



Powering the AI Age - Utilities and the Infrastructure Backbone of Datacenters

Equity Insights
September 2025


LIGHTHOUSE
CANTON

Asia's Rising Tier in the AI & Datacenter Value Chain | Equity Insights

For two years the AI conversation has centered on U.S. hyperscalers and with reason. But the world's ability to ship compute, scale memory, assemble racks, and keep “powered shells” online increasingly runs through Asia. What used to be a manufacturing base is now a set of strategic chokepoints: domestic frontier models and cloud, the foundry/memory core, server assembly and thermals, and most pressingly the electrification stack that turns land into resilient capacity. Asia isn't a follower here; in several layers, it's the center of gravity.

In sum, Asia's role in the AI build-out is no longer peripheral—it is the flywheel. Domestic model and cloud platforms (Alibaba, Tencent, Baidu) are creating on-shore demand that anchors multi-year compute needs; the foundry and memory core (TSMC, Samsung, SK hynix, plus Japan's toolmakers) remains the global choke point; Taiwan's rack-scale and thermal ecosystem (Foxconn, Asia Vital Components) converts orders into deployable capacity; and the electrification stack (Kstar, Hitachi, Mitsubishi Heavy, CATL) turns land into resilient, powered shells.

This four-pillar coalition is mutually reinforcing: local AI platforms pull through chips; chips pull through servers and cooling; and everything ultimately clears through transformers, UPS, and storage. For investors, that means looking beyond headline GPUs to the bottlenecks with pricing power - advanced packaging, High Bandwidth Memory (HBM), liquid cooling, transformers, and large-frame UPS, while tracking catalysts like wafer starts, HBM qualifications, rack throughput, and transformer backlog. As AI scales from proofs to platforms, Asia is where scale gets manufactured, integrated, and powered.



Preferred Names embedded within the themes we are bullish on

1) Frontier Models and Regional Cloud Platforms

- Alibaba (BABA US / 9988 HK): From Qwen's open-source momentum to Model Studio usage, Alibaba is positioning as the default on-shore AI platform for Chinese enterprises. Reports in 2H25 also indicate selective use of in-house accelerators, useful hedging against export risk.
- Tencent (700 HK): The strategy is applied AI: agentic, scenario-based tooling inside a broad cloud catalog (TI Platform). That's good for stickiness and utilization.
- Baidu (BIDU US / 9888 HK): ERNIE continues to iterate, and Kunlun-based training indicates the company can progress even under silicon constraints, strategically valuable for long-run independence.

2) Semiconductors & Memory

- TSMC (TSM US / 2330 TT): Still the keystone. November commentary from Nvidia and TSMC reinforced that AI demand is pulling directly on wafer starts; guidance lifts tracked that reality.
- Samsung (005930 KS): Dual-track exposure: foundry (incl. Tesla AI5/AI6 dual-sourcing with TSMC) and HBM/packaging where demand remains fierce.
- SK hynix (000660 KS): The HBM pace-setter; HBM4 readiness keeps it on the critical path for next-gen accelerators.
- MediaTek (2454 TT): Expanding from edge to data-center collaborations with Nvidia and custom ASICs — a reminder that more XPU's are being architected in Asia.
- SMIC (981 HK / 688981 CH): A strategically important, if costlier, on-shore route for China's accelerators; reports point to 5nm (DUV-based) development by 2025.
- Tokyo Electron (8035 JT): Makes the tools chipmakers need to build the latest chips and high-speed memory; the AI build-out means more production lines and steps, lifting tool orders and ongoing service work.
- Advantest (6857 JT): Makes the testers that verify chips and stacked memory (HBM) before they ship; the surge in AI GPUs and thicker memory stacks requires more and longer testing, driving steady demand.

3) Server Assembly, Thermals & the Rack-Scale Supply Chain

- Foxconn/ Hon Hai (2317 TT): Revenue mix is now AI-heavy; Taiwan's server complex still commands the vast majority of AI server builds.
- Asia Vital Components (3017 TT): Leading thermal provider with growing liquid-cooling exposure into Tier-1 OEM/datacenter builds, benefiting directly from rising rack power density.

