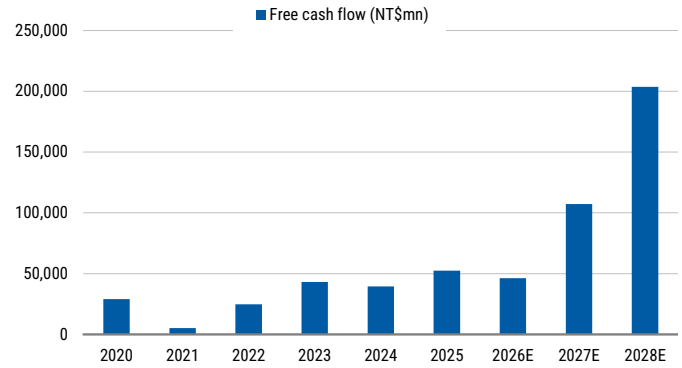


Exhibit 92: Delta: Capex and capex-to-sales ratio

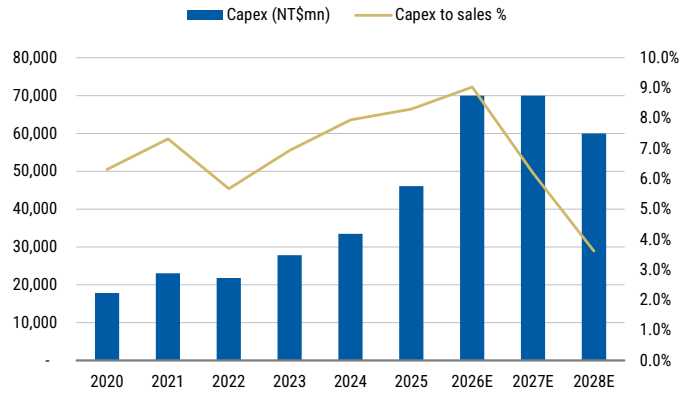


Exhibit 93: Delta: Net cash position

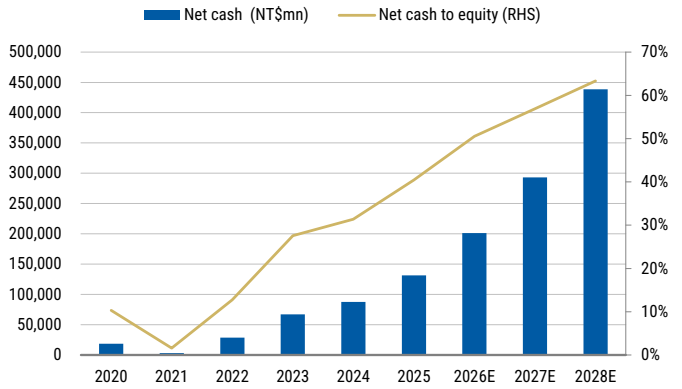


Exhibit 94: Delta: Debt position

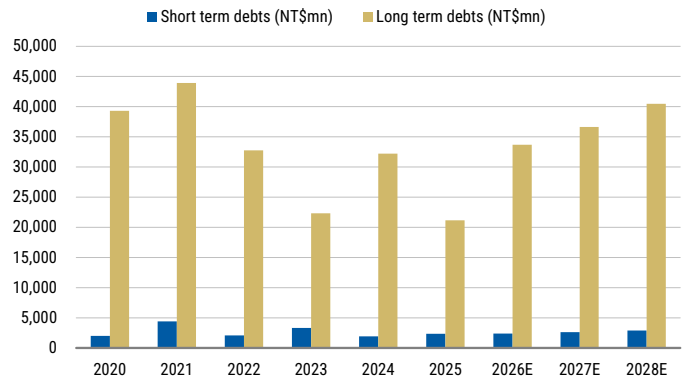
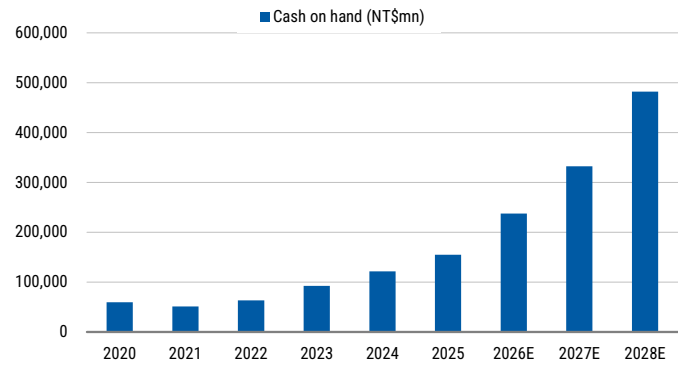


Exhibit 95: Delta: Cash on hand

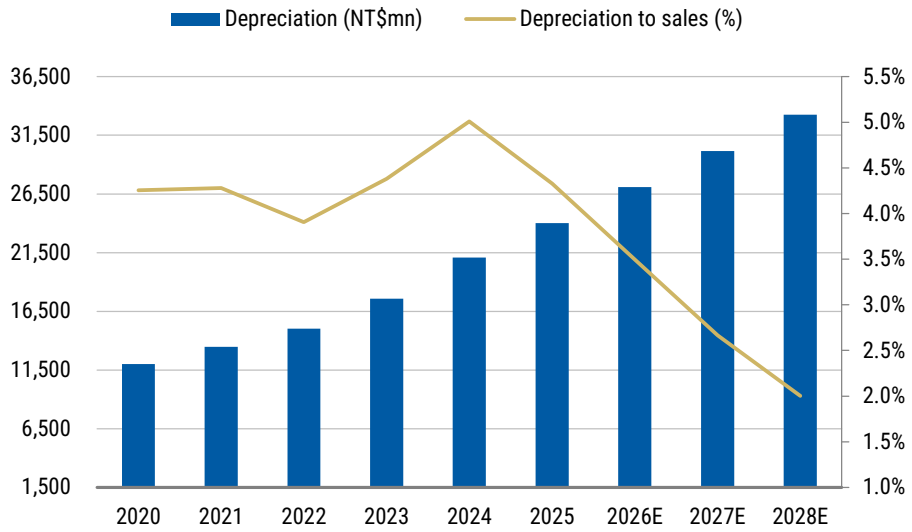


After edging down slightly in 2022 from a relatively high 7.3% in 2021, Delta's capex-to-sales ratio has been rising alongside its aggressive capacity expansion plans since 2022, despite an expanding revenue base. Looking ahead, as we expect capex to increase further in 2026, we project a capex-to-sales ratio of 9.0% for the year on robust revenue growth, before declining to 6.2% in 2027 and 3.6% in 2028.

The depreciation-to-sales ratio hovered around the 4-5% level during 2020-25. We

estimate the ratio will decline sequentially as the revenue base grows faster than capex and depreciation, from 4.3% in 2025 to 3.5% in 2026, 2.7% in 2027, and 2.0% in 2028.

Exhibit 96: Delta: Depreciation and depreciation to sales



Delta's cash conversion cycle (CCC) rose from 77 days in 2020 to a peak of 101 days in 2023 before improving to 80 days in 2025. By 2025, inventory days had increased from 80 days in 2020 to 93 days, while AP days had risen to 82 days after falling from 81 days in 2020 to 75 days in 2023. This was partly offset by improving AR days, which declined from 77 days in 2020 to 70 days in 2025.

Looking ahead, we expect the overall cash conversion cycle to improve further and remain at 77 days throughout 2026-28, mainly supported by improving AR days and a modest decline in inventory days.

Exhibit 97: Delta: AR days

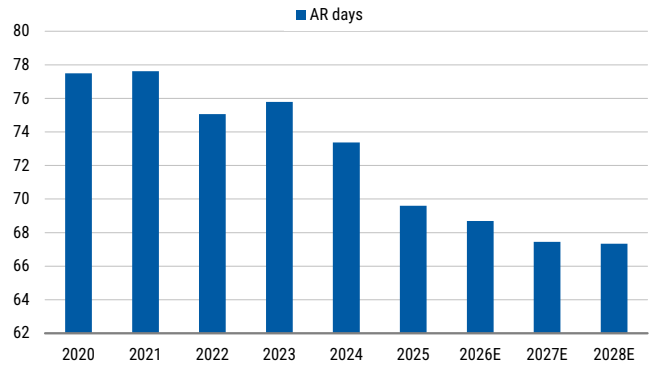


Exhibit 98: Delta: Inventory days

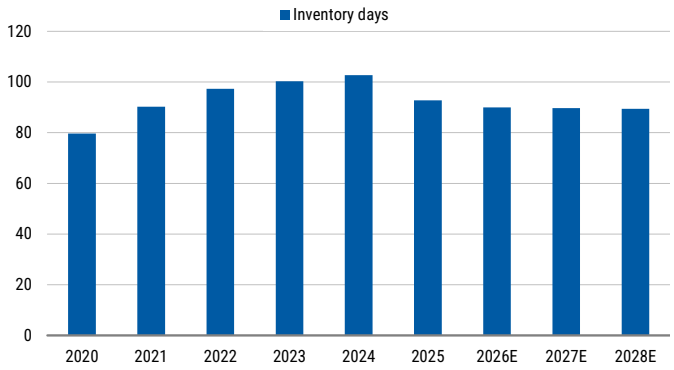


Exhibit 99: Delta: AP days

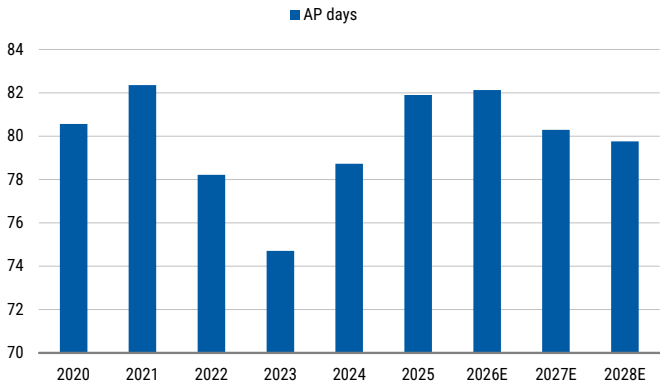
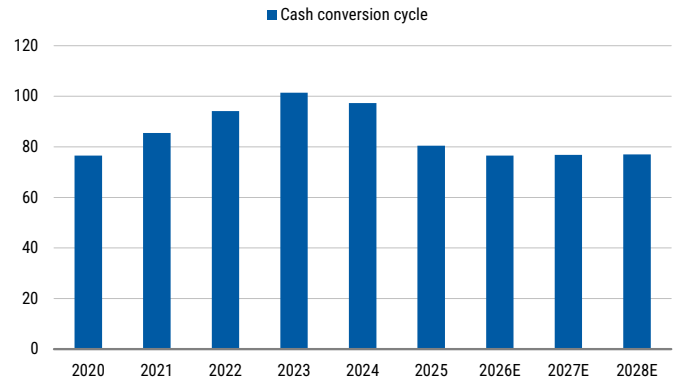


Exhibit 100: Delta: Cash conversion cycle



Company Background – Delta

Company snapshot

Delta Electronics is a global provider of power and thermal management technologies, with a focus on energy-efficient products and integrated solutions. Founded in 1971, the company serves a broad range of applications spanning components, power and system products, thermal management, mobility, industrial and building automation, data centers, telecom, and energy infrastructure, across industries including automotive, healthcare, data centers, and power grids. By combining its capabilities in power electronics, digital technologies, and system integration, Delta aims to support customers in improving operational efficiency and advancing smart manufacturing, smart cities, and more sustainable infrastructure.

The company is headquartered in Taipei, Taiwan. Its global network consists of 293 operations, including 165 sales offices, 55 manufacturing sites, and 73 R&D centers, supporting customers worldwide.

Exhibit 101: Delta: Key milestones

Year	Key Milestones
1971	Delta Electronics, Inc. was established
1976	Achieved >US$1mn of revenue
1977	Moved plant to Taoyuan Guishan Industrial Park
1983	- Mass production of switching power supplies
1988	- Listed on the Taiwan Stock Exchange (code 2308) - Mass production of brushless DC fans - Invested in the establishment of Delta Electronics (Thailand) PCL, with an original stake of 6%
1989	- Established first Japan office in Tokyo - Established Delta Power Electronics laboratory in Raleigh, US
1992	- Inauguration of the first plant in China (Dongguan) - Established Delta Greentech (China) Co., Ltd. in Shanghai
1995	Delta Electronics (Thailand) PCL was listed on Thailand Stock Exchange
1999	- Inaugurated its new Taipei HQ - Established the Delta Power Electronics Center (DPEC) in Shanghai
2000	Established the Delta Power Electronics Science and Education Development program
2001	Inaugurated its Wuhu plant in China
2003	Acquired ASCOM Energy Systems and its plant in Slovakia
2006	Inaugurated Delta's Tainan Plant I
2008	Inaugurated its Rudrapur plant in India and Wuhu plant in China
2010	Inaugurated Chenzhou plant in China
2011	Inaugurated Shanghai R&D and Operations Center and Taoyuan R&D Center in Taiwan
2012	Mr. Yancey Hai was elected Chairman and Mr. Ping Cheng was appointed as CEO
2015	- Acquired Eltek ASA, a leading provider of telecom power, industrial power, and data center power solutions - Inaugurated Delta Americas' headquarters green building and Taoyuan Plant 5
2017	- Announced organizational adjustment to focus on Power Electronics, Automation, and Infrastructure as main business categories - Acquired a controlling stake in VIVOTEK
2019	- Acquired a controlling stake in Delta Electronics (Thailand) Public Company Limited - Acquired lighting-solutions brand Amerlux
2020	Accquired Trihedral, a Canadian SCADA and IIoT software company, to strengthen industrial automation and system-integration businesses
2021	Acquired March Networks, a global video-surveillance provider, and Universal Instruments, a leader in precision automation solutions for electronics manufacturing
2022	- Inaugurated a new facility at the Automotive Campus in Helmond, the Netherlands - Purchased land and a factory in Plano, Texas, United States
2023	Acquired TB&C, a leading provider of automotive high-voltage hybrid components
2024	- Mr. Ping Cheng selected as Chairman while concurrently serves as CEO - Adjusted business categories into Power Electronics, Mobility, Automation, and Infrastructure
2025	- Inaugurated new North East Asia HQ in Tokyo - Acquired ACT Genomics and Noda RF Technologies

Product offerings

Delta's key offerings can be categorized into four key areas: Power Electronics, Infrastructure, Mobility, and Automation.

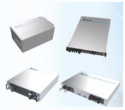
1. Power Electronics

Delta operates in switching power supplies, thermal management, and passive components, supplying products for ICT, cloud computing, EVs, and industrial and commercial applications. The company frames its business around energy efficiency and

sustainability and now also offers grid-to-chip power and thermal management solutions for data centers.

- Components: Includes essential passive and magnetic components used in portable devices, AI cloud computing, automotive electronics, and IoT.
- Power and System: Contains power products and systems for AI, cloud computing, IoT, client devices, and industrial, medical, lighting, e-mobility, and electrical-appliance applications.
- Fans and Thermal Management: This mainly covers air and liquid cooling solutions for AI cloud and edge computing, EVs, as well as industrial, commercial, and residential sectors. Key products includes DC and EC fans, thermal modules, liquid cooling cold plates, and coolant distribution units, etc.

Exhibit 102: Delta: Key offerings for the power electronics segment

POWER ELECTRONICS							
Components							
Inductors	RF Inductors	Standard Power Module (DC/DC Converter & AC/DC Module)	Transformers	Networking	EMI Filters	Solenoids	Current Sensing Resistor
				 Wire-Wound Common Mode Choke			
Power and System							
Embedded Power	Adapters	Industrial Power	Medical power	Industrial Battery Charging	USB Socket Outlets	High Voltage Power	Hydrogen Power
				 Modular Charging System			
Fans and Thermal Management							
Chip Level Cooling Solution	DC Brushless Fans and Blowers	Thermal Module	Cold Plate Loop	In-rack CDU	Rack Manifold	Electical Pump	In-row CDU
							
Communication Equipment Cooling	Energy Infrastructure and Industrial Automation Cooling	Portable Device and Equipment Cooling	Infrastructure Ventilation	Automotive Thermal			
 DC Brushless Fans and Blowers	 DC Brushless Fans and Blowers	 Laptop cooling	 DC Brushless Fans and Blowers				

2. Infrastructure

Delta supplies ICT and energy infrastructure to support demand related to AI, smart cities, and IoT. Its ICT infrastructure solutions target data centers and telecom networks, while its energy infrastructure offerings integrate AIoT, power conversion, and microgrid control across energy generation, storage, and management.

- ICT Infrastructure: It mainly provides infrastructure for AI computing and 5G

communications under this category, ranging from telecom cell sites to data centers. The portfolio includes power, air and liquid cooling, networking equipment, and management software.

- Energy Infrastructure: The company offers renewable energy, AIoT-based energy management, and power conversion technologies for low-carbon transformation. Solutions include solar inverters, wind power converters, energy storage, EV charging, and management platforms for smart grids. It also supplies semiconductor high-voltage power supplies, automated testing instruments, and variable-frequency drives for industrial processes.
- Display and Visualization: Delta's display technologies serve control rooms, public information systems, projection mapping, and exhibition or event spaces. Offerings range from large-scale and 8K projection systems to integrated video-wall solutions and can be customized for mission-critical operations.

Exhibit 103: Delta: Key offerings for the infrastructure segment

INFRASTRUCTURE							
ICT Infrastructure							
Uninterruptible Power Supplies (UPS)	Battery Backup System (BBU)	800 VDC In-Row Power	Power Distribution Unit (PDU)	Rack-Mount Remote Power Panel (rRPP)	Liquid and Air Cooling	Modular Rack & Rack Mount Accessories	Lithium-ion UPS Battery System
							
Ethernet Switches	Indoor/Outdoor Telecom Power Systems	Renewable Hybrid Power Systems	Small Cell Power Systems	Li-ion Battery Backup Systems	Campus/Enterprise Networks		
							
Energy Infrastructure							
EV Charger	Power Conversion System (PCS)	All-in-One ESS solution	Distributed ESS Solutions	Solar Inverter	Wind Power Converter	Solid State Transformer Platforms	Medium Voltage Drives
							
High Voltage Power	DC EV Charger ATS	Power Supply ATS	Solid Oxide Fuel Cell (SOFC)	Hydrogen Energy Solution	Active Power Filter (APF) / Static Var Generator (SVG)		
							
Display and Visualization							
Business and Education Multimedia Projector	High Performance Projector	DLP, LED and LCD Video Wall	Wireless Collaboration and Signage	Vivitek			
				vivitek			

3. Mobility

Delta applies its power electronics and motor technologies to smart mobility, supplying

power systems, drive technologies, and integrated powertrain solutions. It works with automakers to develop electric vehicles, positioning the business as part of the low-carbon transition and sustainable urban mobility.

- EV Powertrain System: The company's powertrain portfolio for EV makers includes On-board Charging Modules (OBCM), DC/DC converters, power electronics combos, and traction inverters, along with X-in-One integrated solutions.

Exhibit 104: Delta: Key offerings for the mobility segment

MOBILITY						
EV Powertrain System						
EV Power Electronics	Traction Products	x-EV Power Management	In-Vehicle Network and AI Driving Computing	Automotive Thermal	Low Voltage Powertrain	EV Charging
						

4. Automation

Delta provides automation solutions for industrial and building applications through integrated software and hardware. Its industrial automation offerings support smart production and operational efficiency, while its building automation solutions target energy-efficient, secure, and sustainable spaces.

- Industrial Automation: This segment provides automation solutions across industries by utilizing AI, big data, IoT, and digitized manufacturing models. Relevant offerings support diverse production needs, including customized, small-volume and high-variety production, as well as mass production.
- Building Automation: Delta's Building Automation Solutions use an IoT architecture and integrate technologies from five brands: Amerlux lighting, Delta Controls and LOYTEC IP-based building control, and March Networks and VIVOTEK surveillance. The management platform enables operators to manage energy use, occupant comfort, and sustainability goals through a single interface.

Exhibit 105: Delta: Key offerings for the automation segment

AUTOMATION							
Industrial Automation							
Inverters - AC Motor Drives	Drive System for Motion Control Solutions	Motors	Servo Systems - AC Servo Motors and Drives	CNC Solution	PC-Based Motion Control Solution	Linear Motion	Embedded Motion Controller
							
Industrial Ethernet Solution	Temperature Controllers	High Voltage Power	SCARA Robot	Electric Servo Press			
							
Building Automation							
Delta Controls	LOYTEC	Indoor Air Quality Monitor	LED Lighting - Industrial/ Road/Fishing/Commercial	Smart Surveillance - VIVOTEK	Healthy Lighting		
							

Monthly sales vs. share price

Delta's monthly sales show that seasonality has been an important factor over the years. Following a substantial increase in monthly sales since 2024, traditional seasonal patterns have become less pronounced as revenue from AI-related offerings has started to outweigh them, while liquid-cooling-related offerings have added a new layer of monthly sales volatility. The correlation between the company's monthly sales and share price movements has remained high, at ~87% since 2023.

The company reported August monthly sales of NT$64,586mn (-4% MoM/+35% YoY). We saw the Power Electronics segment decline sequentially from a high base, offset by growth in Automation and slight growth in the Infrastructure segment.

Exhibit 106: Delta: Monthly sales vs. YoY sales growth

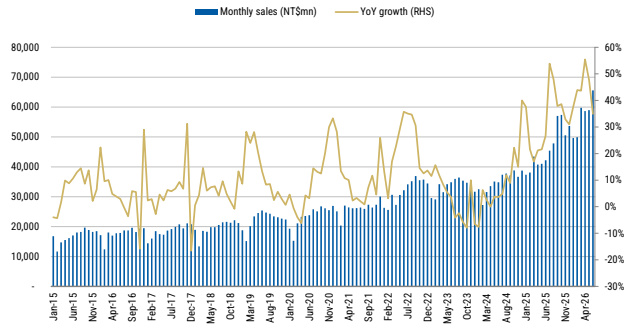
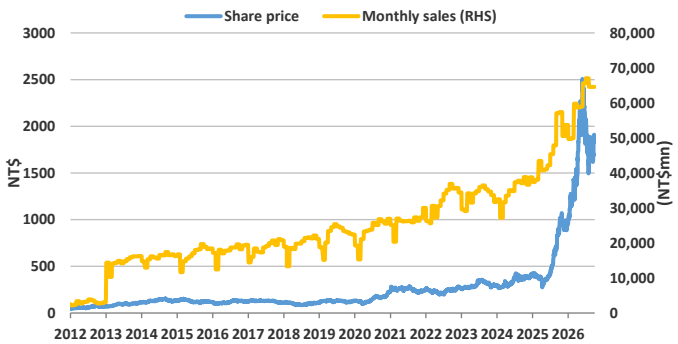


Exhibit 107: Delta: Monthly sales vs. share price



Key management & shareholding structure

Delta's board currently consists of eight directors and five independent non-executive directors. Its founder, Chung-Hua Cheng (Bruce Cheng), and his sons, Ping Cheng and An Cheng, together owned a 7% stake as of March 2026. Excluding the founder and his family, the management team held a stake of ~0.1% as of March 2026.

The FINI holding percentage as of July 14, 2026, was 63%, versus a previous peak of 78% in February 2011 and a historical average (since 2010) of 69%.