4) Electrification

- Kstar (002518 CH): China's #1 UPS brand; shipping modular UPS/MDC with ESS options, exactly where the bottleneck lies.
- Hitachi (6501 JT) / Mitsubishi Heavy (7011 JT): Scaling transformer and HV equipment capacity and standing up dedicated data-center business initiatives. Backlog and industry commentary point to multi-year tightness.
- CATL (300750 CH): ESS volumes and new cell formats support the "electrify the shell" thesis; storage is becoming a standard kit for new campuses.

Chinese AI offtake, following the US Blueprint

China's AI offtake is replaying the U.S. sequence almost beat-for-beat: frontier-model labs kickstarted the demand shock, and everything else reorganized around their compute appetite. In the U.S., OpenAI's GPT pre-training runs forced hyperscalers to aggregate massive clusters, which in turn catalyzed purpose-built silicon (NVIDIA's A100/H100/B100 family, AMD MI series, custom XPU's) and a memory/packaging arms race. China is following the same causal chain. Model platforms like Alibaba's Qwen, Baidu's ERNIE, and Tencent's Hunyuan—plus upstarts like DeepSeek—are scaling pre-training and reinforcement phases, pulling through enormous on-shore compute.

That demand is driving both a rush for accelerated hardware and a local silicon response: domestic accelerators (e.g., Baidu's Kunlun, Huawei's Ascend) manufactured via on-shore routes such as SMIC (semiconductor Manufacturing International Corp.) where feasible, alongside expanded deployments of AI-optimized servers assembled by Taiwan's original device manufacturers (ODMs) and fed by high bandwidth memory (HBM) from SK hynix and Samsung.

The plot points differ: export controls shape China's hardware mix more than America's, but the loop is identical: frontier models → pre-training compute surge → accelerator and memory specialization → datacenter scale-out, with the same second-order pull on packaging, thermals, and power, all of which are industries where multiple Asian companies are at the forefront and likely to benefit tremendously from the advent of AI, not only in China, but across the world, setting the stage for intense competition in neutral territories—open to using whichever technology offers the best return.

1) Frontier Models and Regional Cloud Platforms


Across China, the internet platforms are evolving from “Alusers” into “AI platforms.” Alibaba's (BABA US / 9988 HK) cloud unit has pushed squarely into full-stack model ops—Qwen's open-source line now spans flagship reasoning models through code-specialists, all delivered through Model Studio to tens of thousands of enterprises building onshore AI. In late-2025, Alibaba highlighted hundreds of Qwen variants and heavy enterprise usage—evidence the stack is not a demo lab but a commercial platform.

Baidu (BIDU US / 9888 HK) has taken the complementary path: pairing ERNIE's rapid iteration with a push into domestic accelerators (KunlunP800), reducing exposure to export-controlled silicon. Reports through 2H25 indicated Baidu training ERNIE variants on its own chips, and even academic work noted large-scale training on Kunlun clusters—an important proof point for self-reliance.

Tencent (700 HK) has leaned into “scenario-based” enterprise AI — agentic workflows embedded in cloud services (TI Platform) and vertical SaaS, which matters because sticky usage, not model benchmarks, drives utilization and monetization of domestic cloud.

Finally, DeepSeek has changed industry pricing psychology. Its R1/V-series models catalyzed a domestic price war and set a new reference for low-cost inference (including off-peak pricing), forcing incumbents to rethink unit economics. That matters for Asia’s clouds: cheaper inference expands the TAM of AI features across SMB and mid-market customers that prefer to keep data and workloads local.

Why it matters: The U.S. hyperscaler playbook (models → apps → utilization → recurring cloud) is now playing out onshore in Asia. As domestic models mature and platform usage grows, these clouds generate their own “backlog of compute,” anchoring multi-year demand for chips, racks, and power within the region.



2) Semiconductors & Memory: Where Asia Is the Bottleneck

Foundry & advanced packaging

TSMC (TSM US / 2330TT) remains the non-fungible fulcrum for bleeding-edge AI accelerators. Into November 2025, Nvidia's CEO publicly asked TSMC for "more wafers" as Blackwell demand surged; TSMC simultaneously lifted its growth outlook on persistent AI orders. The read-through is simple: if the world needs more transformers and attention heads, it needs more TSMC.

Samsung's (005930 KS) dual exposure—foundry plus advanced memory/packaging—adds a second Asian anchor. A widely reported \$16.5bn Tesla program underscored Samsung's traction in AI-oriented foundry and packaging, with recent comments confirming dual-sourcing between Samsung and TSMC for Tesla's next generations (AI5/AI6). It's a mitigation of single-source risk and a signal that leading AI silicon will remain at Asian fabs.