Exhibit 108: Delta: Shareholding structure

Shareholders	%
Dejico International LTD	10.30%
Deltron Holding LTD	8.40%
New Labor Retirement Pension Fund	3.75%
Chung-Hua Cheng (Bruce CH Cheng)*	3.15%
Government of Singapore-GOS-EFM C	2.28%
Ping Cheng*	2.14%
An Cheng (Victor Cheng)*	1.94%
Yuanta/P-shares Taiwan Top 50 ETF	1.43%
Labor Retirement Reserve Fund (The Old Fund)	1.42%
Labor Insurance Fund	1.32%
Total	**36.13%**

Exhibit 109: Delta: FINI holdings vs. share price since 2010



Exhibit 110: Delta: FINI holdings vs. share price since 2023

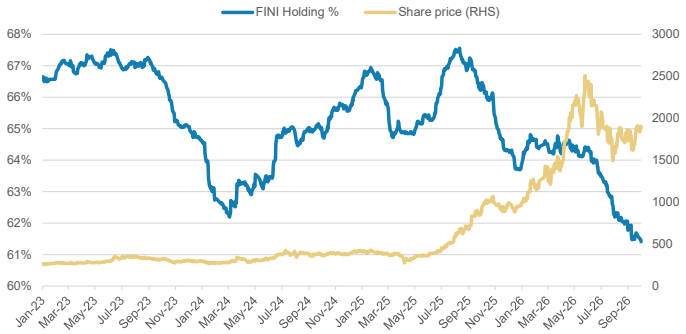


Exhibit 111:

Delta: Management profiles

Name	Title	Education	Work Experience
Bruce C.H. Cheng	Founder and Honorary Chairman Chairman, Delta Electronics Foundation	Bachelor of Electrical Engineering — National Cheng Kung University	- 2012–present — Honorary Chairman, Delta Electronics - 1990–present — Founded and chairs the Delta Electronics Foundation - 1971–2012 — Founder and Chairman, Delta Electronics - Before 1971 — Worked at TRW and Air Asia Co.
Ping Cheng	Chairman and CEO	Bachelor of Business Administration — California State University, East Bay	- Joined Delta in 1988 - 2024–present — Chairman - 2012–present — CEO of Delta Electronics - 2010–2012 — Corporate Senior Vice President and President of China Region - 2008–2010 — Senior Vice President & General Manager of Power Supply Business Group - 1998–2008 — Deputy General Manager of Component Business Group (Delta China); Executive Vice President of China Region
Mark Ko	Vice Chairman	Bachelor of Control Engineering — National Chiao Tung University	- Joined Delta in 1988 - 2012–present — Vice Chairman of Delta Electronics - 2004–2012 — President and Chief Operating Officer of Delta Group - 1994–2004 — General Manager of Power System Business Group
Simon Chang	Board Member President and COO	EMBA — Chung Yuan Christian University	- Joined Delta in 1981 - 2017–present — Board Member; President and Chief Operating Officer, Delta Electronics - Before 2017 — Senior Vice President and General Manager of Industrial Automation Business Group - From 1995 — Director of Industrial Automation Division
Victor Cheng	Board Member CEO of Delta Electronics (Thailand) PCL.	-Master of Electrical Engineering — Santa Clara University -Bachelor of Electrical Engineering — Santa Clara University	- Joined Delta in 1993 - 2024–present — Board Member of Delta; CEO of Delta Electronics (Thailand) PCL. - 2023–2024 — Executive Vice President of Infrastructure Business - 2017–2023 — General Manager of Information and Communications Technology Business Group - 2014–2017 — General Manager of Power System Business Group - 2002–2019 — President, then Chairman of Delta Networks - 1999–2002 — Head of Video Display Business Unit - 1993–1999 — International Sales Director to Assistant General Manager of Video Display Business Unit
Shan-Shan Guo	Board Member Chief Brand Officer Vice Chairman, Delta Electronics Foundation	- Master of Marketing and Communications — University of Michigan - Bachelor of Foreign Languages and Literatures — National Taiwan University	- Joined Delta in 2010 - 2011–present — Chief Brand Officer of Delta
Johnson Lee	Executive Vice President Infrastructure Business (IFB)	- EMBA — National Taiwan University - Bachelor of Industrial Engineering and Systems Management — Feng Chia University	- Joined Delta in 1985 - 2024–present — Executive Vice President of Infrastructure Business - 2012–2017 — Chief Operating Officer of Delta Electronics - 2004–2012 — General Manager of Power Supplies Business Group - Earlier — Procurement manager; General Manager of materials management
Ted Shyy	Executive Vice President Power Electronics Business (PEB)	Master in Materials Science and Engineering — National Cheng Kung University	- Joined Delta in 1988 - 2024–present — Executive Vice President of Power Electronics Business - 2014–present — Chief Executive Officer of Cyntec (concurrent role) - 2011–2014 — General Manager of Delta Components Business Group - 2003–2014 — General Manager of Cyntec - 1991–2003 — Rotation in Research and Development, production and marketing, under Cyntec, a joint venture between Delta and Japan's SUSUMU
James Tang	Executive Vice President Mobility Business (EVB)	- EMBA — National Taiwan University - Bachelor of Management Science — National Cheng Kung University	- Joined Delta in 1988 - 2024–present — Executive Vice President of Mobility Business; continues as General Manager of Electric Vehicle Solutions Business Group - 2017–present — General Manager of Electric Vehicle Solutions Business Group - 2008–2016 — General Manager across Power Supplies, Consumer and Lighting Solutions, and Embedded Power Supplies business groups - 1988–2008 — Procurement; Materials Manager; Director of Personal Computer Business Group
Jimmy Yiin	Executive Vice President Global Business Operations (GBO)	Master of Business Administration — National Taiwan University and Fudan University	- Joined Delta in 1992 - 2023–present — Executive Vice President of Global Business Operations - 2019–2022 — General Manager of Power System Business Group - 2014–2019 — Vice Chairman of Delta Greentech (China); General Manager of Delta Electronics (Shanghai) - 2008–2013 — General Manager of Delta Electronics (Jiangsu) - 1999–2007 — Purchasing manager/materials section chief of Delta Dongguan and Delta Networks
Rock Huang	Executive Vice President Global Manufacturing	Master of Industrial Engineering and Management — National Taipei University of Technology	- Joined Delta in 1987; ~40 years of experience in the manufacturing field - 2026–present — Executive Vice President of Global Manufacturing - 2020–2026 — President of Global Manufacturing - 2019–2020 — President of Delta China - 2018–2019 — Executive Vice President of Delta China - 2007–2018 — General Manager of Southern China, leading Dongguan team

Risk Reward – Lite-On Technology (2301.TW)

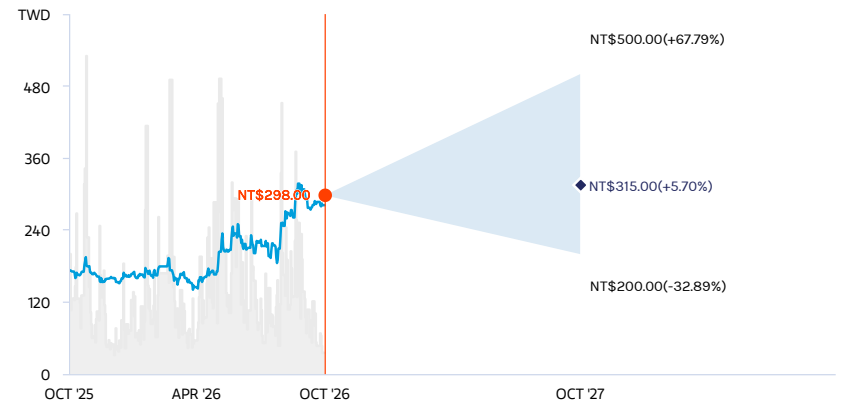
Stronger ASIC server power supply but fair risk-reward

PRICE TARGET NT$315.00

Base case: P/E multiple of 17x based on average 2027-28 EPS. We believe our target multiple is justified by the 47% earnings CAGR we estimate for 2025-28, implying a PEG of 0.4, versus the stock's 16x P/E and 2% earnings CAGR during 2022-25, which implied a PEG of 7.2.

Consensus Price Target Distribution NT$176.00 — NT$300.50 (MS PT, Mean) — NT$400.00

RISK REWARD CHART



EQUAL-WEIGHT THESIS

- We expect Lite-On Tech to benefit from strong AI server power demand in 2026-27, with the company gaining more traction on ASIC platforms.
- We are mindful of weakness in consumer/IT demand amid rising input costs, which could weigh on Lite-On Tech's profit in 2026-27.
- Liquid-cooling sidecar CDU system shipments may take more time to ramp up.

Consensus Rating Distribution

77%	Overweight
23%	Equal-weight
0%	Underweight

MS Rating

Risk Reward Themes

Contrarian:	*Negative*
New Data Era:	*Positive*
Technology Diffusion:	*Positive*

View descriptions of Risk Rewards Themes here

BULL CASE NT$500.00

27x average 2027-28e P/E

Better-than-expected margin expansion: Operating efficiency and accelerating production automation drive margin expansion, while new-business revenue comes in earlier than expected. Growth in cloud computing and automotive electronics is stronger than expected.

BASE CASE NT$315.00

17x average 2027-28e P/E

AI server projects progress, but some business uncertainties remain: We expect margin improvement from a better product mix as AI server power supply projects ramp up. However, we are mindful of weakness in consumer/IT demand amid rising input costs.

BEAR CASE NT$200.00

11x average 2027-28e P/E

Operating margin contraction: Operating efficiency improves less than expected, while pricing competition intensifies amid macro volatility; automotive electronics take longer to ramp up.

Risk Reward – Lite-On Technology (2301.TW)

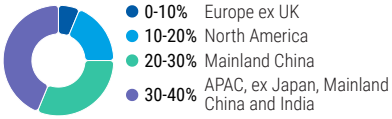
KEY EARNINGS INPUTS

Drivers	Dec 2025	Dec 2026e	Dec 2027e	Dec 2028e
Opto-electronics OPM (%)	8.3	7.9	7.5	7.5
Information Technologies OPM (%)	9.2	7.7	7.8	8.6
Opto-electronics revenue YoY (%)	0.2	8.0	1.7	2.3
Information Technologies revenue YoY (%)	7.7	(5.2)	0.5	(0.2)

INVESTMENT DRIVERS

- LED adoption (outdoor lighting/auto)
- Margin expansion
- Raw material prices, particularly for steel and copper

GLOBAL REVENUE EXPOSURE



View explanation of regional hierarchies here

MS ALPHA MODELS

2/5 **MOST** 3 Month Horizon

the highest favored Quintile and 5 is the least favored Quintile

RISKS TO PT/RATING

RISKS TO UPSIDE

- Stronger optical business growth
- Stronger data center/server demand
- PC/NB demand growth
- Downtrend in raw material prices

RISKS TO DOWNSIDE

- Macro demand weakness
- Weaker data center/server demand
- PC/NB demand slowdown
- Uptrend in raw material prices

OWNERSHIP POSITIONING

Inst. Owners, % Active 65.2%

MS ESTIMATES VS. CONSENSUS

FY Dec 2026e

Sales / Revenue (NT$, mn)	210,118	224,829 / 219,960	229,691
EBITDA (NT$, mn)	31,349	33,549 / 34,893	38,115
Net income (NT$, mn)	23,541	25,187 / 24,910	26,835
EPS (NT$)	10.30	11.07 / 10.88	11.79

◆ Mean ◆

Financial Summary – Lite-On

Exhibit 112: Financial Summary – Lite-On

Consolidated Income Statement

NT$ mn	2024	2025	2026E	2027E	2028E
Net sales	**137,134**	**166,085**	**224,829**	**297,964**	**357,670**
COGS	-107,504	-128,046	-168,703	-219,712	-259,785
Gross profit	**29,630**	**38,039**	**56,126**	**78,252**	**97,885**
Operating expenses	-16,648	-21,301	-26,044	-33,189	-37,645
- Promotion	-5,184	-7,629	-6,902	-8,782	-10,247
- ADM	-3,990	-4,965	-7,384	-9,340	-10,162
- R&D	-7,476	-8,721	-11,787	-15,037	-17,200
Operating income	**12,981**	**16,738**	**30,082**	**45,063**	**60,240**
Non-operating income	2,653	2,975	2,876	2,500	2,500
Interest income	1,491	1,119	799	800	800
Investment income	2,808	-956	1,436	1,200	1,200
Disposal of investment	75	-537	-77	0	0
Exchange gain	2,808	1,091	1,390	0	0
Other	-4,528	2,258	-672	500	500
Pre-tax income	**15,635**	**19,713**	**32,958**	**47,563**	**62,740**
Income tax	-3,674	-4,601	-7,767	-11,227	-14,793
Minority	19	-4	3	0	0
Reported net income	**11,942**	**15,116**	**25,187**	**36,337**	**47,947**
EPS (NT$)	**5.21**	**6.64**	**11.07**	**15.97**	**21.07**

Consolidated Cash Flow Statement

NT$ mn	2024	2025	2026E	2027E	2028E
Cash flow from Operations	**12,411**	**14,533**	**34,885**	**43,140**	**51,224**
Net profits	11,942	15,116	25,187	36,337	47,947
Depreciation + Amortization	4,047	3,668	3,994	4,065	4,137
Inv. Gain/Loss (Equity)	-15	-86	0	0	0
Disposal of Investment	9	568	0	0	0
Working Capital Change	-3,886	-2,099	-12,395	-12,092	-10,109
Other adjustments	313	-2,635	18,098	14,830	9,248
Cash flow from Investing	**-4,843**	**-7,395**	**-13,287**	**-8,993**	**-11,919**
Capex	-3,283	-6,795	-13,287	-8,993	-11,919
Change of LT investment	-2,234	-43	0	0	0
Change of ST investment	1,453	-417	0	0	0
Other adjustments	-779	-140	0	0	0
Cash flow from financing	**-3,363**	**-23,099**	**-19,866**	**-28,659**	**-37,818**
Increase in L/T debt	0	0	0	0	0
Increase in S/T debt	7,584	-10,116	0	0	0
Cash Dividend Paid	-10,372	-10,297	-19,866	-28,660	-37,818
Issuance of stock	0	0	0	0	0
Other adjustments	-574	-2,686	0	0	0
Exchange rate adjustment	3,736	-3,748	0	0	0
Net change in cash	**7,942**	**-19,709**	**1,732**	**5,487**	**1,487**