High bandwidth memory & the memory throttle

In AI systems the need for memory is immense and therefore bandwidth has become the constraint, without which the many billions being spent on accelerators (GPUs and XPU) are rendered useless and SK hynix (000660 KS) sits at the sharp end. The company completed HBM4 development and prepared mass production, with industry coverage highlighting ongoing negotiations for next-gen Nvidia platforms. The company has also guided to sustained AI-memory growth into the decade. Memory isn't a side-story anymore; it's the performance and TCO lever.

Domestic chip design & sovereignty

Chip designers such as Baidu (BIDU US / 9888 HK), Xiaomi (1810 HK), Cambricon (688256 CH), Huawei and more recently even Chinese cloud service providers such as Alibaba (BABA US / 9988 HK) are working on developing their own indigenous chips for the development of AI accelerators, which are concurrently being manufactured by domestic Chinese fabs such as SMIC (981 HK / 688981 CH). Additionally, Taiwanese players such as MediaTek (2454 TT) seem to have captured a share of the XPU / application specific integrated circuits market, which are used by the hyperscalers to design their in-house chips.

Wafer Fab Equipment manufacturers

There are multiple wafer fab equipment companies across Asia, such as Tokyo Electron (8035 JT), Advantest (6857 JT) etc. supplying critical packaging technology / process control equipment which is imperative for advancing chip technology going forward - supplying to the leading players such as TSMC.

Investor angle: Foundry and HBM are true chokepoints—capacity, yield and customer qualifications create multi-year visibility; custom ASIC/XPU work widens the funnel beyond GPUs, tying more designs back to Asian fabs and OSATs.

Value Flow: WFE enables → Foundry manufactures → Memory adds bandwidth → Packaging integrates → Designers deploy

• Foundry Chokepoint

TSMC + Samsung control 90%+ of advanced node capacity (<5nm). Nvidia's November 2025 request for "more wafers" and Samsung's \$16.5bn Tesla program confirm: **leading AI silicon remains at Asian fabs.**

• Memory Throttle

SK hynix leads HBM4 mass production, with ongoing negotiations for next-gen Nvidia platforms. **Bandwidth is now the constraint** — memory isn't a side-story, it's the performance and TCO lever.

• Domestic Sovereignty Push

Chinese designers (Baidu, Alibaba, Huawei, Cambricon) + SMIC manufacturing creates **onshore AI accelerator pipeline**. MediaTek captures XPU/ASIC share for hyperscaler custom chips.

• Investor Angle

Capacity, yield, and customer qualifications create multi-year visibility. Custom ASIC/XPU work widens the funnel beyond GPUs, tying more designs back to **Asian fabs and OSATs**.

3) Server Assembly, Thermals & the Rack-Scale Supply Chain

When orders become infrastructure, Taiwan's ODM/EMS complex does the heavy lifting. Foxconn (2317 TT) has pivoted from "iPhone assembler" to "AI rack assembler," with management and industry data consistently pointing to Taiwan supplying ~90% of AI servers and Foxconn itself taking a very large share. Through 2025, AI servers and networking overtook consumer devices in Foxconn's revenue mix, with quarterlies flagging >40% mix from cloud/AI and growth well ahead of plan.

Around Foxconn is a dense ecosystem of specialists. Asia Vital Components (3017 TT) is a leading thermal provider shipping fans, heat-sinks, liquid-cooling modules and cold-plates into Tier-1 Original equipment manufacturers and, by extension, the western clouds. As rack power density rises and liquid cooling normalizes, this "thermal layer" becomes a performance and uptime determinant — not a commodity.

Why it matters: This layer is less exposed to "which model wins" and more exposed to volume ramp, thermal envelopes, and integration timelines. In other words, it's a cleaner way to play AI infrastructure scale — especially as customers de-risk with multi-region builds.

• Taiwan's 90% Market Share

Foxconn has pivoted from "iPhone assembler" to "AI rack assembler." Through 2025, AI servers and networking overtook consumer devices in revenue mix, with **>40% mix from cloud/AI** and growth ahead of plan.