Consolidated Balance Sheet

NT$ mn	2024	2025	2026E	2027E	2028E
Cash	100,684	80,974	82,706	88,193	89,680
Mkt Securities	933	1,009	1,009	1,009	1,009
AR/NR	37,182	42,318	60,960	80,790	96,978
Inventory	26,817	34,262	42,083	54,807	64,804
Others	3,501	5,458	5,458	5,458	5,458
Current Assets	**169,117**	**164,021**	**192,216**	**230,257**	**257,928**
Long-term investments	8,195	7,748	7,748	7,748	7,748
Fixed assets	18,776	24,752	34,495	39,873	48,104
Deferred Asset	1,628	2,246	2,246	2,246	2,246
Intangible Assets	3,156	3,060	2,610	2,160	1,710
Other assets	3,811	3,860	3,860	3,860	3,860
Total Assets	**204,683**	**205,687**	**243,175**	**286,144**	**321,597**
S/T borrowings	30,187	20,029	20,029	20,029	20,029
AP/NP	43,126	53,608	67,676	88,138	104,213
Other ST liabilities	34,556	35,960	54,059	68,888	78,137
Total Current Liabilities	**107,869**	**109,597**	**141,764**	**177,056**	**202,379**
LT debt(include CB)	3,000	3,000	3,000	3,000	3,000
Other LT liabilities	2,903	3,111	3,111	3,111	3,111
Total Liabilities	**113,772**	**115,709**	**147,875**	**183,167**	**208,491**
Common shares	23,473	23,168	23,168	23,168	23,168
Retained Earnings	46,277	42,393	47,714	55,390	65,519
Other shareholders' equity	21,161	24,418	24,418	24,418	24,418
Total Equity	**90,911**	**89,979**	**95,300**	**102,976**	**113,105**
Total Liab. & Shldr's Equity	**204,683**	**205,687**	**243,175**	**286,143**	**321,596**

Consolidated Financial Ratios

	2024	2025	2026E	2027E	2028E
Margins					
Gross margin	21.6%	22.9%	25.0%	26.3%	27.4%
Operating margin	9.5%	10.1%	13.4%	15.1%	16.8%
Pretax margin	11.4%	11.9%	14.7%	16.0%	17.5%
Net margin	8.7%	9.1%	11.2%	12.2%	13.4%
YoY growth					
Sales	-7.6%	21.1%	35.4%	32.5%	20.0%
Operating profits	-10.6%	28.9%	79.7%	49.8%	33.7%
Pretax profits	-14.9%	26.1%	67.2%	44.3%	31.9%
Net profits	-18.0%	26.6%	66.6%	44.3%	32.0%
EPS	-18.1%	27.5%	66.6%	44.3%	32.0%
Net Debt/Equity	-75.3%	-65.5%	-63.7%	-64.3%	-59.8%
Liabilities/Equity	125.1%	128.6%	155.2%	177.9%	184.3%
Liabilities/Assets	55.6%	56.3%	60.8%	64.0%	64.8%
ROAE	13.6%	16.7%	27.2%	36.7%	44.4%
ROAA	6.1%	7.4%	11.2%	13.7%	15.8%
AR/NR Turnover (days)	92.2	87.4	83.8	86.8	90.7
AP/NP Turnover (days)	142.7	137.9	131.2	129.4	135.1
Inventory Turnover (days)	89.3	87.1	82.6	80.5	84.0
Cash conversion cycle (days)	38.9	36.5	35.2	37.9	39.6

Exhibit 113: Lite-on: Quarterly Financials

NT$ mn	1Q25	2Q25	3Q25	4Q25	1Q26	2Q26	3Q26E	4Q26E	2025	2026E	2027E
Net sales	**36,417**	**40,421**	**44,891**	**44,356**	**43,407**	**52,703**	**59,526**	**69,193**	**166,085**	**224,829**	**297,964**
COGS	-28,195	-31,479	-33,648	-34,724	-34,004	-38,355	-44,867	-51,477	-128,046	-168,703	-219,712
Gross profit	**8,222**	**8,942**	**11,243**	**9,633**	**9,403**	**14,348**	**14,659**	**17,716**	**38,039**	**56,126**	**78,252**
Operating expenses	-4,546	-5,215	-6,560	-4,980	-5,304	-6,146	-6,851	-7,743	-21,301	-26,044	-33,189
- Promotion	-1,532	-2,051	-2,535	-1,512	-1,500	-1,623	-1,786	-1,993	-7,629	-6,902	-8,782
- ADM	-998	-1,048	-1,792	-1,127	-1,457	-1,783	-1,964	-2,180	-4,965	-7,384	-9,340
- R&D	-2,023	-2,110	-2,244	-2,345	-2,341	-2,787	-3,095	-3,563	-8,721	-11,787	-15,037
Operating income	**3,675**	**3,727**	**4,683**	**4,653**	**4,099**	**8,203**	**7,807**	**9,973**	**16,738**	**30,082**	**45,063**
Non-operating income	994	536	886	560	871	1,166	445	395	2,975	2,876	2,500
Net interest income	424	323	184	188	214	195	195	195	1,119	799	800
Net investment income	596	-3,235	1,389	293	99	1,237	50	50	-956	1,436	1,200
Net gain of investment dispos	137	-702	23	5	1	-78	0	0	-537	-77	0
Exchange gain	379	2,540	-751	-122	284	1,106	0	0	1,091	1,390	0
Others	-542	1,609	41	195	272	-1,295	200	150	2,258	-672	500
Pre-tax income	4,669	4,262	5,569	5,213	4,970	9,368	8,252	10,368	19,713	32,958	47,563
Income tax	-1,214	-1,108	-924	-1,355	-1,195	-2,248	-1,939	-2,385	-4,601	-7,767	-11,227
Net income	**3,456**	**3,155**	**4,645**	**3,860**	**3,775**	**7,116**	**6,313**	**7,983**	**15,116**	**25,187**	**36,337**
EPS (NT$)	**1.51**	**1.39**	**2.05**	**1.70**	**1.66**	**3.14**	**2.77**	**3.51**	**6.64**	**11.07**	**15.97**
Margins											
Gross margin	22.6%	22.1%	25.0%	21.7%	21.7%	27.2%	24.6%	25.6%	22.9%	25.0%	26.3%
Operating margin	10.1%	9.2%	10.4%	10.5%	9.4%	15.6%	13.1%	14.4%	10.1%	13.4%	15.1%
Pre-tax margin	12.8%	10.5%	12.4%	11.8%	11.4%	17.8%	13.9%	15.0%	11.9%	14.7%	16.0%
Net margin	9.5%	7.8%	10.3%	8.7%	8.7%	13.5%	10.6%	11.5%	9.1%	11.2%	12.2%
QoQ Growth											
Sales	-4.9%	11.0%	11.1%	-1.2%	-2.1%	21.4%	12.9%	16.2%			
Gross profit	1.0%	8.8%	25.7%	-14.3%	-2.4%	52.6%	2.2%	20.9%			
Operating profit	5.1%	1.4%	25.7%	-0.6%	-11.9%	100.1%	-4.8%	27.7%			
Pre-tax profit	17.1%	-8.7%	30.7%	-6.4%	-4.7%	88.5%	-11.9%	25.6%			
Net profit	13.3%	-8.7%	47.2%	-16.9%	-2.2%	88.5%	-11.3%	26.5%			
YoY Growth											
Sales	26.6%	21.4%	22.1%	15.8%	19.2%	30.4%	32.6%	56.0%	21.1%	35.4%	32.5%
Gross profit	40.6%	21.0%	36.3%	18.3%	14.4%	60.5%	30.4%	83.9%	28.4%	47.5%	39.4%
Operating profit	64.8%	13.6%	19.3%	33.1%	11.5%	120.1%	66.7%	114.3%	28.9%	79.7%	49.8%
Pre-tax profit	49.2%	4.4%	25.6%	30.7%	6.4%	119.8%	48.2%	98.9%	26.1%	67.2%	44.3%
Net profit	44.7%	1.3%	37.0%	26.6%	9.3%	125.5%	35.9%	106.8%	26.6%	66.6%	44.3%

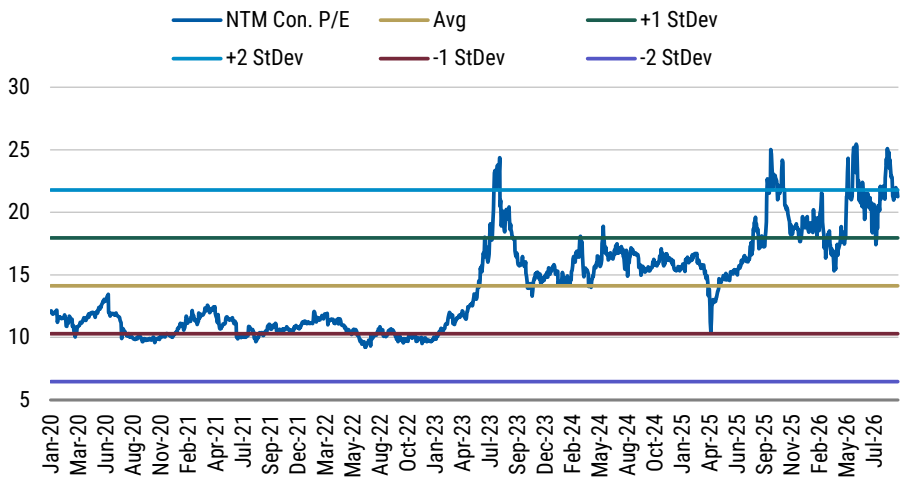
Valuation Methodology – Lite-On

Our price target (base case scenario value) is NT$315, derived from 17x 2027-28e P/E: We use price/earnings as our valuation methodology because we believe it appropriately captures Lite-On's earnings growth potential. We believe our target P/E multiple of 17x is justified by an earnings CAGR of 47% over 2025-28, implying a PEG of 0.4. This is lower than the average PEG of 7.2 in 2022-25, when the stock was trading at an average P/E of 16x for a 2% earnings CAGR.

Bull case scenario value – NT$500: We apply a 27x P/E multiple to our average base case estimate for 2027-28. This scenario assumes better-than-expected margin expansion, driven by improved operating efficiency and accelerating production automation, while new business revenue comes in earlier than expected. Growth in cloud computing and automotive electronics is stronger than expected.

Bear case scenario value – NT$200: We apply an 11x P/E multiple to our average base case estimate for 2027-28. This scenario assumes weaker-than-expected improvement in operating efficiency, while pricing competition intensifies amid macro volatility; automotive electronics take longer to ramp up.

Exhibit 114: Lite-on: NTM P/E



Global Footprint – Lite-On

More capacity expansion to come

Lite-On's production footprint spans the globe, with 15 sites in total, although it remains heavily weighted toward China. The company has been expanding aggressively, both to meet strong demand for AI power supplies and to support its strategy of geographic risk diversification and proximity to customers.

Looking at its expansion plans in recent years, the company has continued to demonstrate its commitment to diversifying production away from China and increasing proximity to customers. It completed its Thailand plant expansion and announced the second-phase expansion of its Vietnam facility in 2021. It then established automotive electronics and cloud power production lines in Dallas, Texas, in 2H22, along with the official launch of the second phase of its Hai Phong plant in Vietnam, which focuses on networking and consumer products. In 2023, it began expanding its second Kaohsiung plant and announced additional capacity expansion in Vietnam. In 2024, it committed a further US $120mn to establish a wholly owned subsidiary in Vietnam and began progressively shifting production of high-end power products to that facility. In 2025, the company aimed to double capacity as demand for AI server power supplies and battery backup units (BBUs) surged, focusing primarily on capacity expansion in the US, Kaohsiung, and Vietnam.

For 2026, the company expects capacity allocation of ~40% in China, 27% in Vietnam, 20% in Taiwan, and 10% in the US, with the remaining non-China footprint (i.e., Mexico, Brazil, and Thailand) accounting for ~3% of total capacity. It has also announced investments of US$919mn and US$149mn in its US and Vietnam subsidiaries, respectively, to continue expanding power-related capacity, mainly for PSUs. The capex budget for 2026 has also been raised to ~NT$13bn, versus ~NT$7bn in 2025.

Exhibit 115: Lite-on's global footprint



Source: Lite-on.

Financial Analysis – Lite-On

29% revenue CAGR in 2025-28

Lite-On's revenue growth was relatively moderate during 2020-22, with revenue expanding at a CAGR of 5% over the period. The company benefited from surging cloud computing demand, which drove increasing demand for power supplies, along with rising demand from EV/automotive electronics, although growth was offset by elevated component costs and supply-chain disruptions during COVID-19. Revenue then fell by 14% in 2023 and 8% in 2024 due to a post-pandemic demand pullback across the consumer electronics market and PC/NB model transitions, coupled with softening demand from EV-related businesses. The IT segment declined 20%/11% YoY, while the opto-electronics segment declined 11%/9% YoY in 2023/2024, respectively.

Revenue rebounded in 2025, growing 21% YoY, as the company began to benefit from surging demand for AI/cloud server power supplies. Contributions from ramping high-voltage/AI power infrastructure products began to increase, coupled with a partial recovery in the IT segment, mainly fueled by power adapter shipments for game consoles and gaming PCs/NBs, as well as related power offerings for low-Earth-orbit applications.

Looking ahead, we estimate revenue to expand at a 29% CAGR in 2025-28e, driven by the AI server power demand boom, while we remain mindful of weakness in consumer/IT demand amid elevated component costs.

By product segment, IT-related products were the primary revenue driver in past years, while 2025 marked a crossover point, with the contribution from the Cloud & AIoT segment surpassing that of the IT segment, mainly supported by AI power-related offerings. IT revenue contribution declined from 72% in 2020 to 38% in 2025, and we estimate it will reach 17% by 2028. For the Cloud & AIoT segment, revenue contribution grew substantially from 27% in 2021 (0% in 2020) to 45% in 2025; we estimate it will account for 74% of total revenue by 2028, on the back of surging AI power-related demand. For opto-electronics, we expect revenue contribution to decline from 17% in 2025 to 9% in 2028.

Exhibit 116: Lite-on: Revenue CAGR, 2025-28e

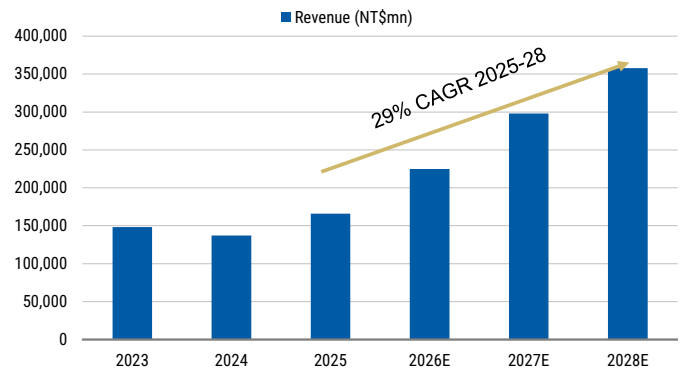


Exhibit 117: Lite-on: Annual revenue growth

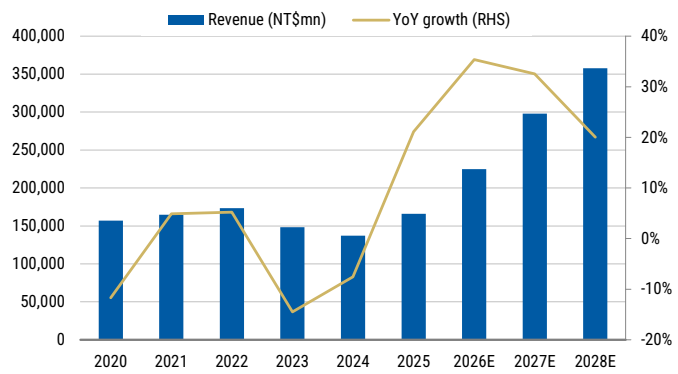


Exhibit 118: Lite-on: Cloud & AIoT revenue

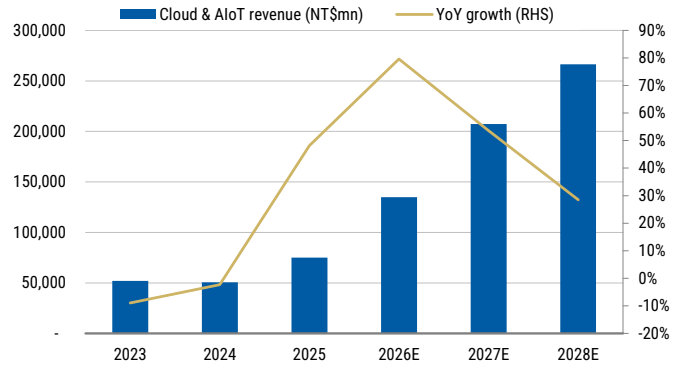


Exhibit 119: Lite-on: Opto-electronics revenue

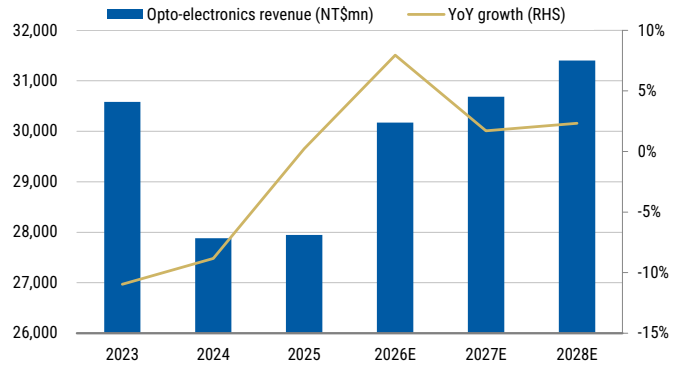


Exhibit 120: Lite-on: IT revenue

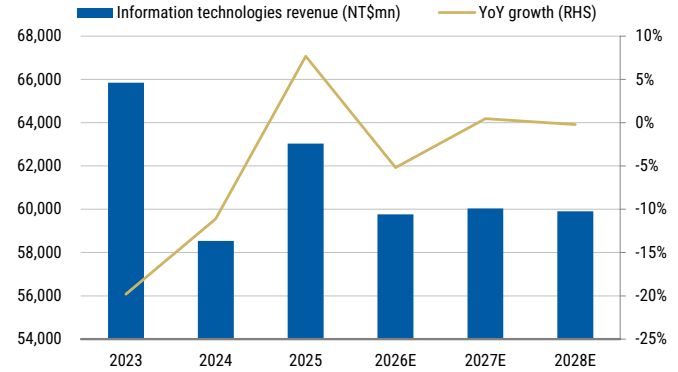
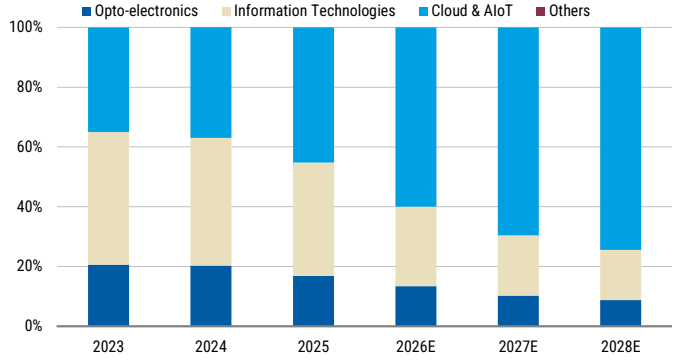


Exhibit 121: Lite-on: Revenue breakdown by segment



Continuing margin expansion

Lite-On's gross margin rose steadily over the years, from 17.4% in 2020 to 22.0% in 2023, thanks to a continually improving product mix. Gross margin dipped to 21.6% in 2024, mainly due to a smaller revenue base, while contributions from AI/cloud had yet to scale and margins were weighed down by learning-curve costs at the new Vietnam site. It then rose to 22.9% in 2025 as the AI server power business began to scale alongside rising BBU (battery backup unit) shipments, resulting in a more favorable product mix.

The opex ratio was largely stable at 10.5-11.0% during 2020-22. It then increased to 12.2% in 2023 and 12.1% in 2024, mainly due to operating deleverage and higher R&D costs. The opex ratio edged up further to 12.8% in 2025, mainly due to ongoing investment in R&D to support new business development in AI server power supplies, BBUs, liquid cooling, etc., along with a one-off tariff prepayment.

Looking ahead, we expect margins to continue expanding on a favorable product mix. We now pencil in a 2.1ppt gross margin improvement in 2026 to 25.0%, followed by expansions of 1.3ppt and 1.1ppt in 2027 and 2028, to 26.3% and 27.4%, respectively. For the opex ratio, we expect it to decline sequentially as the revenue base expands. Accordingly, we estimate operating margin to improve by 3.3ppt to 13.4% in 2026, followed by expansions of 1.7ppt in both 2027 and 2028, reaching 15.1% and 16.8%.

Exhibit 122: Lite-on: Margin trend

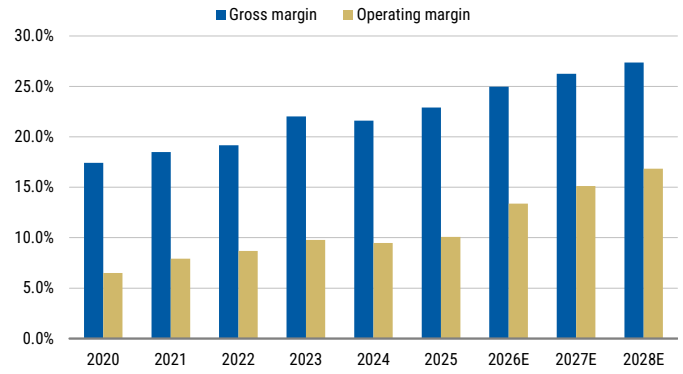
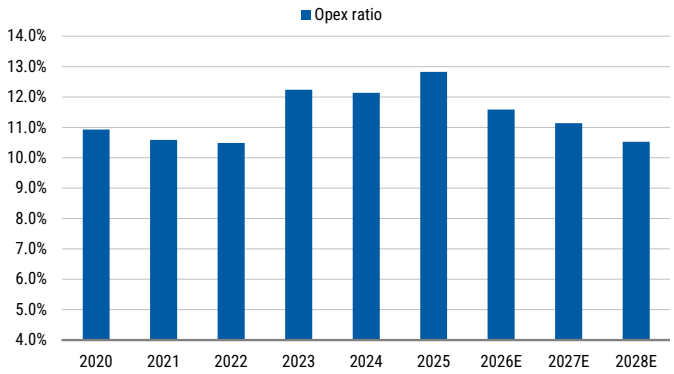


Exhibit 123: Lite-on: Opex ratio



Net profit to grow at a 47% CAGR in 2025-28

Operating profit grew steadily alongside gross margin expansion during 2020-22, from NT $10,206mn to NT$15,073mn. It then fell 4% YoY in 2023 and a further 11% YoY in 2024 due to higher operating expenses, mainly from R&D costs, despite gross margin expansion in 2023. Operating margin increased by 1.1ppt YoY to 9.8% in 2023 amid a smaller revenue base, before declining by 0.3ppt YoY to 9.5% in 2024 on lower gross margin and an elevated opex ratio.

Operating profit rebounded by 29% YoY to NT$16,738mn in 2025, supported by both an expanding revenue base and strengthening gross margin. Operating margin also improved, increasing by 0.6ppt YoY to 10.1%, as the recovery in gross margin was partly offset by higher opex.

With the tax rate remaining relatively stable at 20-24%, Lite-On's net profit growth has generally tracked its operating profit growth. An exception occurred in 2023, when net profit grew 3% YoY despite a 4% decline in operating profit, mainly due to higher FX gains during the year.

Net margin has broadly tracked operating margin in recent years. It reached a record high of 9.8% in 2023 on improving gross margin, despite a higher opex ratio. It then declined to 8.7% in 2024 on gross margin erosion before rebounding to 9.1% in 2025, benefiting from increasing demand for AI power supplies.

Looking ahead, we forecast net profit to grow 67% YoY in 2026e, followed by 44% YoY growth in 2027e and 32% YoY growth in 2028e, implying a 47% CAGR in 2025-28e. Net margin is likely to rise to 11.2% in 2026, 12.2% in 2027, and 13.4% in 2028.