Taiwan supplies ~90% of global AI servers

• Thermal Layer: From Commodity to Critical

Asia Vital Components (AVC) ships fans, heat-sinks, liquid-cooling modules and cold-plates into Tier-1 OEMs. As rack power density rises and liquid cooling normalizes, the thermal layer becomes a **performance and uptime determinant**.

No longer a commodity component



Why This Layer Matters for Investors

This layer is **less exposed to "which model wins"** and more exposed to **volume ramp, thermal envelopes, and integration timelines**. In other words, it's a **cleaner way to play AI infrastructure scale**—especially as customers de-risk with multi-region builds. The thermal ecosystem around Foxconn represents predictable demand tied to compute density, not model benchmarks.

4) Electrification: Solving the Powered-Shell Bottleneck

The immediate constraint in many markets is no longer “power in the zip code” but the powered shell: UPS, MV/HV transformers, switchgear, cabling, and behind-the-meter storage to deliver resilient, high-quality power to high-density halls. Here too, Asia’s manufacturers are decisive.

Shenzhen Kstar (002518 CH), China’s leading UPS provider by share, has expanded aggressively into modular datacenter power, with product lines spanning large-frame UPS, precision cooling and containerized MDC — exactly the kit operators need when turning land into Tier III/IV capacity under tight timelines.

Hitachi Energy (6501 JT) and Mitsubishi Heavy Industries (7011JT) bring the heavy electricals: large power transformers, HV switchgear, and grid integration. Throughout 2025 they announced multi-billion capacity expansions and new data-center-focused business lines, explicitly tying investments to AI-driven grid demand and transformer bottlenecks. Industry commentary and reporting across Q3–Q4 2025 pointed to a structurally tight transformer market with backlogs stretching years; both companies are adding factories and U.S. footprint to meet hyperscale timetables.

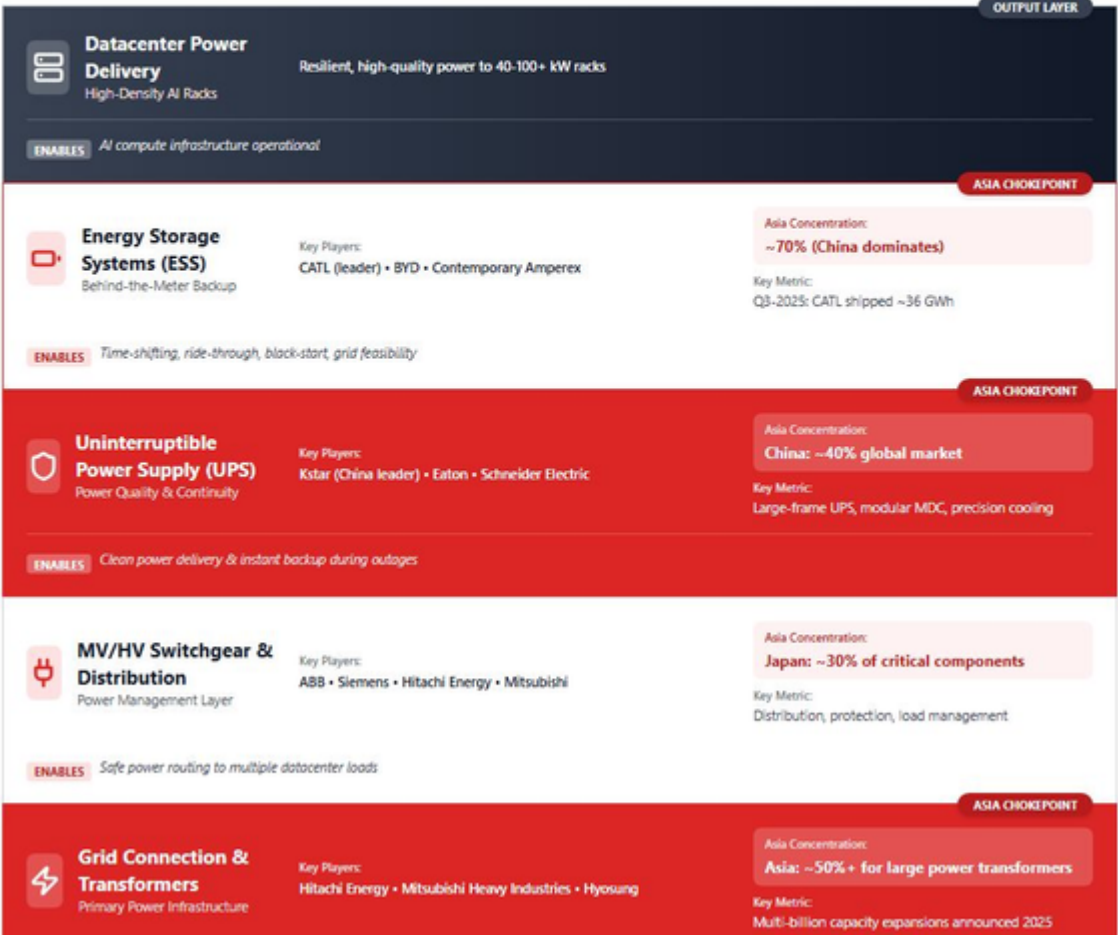
On the storage side, CATL’s (3750 HK / 300750 CH) ESS shipments have scaled rapidly, with Q3-2025 reporting ~36 GWh and fresh supply agreements into year-end — important because storage now underwrites datacenter operability (time-shifting, ride-through, black-start) and grid-connection feasibility. Kstar even markets integrated ESS built on CATL cells, illustrating the vertical tie-ups forming in Asia’s power stack.