Exhibit 124: Lite-on: 35% CAGR for operating profit in 2024-27

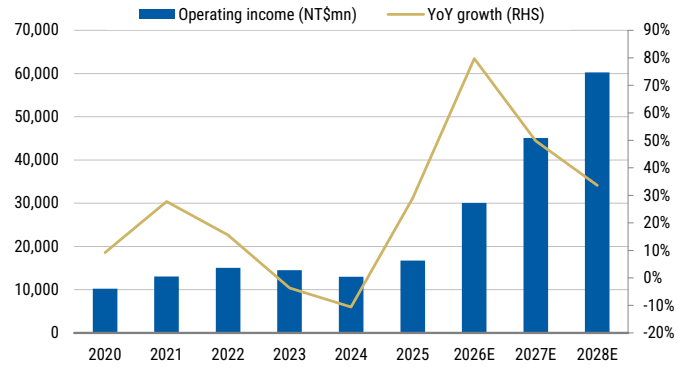


Exhibit 125: Lite-on: Operating margin

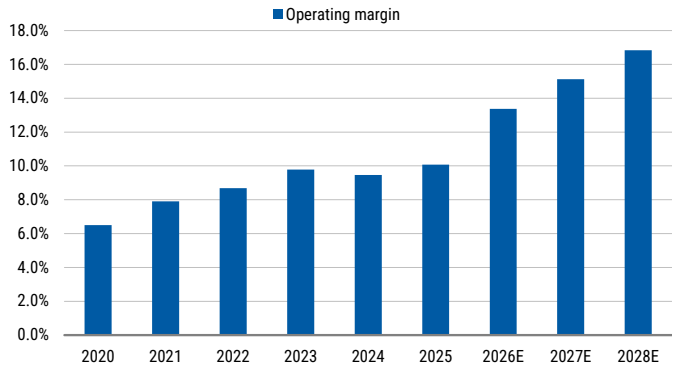


Exhibit 126: Lite-on: Net profit trend

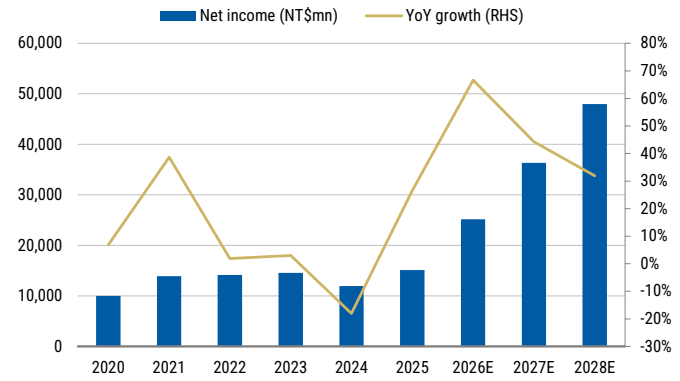
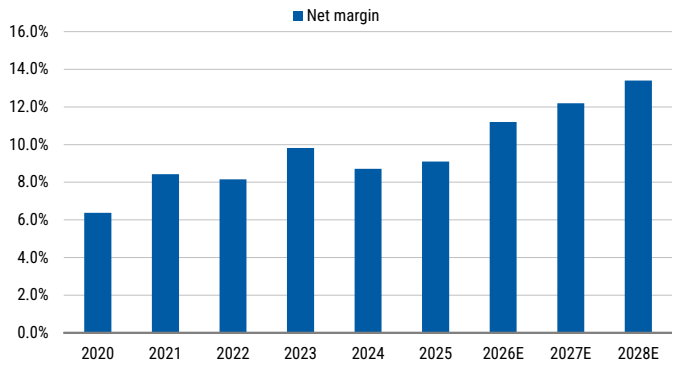


Exhibit 127: Lite-on: Net margin



ROE to rise

ROE movements were highly correlated with changes in net margin during 2020-22, as asset turnover and financial leverage were largely stable over the period. ROE edged down slightly in 2023 despite higher net margin, which increased from 8.2% to 9.8%, due to lower asset turnover and financial leverage. ROE fell further to a trough of 13.1% in 2024 as net margin and asset turnover declined further. It then rebounded to 16.8% in 2025, primarily supported by net margin expansion, fueled by a favorable product mix, along with an improvement in asset turnover back to the 80% level.

Looking ahead, we expect ROE to continue rising, reaching 26.4% in 2026, 35.3% in 2027, and 42.4% in 2028, primarily driven by continued margin expansion and improving asset turnover. Net margin is projected to increase from 9.1% in 2025 to 13.4% by 2028, supported by gross margin improvement from 22.9% to 27.4% and a decline in the operating expense ratio from 12.8% in 2025 to 10.5% in 2028.

Exhibit 128: ROE trend

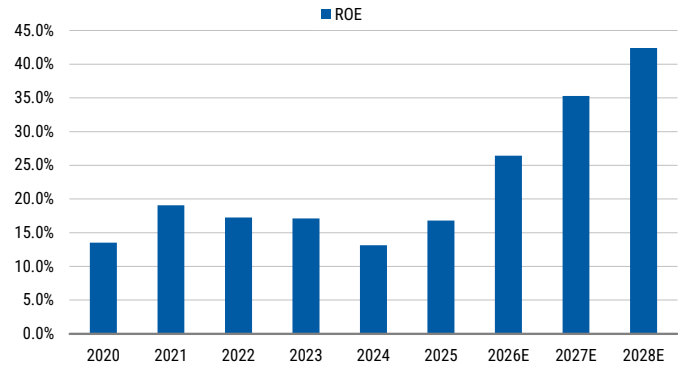
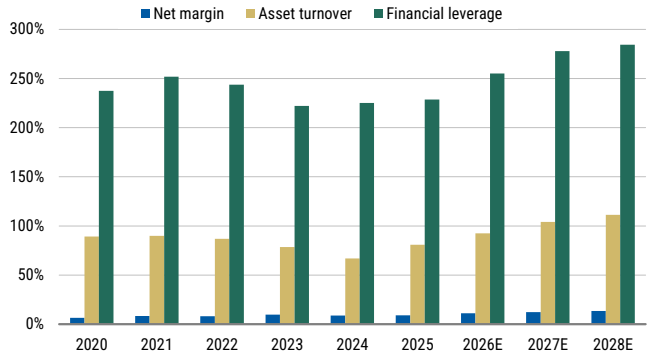


Exhibit 129: DuPont analysis



Improving balance sheet and continuous free cash flow generation

Lite-On has consistently generated positive operating cash flow in recent years, although cash flow dropped notably in 2021, 2024, and 2025, mainly due to higher working-capital-related cash outflows. The company has also consistently generated positive free cash flow, with capital expenditures generally at the NT$3-4bn level over the years. An exception occurred in 2025, when capital expenditure rose to NT$6.7bn, mainly for power-related capacity expansion in Vietnam, Taiwan, and the US.

As the company has announced investments of US$919mn and US$149mn in its US and Vietnam subsidiaries, respectively, alongside continued expansion of power-related capacity, mainly for PSUs, we now estimate capex to rise to ~NT$13bn in 2026, versus ~NT $7bn in 2025. Capex should then decline to NT$9bn in 2027 as investment levels normalize.

Lite-On's net cash position has remained relatively healthy, with its net cash-to-equity ratio improving from 65% in 2020 to 80% in 2023. It then declined to 75% in 2024 and 66% in 2025, as total debt increased by NT$7.7bn in 2024, while cash on hand declined in 2025, likely due to increasing capital investment.

Looking ahead, we assume the total debt position will remain stable despite ongoing capex investment. The net cash position should continue to grow sequentially on the back of positive free cash flow generation.

Exhibit 130: Lite-on: Positive operating cash flow

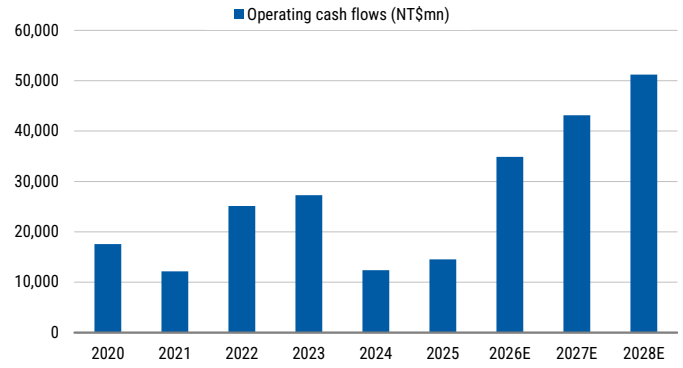


Exhibit 131: Lite-on: Improving free cash flow generation

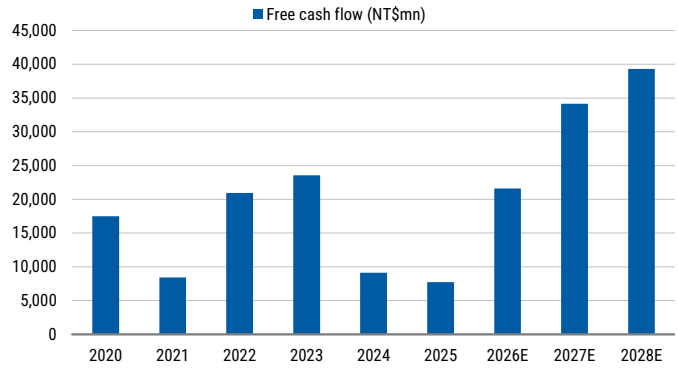


Exhibit 132: Lite-on: Capex and capex-to-sales ratio

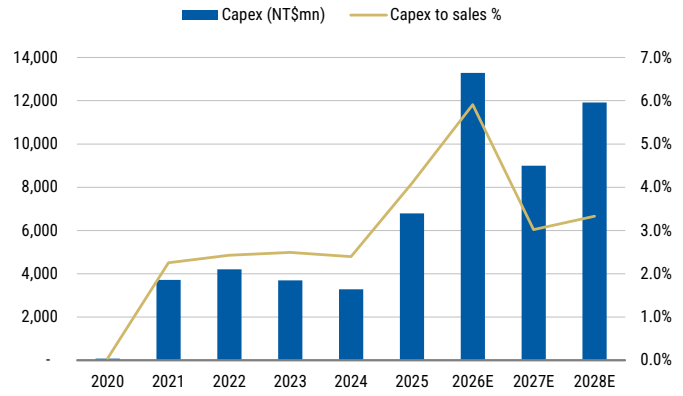


Exhibit 133: Lite-on: Net cash position

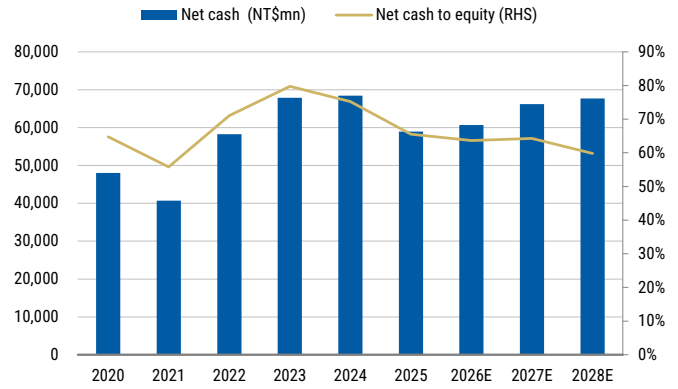


Exhibit 134: Lite-on: Debt position

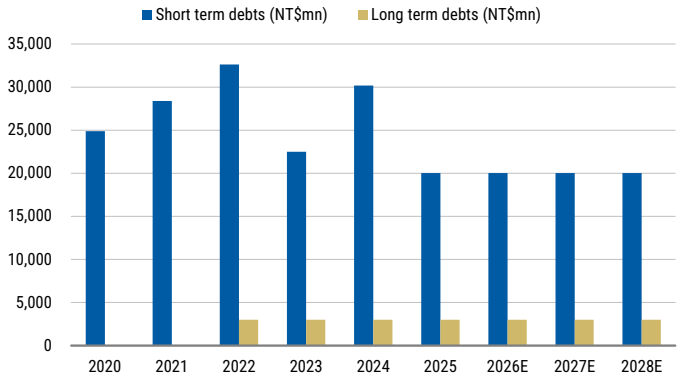
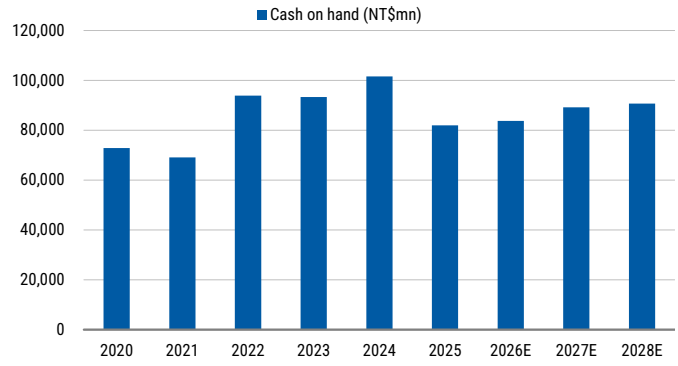