The punchline: Chips grab headlines, but transformers, UPS and ESS clear the land. Asia owns crucial capacity in each.

Electrification: Solving the "Powered Shell" Bottleneck

Complete Power Infrastructure Stack for AI Datacenters

Beyond "power in the zip code" — Asia dominates every layer needed to deliver resilient power



LIGHTHOUSE CANTON

Singapore

16 Collyer Quay #11-02
Collyer Quay Centre
Singapore 049318

☎ +65 67130570

India

Unit No-104A,Worldmark
2 Asset Delhi Aerocity,
New Delhi 110037

☎ +91 9650473961

Unit No 507/508, A Wing,
INS Tower, G Block, BKC,
Mumbai- 400051

☎ +91 9650473961

UAE

The Exchange Gate Village 11,
Unit 802 Dubai International
Financial Centre PO Box 507026
Dubai, UAE

☎ +971 45 861500

1st Floor, WeWork37,
Cunningham Cross Rd,
SRT Road Vasant Nagar,
Bengaluru-560001

☎ +91 9900096873

RK Swamy Centre, Hansa
Building, Door No:3, Thousand
Lights,
Chennai-600006

☎ +91 9650473961

UK

24 Hanover Square,
London W1S 1JD

☎ +44 164 2843 487

Suite 502, Building 450, Central
Plaza, Genome Valley,
Shameerpet,
Hyderabad 500 078

☎ +91 9650473961

Unit No FF-10, FF Floor, Pragya
Accelerator, Block 15T GIFT CITY,
Gandhinagar
Gujrat-382355

☎ +91 9650473961

✉ info@lighthouse-canton.com

✉ service@lighthouse-canton.com

in Lighthouse Canton

DISCLAIMER

The contents of this document are confidential and are meant for the intended recipient only. If you are not the intended recipient, please delete all copies of this document and notify the sender immediately.

This document, provided as a general commentary, is for informational purposes only and is not to be construed as an offer to sell or solicit an offer to buy any financial instruments in any jurisdiction. This does not constitute any form of regulated financial advice, and your independent financial advisor should be consulted prior to taking any investment decision(s).

This document is based on information from Sources which are reliable but has not been independently verified by Lighthouse Canton Pte Ltd and its affiliate companies ("LC"). LC has taken the reasonable steps to verify the contents of this document and accept no liability for any loss arising from the use of any information contained herein. Please also note that past performances are not indicative of future performance.

Information contained herein are those of the author(s) and does not represent the views held by other parties. LC is also under no obligation to update you on any changes made to this document.

This document is prepared by Lighthouse Canton Pte Ltd and its affiliate company, Lighthouse Canton Capital (DIFC) Pte Ltd, which are regulated by Monetary Authority of Singapore ("MAS") and Dubai Financial Services Authority ("DFSA") respectively. MAS and DFSA has no responsibility for reviewing, verifying and approving the contents of this document and/or other associated documents. The contents of this document may not be reproduced or referenced, either in part or in full, without prior written permission from LC.

This document is confidential and is only intended for Accredited Investors and/or Professional Clients, as defined by MAS and DFSA.