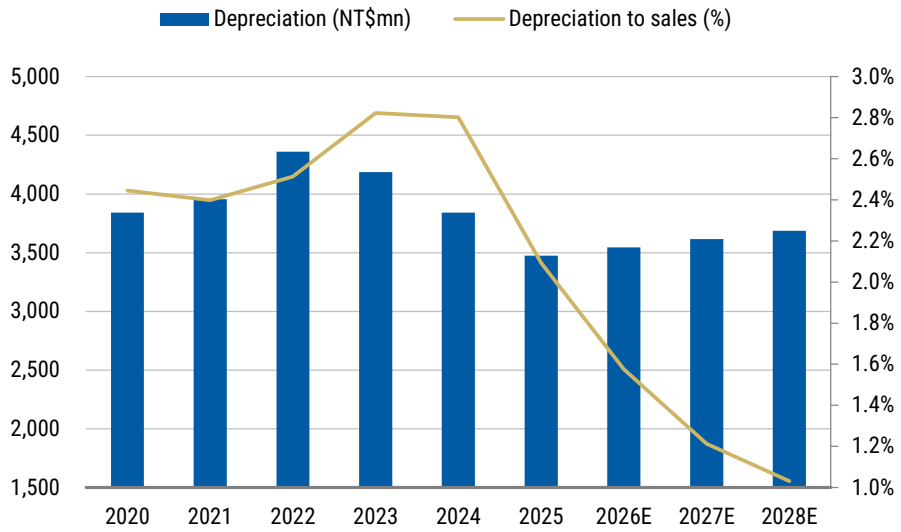
Exhibit 135: Lite-on: Cash on hand



Lite-On's average capex-to-sales ratio was around 2.3-2.4% during 2021-24, then rose to 4.1% in 2025 as the company began aggressively expanding global capacity, despite a larger revenue base during the year. Looking ahead, we project the capex-to-sales ratio to rise to 5.9% in 2026 amid continued production-site development, before declining to 3.0-3.3% in 2027-28 as investments normalize.

The depreciation-to-sales ratio remained at 2.4-2.5% during 2020-22 before rising to 2.8% in 2023, in line with higher capex during the year. The ratio then declined to 2.1% in 2025 on a larger revenue base. Looking ahead, we expect the ratio to decline further as the revenue base grows faster than capex and depreciation, falling to 1.6% in 2026 and 1.0% in 2028.

Exhibit 136: Lite-on: Depreciation and depreciation to sales



Lite-On's cash conversion cycle (CCC) fluctuated within a relatively narrow band of roughly 35-44 days between 2020 and 2025, reflecting generally stable but modestly shifting working-capital efficiency. CCC was at a relatively low level of 35 days in 2020, supported by lower inventory days and higher AP days, before rising steadily to 44 days in 2022 as both inventory days (78) and AR days (85) increased, while AP days were slightly lower (119). From 2023 onward, the cycle improved meaningfully, falling to 36 days as AP days extended sharply from 119 in 2022 to 135. This improvement partially reversed in 2024, with CCC ticking up to 39 days as inventory and AR days both increased to 89 and 92 days, respectively, even as AP days extended further to 143. CCC then eased again to 37 days in 2025 as AP days moderated slightly to 138, alongside stable inventory and receivables levels.

Looking ahead, we expect the overall cash conversion cycle to edge up slightly from 37 days in 2025 to 40 days in 2028, as increasing AR days and lower AP days are partially offset by improving inventory days.

Exhibit 137: Lite-on: AR days

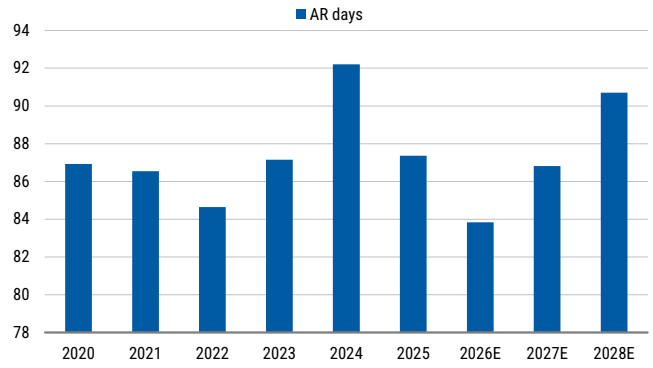


Exhibit 138: Lite-on: Inventory days

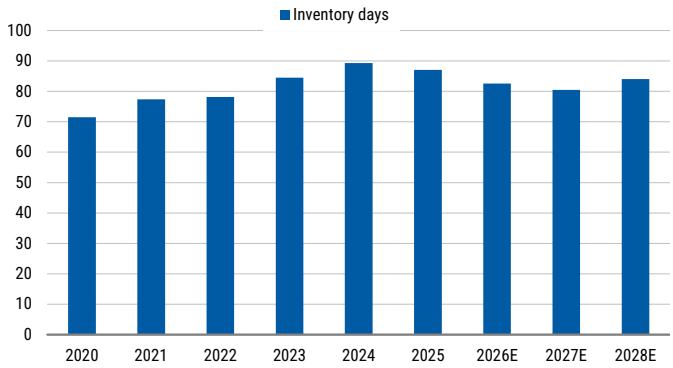


Exhibit 139: Lite-on: AP days

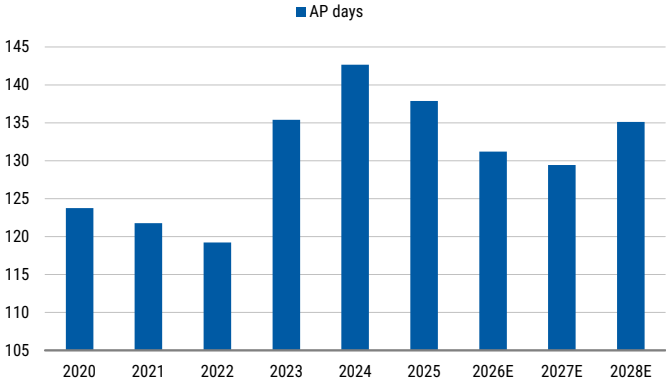
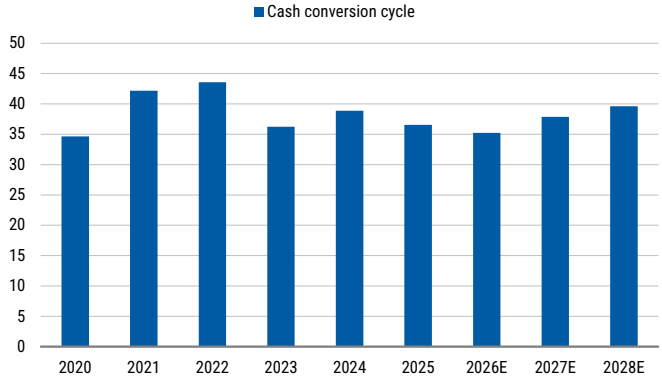


Exhibit 140: Lite-on: Cash conversion cycle



Company Background – Lite-On

Company snapshot

Founded in 1975, Lite-On Technology is a Taiwan-based manufacturer of optoelectronic components, power supplies, and information technology products. The company produced Taiwan's first LEDs and became the first technology company listed on the Taiwan Stock Exchange in 1983. It mainly operates through four business segments: Information Products (power supplies, servers, networking equipment, keyboards, and chassis), Optoelectronics (LED packaging, LED lighting, automotive electronic components, and camera modules), Storage (storage components), and other supporting businesses. Additional product lines include energy management solutions, such as data center power supply systems and cloud server chassis; networking and surveillance products; and automotive electronics, including ADAS and EV charging solutions.

Lite-On is headquartered in Taipei, Taiwan. The company operates manufacturing sites, R&D centers, and sales offices across Asia (Taiwan, China, Japan, and Singapore), the US, and Europe (the UK, the Netherlands, and Germany).

Exhibit 141: Lite-on: Key milestones

Year	Key Milestones
1975	Established LITEON's first LED factory in Zhonghe, Taiwan
1983	Became the first electronics company listed on the Taiwan Stock Exchange (Ticker: 2301)
1985	Began manufacturing power supply units.
1989	Opened first overseas plant in Bangkok, Thailand, for optoelectronic products.
1990	• Set up a switching power supply plant in Penang, Malaysia. • Silitek acquired Maxi Switch, the No. 3 U.S. keyboard maker.
1991	Acquired NPE's switching power supply plant in the UK.
1994	Established a subsidiary in California, USA.
1999	Spun off the Optoelectronics business unit as LITEON IT Corp.
2002	LITEON Electronics, Silitek, LITEON Technology, and GVC merged into LITEON Technology Corporation (Ticker: 2301)
2003	Moved global headquarters to Neihu Science Park.
2006	• Acquired BenQ's Optical Disk Drive business, becoming a top-3 global ODD maker. • Partnered with Philips to enter the automotive disc drive market. • Acquired 100% of LITEON Automotive Corp., expanding into automotive electronics.
2007	• Acquired 100% of Li Shin International via stock swap, strengthening power supply resources for LCD TVs and other products. • Completed tender offer for Perlos, the world's largest handset casing maker.
2009	Opened East China Regional HQ in Wujin, Changzhou.
2010	• Invested in Epyon Power, a supplier of fast-charging technology. • Completed the world's first Optoelectronic Semiconductor Carbon Footprint and Type III Environmental Product Declarations.
2011	Completed the world's first Product Category Rules for Circuit Boards.
2013	Acquired Power Innovations (US), strengthening cloud and power management capabilities.
2014	Selected for the Carbon Disclosure Project's Climate Disclosure Leadership Index.
2016	Opened a new operations center in Kaohsiung's Nanzih Export Processing Zone.
2018	Divested the mobile camera module and mechanical business to focus resources.
2019	Acquired 100% of LOJ via tender offer.
2020	• Launched succession plan: Tom Soong named as Chairman, Anson Chiu named as President. • Completed sale of the SSD business to KIOXIA Holdings.
2021	• Launched an Employee Stock Ownership Trust (ESOT). • Joined the O-RAN Alliance and launched 5G small cells.
2022	• Began reconstructing the Zhonghe office into the "LITEON Zhonghe Digitech Center." • Inaugurated Phase II of the LITEON Vietnam plant in Haiphong.
2023	Opened its first corporate store, 'LITEON Collection,' in Neihu Science Park.
2024	• Took a 19.99% strategic stake in COSEL. • Broke ground on a new low-carbon building in Kaohsiung.
2025	• Returned to COMPUTEX with next-gen AI solutions, winning the Sustainable Design Award (Bronze). • Reopened its founding Zhonghe headquarters after a three-year reconstruction.

Exhibit 142: Lite-on: Global operation locations



Source: Company data

Exhibit 143: Lite-on: Global footprint

Headquarter	
Taipei City, Taiwan	
Sales Offices	
Milpitas, California, U.S.A.	
Plano, Texas, U.S.A.	
Tokyo, Japan	
Shanghai, China	
Singapore	
Hoofddorp, NL	
Cramlington, Northumberland, UK	
Wetzlar, DE	
Subsidiary	**Entity name**
American Fork, Utah, U.S.A.	Power Innovations International

Product offerings

Lite-On's key offerings can be categorized into six key areas: Optoelectronic Semiconductors, Energy Management, Consumer Electronics, Networking and Surveillance, Automotive Electronics, and Roadway Intellimation System.

1. Optoelectronic semiconductors

Its product portfolio spans a broad technological range, encompassing illumination and sensing applications, low- to high-power devices, standard and custom-made solutions, and components ranging from LEDs to lasers and from visible light to infrared and ultraviolet spectra.

Exhibit 144: Lite-on: Offerings for optoelectronic semiconductors

Optoelectronic Semiconductors						
Visible, UV & Automotive LEDs				Couplers	Optical Sensors	IR Emitters & Detectors
Illumination / UV-C / Automotive				Isolation & Signal Transfer	Ambient / Proximity / Gesture	Infrared Sensing
						
Automotive LED	SMD LED	Flash LED	UV LED	Photocouplers	Optical Sensors	IR Emitters & Detectors

2. Energy management

The product range serves the communications, industrial automation, smart home, medical, office, and automotive electronics industries. It offers a comprehensive suite of platinum- to titanium-efficiency-grade power supplies (CRPS, GPU, and SLIM series) suitable for both standard and AI servers, supporting AC or DC input with high power density, high reliability, and modular scalability. Power-shelf and battery-backup-unit (BBU) solutions, along with rack-cooling and liquid-cooling technologies for data centers, cloud infrastructure, and AI/machine-learning platforms, are also included in this segment.

Exhibit 145: Lite-on: Offerings for energy management

Energy Management			
Data Center Power Supply System	**Cloud Server Chassis**	**Power Supply Module**	**Magnetic Components**
CRPS / GPU / SLIM, Platinum-Titanium	Open Rack Standard V3	PoE / Industrial / LED Driver	Transformers & Inductors
			Includes components apply for cloud infrastructure, industrial power supply, white appliances, electric vehicle and consumer electric, etc
Server Power Networking Power Open Compute Project	Server Rack Server Chassis	POE Injector Industrial Power Adapter LED Driver	Magnetic Components

3. Consumer electronics

This segment is committed to advancing home entertainment and office solutions through feature-rich product design intended to enhance the overall user experience. Offerings include office/gaming solutions, laptop keyboard modules and touchpads, laptop power adapters, USB Type-C/PD chargers, etc.

Exhibit 146: Lite-on: Offerings for consumer electronics

Consumer Electronics		
3C Accessories	**Power Supply Module**	
Office / Gaming / Laptop peripherals	PC & Networking Power	
		
Office Solutions Gaming Solutions Laptop Keyboard Module Touchpad Laptop Power Adapters USB Type C/PD Charger	PC Power Supply	Networking Adapter

4. Networking and surveillance

The product portfolio covers networking solutions, video surveillance, AIoT solutions, and server-related offerings. The solutions are designed to be user-friendly and adaptable to emerging trends in digital intelligence.

Exhibit 147: Lite-on: Offerings for networking and surveillance

Networking and Surveillance		
Video Surveillance	**Access Control & Alarms**	**Networking Solutions**
Home / Building / Module	Smart Hub / Access / Alarm	5G / WLAN / Switch / Module
		
Home Security Camera Building Surveillance Camera Camera Module	Smart Hub Access Control Alarm Kits	5G Total Solution Wireless LAN Wired Switch Wireless Module

5. Automotive electronics

Lite-On is able to leverage its core competencies in optoelectronics and power management, along with more than four decades of expertise in automotive electronics. This segment is dedicated to advancing vehicle-charging technology, ADAS (Advanced Driver Assistance Systems) solutions, and automotive lighting modules through high-end engineering. AC charging solutions are offered across multiple current ratings (32/40A, 48A, 80A, and an EU-specific variant), while DC charging solutions include a 30kW fast-charging offering.

Exhibit 148: Lite-on: Offerings for automotive electronics

Automotive Electronics		
ADAS Solutions	**AC Charging Solutions**	**DC Charging Solutions**
Vision & Lighting	32/40A - 48A - 80A - EU	Fast Charging
		
Automotive Vision Solutions Vehicle Lighting Solutions	IC 32/40A IC 48A IC 80A IC EU	DC 30kW

6. Roadway Intellimation System

This segment is mainly operated under the LEOTEK brand, a subsidiary of the LITEON Technology Group with dual headquarters in Silicon Valley and Taipei. It has served road infrastructure across more than 30 countries—including North America, Europe, the Middle East, and the Asia-Pacific region—since 1992. The segment comprises intelligent roadway lighting, traffic signal systems, and AI-driven predictive maintenance solutions, all developed to support smarter and more efficient road infrastructure.

Exhibit 149: Lite-on: Offerings for Roadway Intellimation System

Roadway Intellimation System		
Roadway Lighting Intellimation	**Intelligent Traffic Signal System**	**AI Predictive Maintenance Solution**
Smart LED Streetlights	AI-driven Signal Control	Condition Monitoring
		Includes AI warning and maintenance management, traffic flow data analysis and computation, traffic-adaptive signal control, and priority passage for emergency vehicles, etc
Roadway Lighting Intellimation	Intelligent Traffic Signal System	AI Predictive Maintenance Solution

Monthly sales vs. share price

Lite-On's monthly sales have hovered around the NT$10-20bn level since 2013. Notably, we see a high correlation of ~76% between the company's monthly sales and its share price movements since 2023.

The company reported August monthly sales of NT$19,592mn (+3% MoM/+25% YoY).

Exhibit 150: Lite-on: Monthly sales vs. sales YoY growth

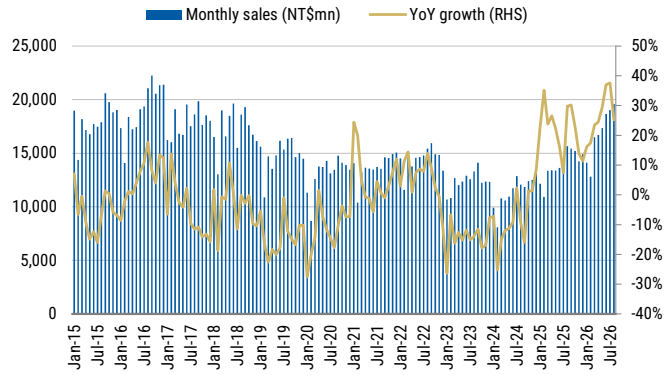
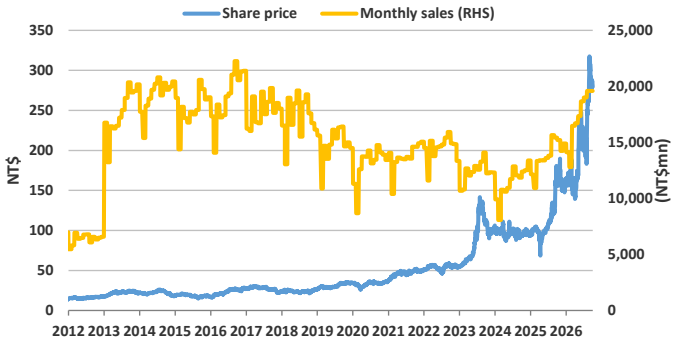


Exhibit 151: Lite-on: Monthly sales vs. share price



Key management & shareholding structure

Lite-On's board currently consists of five directors and four independent non-executive directors.

Its founder, Raymond Soong, and current Chairman, Tom Soong, together owned a stake of ~5% as of June 2026. Excluding the founder and Chairman, the management team owned a stake of ~0.1% as of June 2026.

The FINI holding percentage as of July 7, 2026, was 41%, versus a previous peak of 62% in April 2017 and a historical average (since 2010) of 46%.

Exhibit 152: Lite-on: Shareholding structure

Shareholders	Ownership
Cathay MSCI Taiwan ESG Sustainability High Dividend Yield ETF	5.6%
Yuanta/P-shares Taiwan Dividend Plus ETF	5.5%
Ta-Rong Investment Co., Ltd	3.7%
Raymond Soong*	3.4%
Fuh Hwa Taiwan Technology Dividend Highlight ETF	3.3%
Labor Pension fund	2.6%
Ming-Hsing Investment Co., Ltd	2.6%
Ta-Sung Investment Co., Ltd*	2.1%
Chunghwa Post Co., Ltd	2.0%
Lite-On Technology Corporation (Treasury shares)*	1.7%
Total	**32.6%**

Exhibit 153: Lite-on: FINI holdings vs. share price since 2010

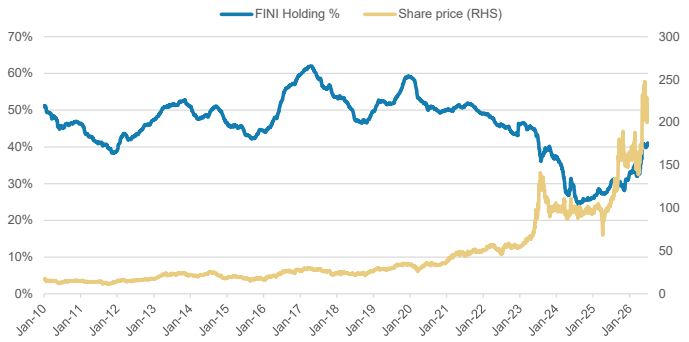


Exhibit 154: Lite-on: FINI holdings vs. share price since 2023



Exhibit 155:

Lite-on: Management profiles

Name	Title	Education	Work Experience
Tom Soong	Chairman and Chief Sustainability Officer, LITEON Technology	Electrical Engineering, University of Southern California, USA; International Business Program Attendance, NTU-FUDAN EMBA	- 20+ years at LITEON - Special Assistant to Chairman & Vice Chairman, LITEON - CEO of LITEON Smart Life and Applications Business Group - General Manager, LITEON Shanghai Operational Center
Anson Chiu	President and Director, LITEON Technology	Bachelor of Industrial Management, Lunghwa University of Science and Technology	- Joined LITEON Technology in 1987 - 35+ years in global supply chain and business development' - CEO, LITEON Power Conversion Business Group - General Manager, LITEON Power Conversion Business Group - Director, NB Business Unit, Power Conversion Business Group
Raymond Soong	Founder and Director, LITEON Technology and LITEON Cultural Foundation	Fellow, Industrial Technology Research Institute; Honorary Doctor of Management, National Chiao Tung University; Honorary Doctor of Management, Taipei University of Technology	- Chairman, LITEON Group - Chairman & Founder, LITEON Cultural Foundation - Chief Engineer, Texas Instruments Taiwan Ltd. - Board of Councilors member, Doctorate College of Technology, USC
Harvey Chang	Director Representative, Lite-On Technology Corp.	MBA, The Wharton School, University of Pennsylvania	- Independent Director, LITEON Technology - Chairman, IC Broadcasting - Chairman, TVBS and Catchplay Media Holdings Ltd. - President & CEO, Taiwan Mobile - Senior VP & CFO, TSMC
Sarina Lin	Director Representative, Lite-On; Advisor to Chairman's Office, Lite-On	Doctor of Laws, National Taiwan University	- Chief Legal Officer, Lite-On - Special Assistant to Chairman and Attorney, MediaTek Inc. - Director, MediaTek Foundation - SVP, CHRO & Chief Legal Officer, MStar Semiconductor - Director & CEO, MStar Semiconductor Charity Foundation - Director and Supervisor, MStar Semiconductor
Albert Hsueh	Independent Director, LITEON Technology;	MBA, Bloomsburg University, Pennsylvania, U.S.A.	- Chairman, PricewaterhouseCoopers Taiwan - Director, Corporate Governance Association in Taiwan - Professor, National Taiwan University and Tsinghua University, College of Technology Management
Mike Yang	Independent Director, Lite-On Technology Corp.	Master of Science, Arizona State University	- Associate Vice President, ADI Corp. Procurement Dept. - Director, LIONS Taiwan Technology Inc. - Director, Quanta Culture & Education Foundation - Cloud computing and global technology leadership
MK Lu	Independent Director, Lite-On Technology Corp.	Honorary Doctor, Tatung University; Honorary Doctor of Engineering, National Chiao Tung University; Department of Electrical Engineering, Tatung University	- Chairman, Tatung Company - Chairman, Sino-American Silicon Products Inc. - Chairman, Actron Technology Corporation - President, Lite-On Semiconductor Corp. - President, Vishay Lite-On Power Semiconductor Corp. - Vice President, Silitek Corp.
Jesse Ding	Independent Director, Lite-On Technology Corp.	PhD program, The Wharton School, University of Pennsylvania; MBA, University of Detroit, USA	- Chairman, EnTie Commercial Bank - General Manager, EnTie Commercial Bank - General Manager, TAIPEI FUBON COMMERCIAL BANK CO., LTD - Financial accounting and leadership decision-making

Valuation Methodology and Risks

Panasonic Holdings (6752.T)

Our target price is based on P/E 19.5x our FY3/28 EPS forecast.

We apply: 36x to EPS generated by AI infrastructure-related businesses, 5x to EPS attributable to IRA subsidies in the automotive battery business, and 13.5x to EPS from all other businesses.

Risks to Upside

- Acceleration of restructuring and exit/scale-down of unprofitable businesses
- Stronger than expected growth of datacenter BBU, automotive battery business

Risks to Downside

- Deteriorating visibility for energy business
- Expectations for business portfolio revisions sink

Risk Reward Reference links

1. View explanation of Options Probabilities methodology - Options_Probabilities_Exhibit_Link.pdf

2. View descriptions of Risk Rewards Themes - RR_Themes_Exhibit_Link.pdf

3. View explanation of regional hierarchies - GEG_Exhibit_Link.pdf

4. View explanation of Theme/Exposure methodology - ESG_Sustainable_Solutions_External_Link.pdf

5. View explanation of HERS methodology - ESG_HERS_External_Link.pdf