

July 27, 2026 09:00 PM GMT

Greater China Semiconductors | Asia Pacific

AI Networking: Scale-up Technology Enables China's AI Supernodes

Following our [WAIC takeaways](#), we further examine China's emerging scale-up landscape, the role of optical interconnect, and implications for the domestic interconnect supply chain.

AI Supernodes enabled by scale-up are becoming the focus of China's AI computing landscape. Alongside our China [WAIC takeaways](#), we provide an in-depth follow-up study of China's scale-up technology development. As we wrote in [China's AI Accelerators – Who's Poised to Win?](#), Chinese AI GPUs are strong in optical networking and server rack design at the system level, although they remain constrained by wafer process technology at the chip level ([Exhibit 6](#)). These chip-level constraints are shifting competition in China's AI computing market from standalone chip specs to system solutions. WAIC 2026 featured fewer chip launches and more supernode solutions. Leading domestic accelerator vendors displayed rack-scale or multi-rack supernodes with scale-up domains of 64 accelerators or more, mostly enabled by proprietary scale-up technologies.

China's scale-up landscape is rapidly broadening, led by proprietary fabrics but spanning multiple system architectures ([Exhibit 4](#)). A year ago, our US semis team published [Semiconductors: Scaling the AI Opportunity: A Scale-Up Network Primer \(29 Aug 2025\)](#); China is now catching up quickly in scale-up technology. Huawei remains the domestic scale leader, with Atlas 950 using UnifiedBus 2.0 to connect 1,024 NPUs in the demonstrated configuration and designed to scale to 8,192 NPUs. Biren's next-gen BR2xx NPO architecture targets up to 1,024 GPUs through BLink 2.0. Sugon's scaleX640 connects 640 accelerators in a single rack. Moore Threads(not covered) MTT C256 supports 128 GPUs in one rack and 256 across two racks through MT-Link 2.0, while MetaX, Alibaba and Enflame support 64- to 128-accelerator domains through MetaXLink-E, ICN, and GCU-LARE, respectively.

Stock implications: We believe **Montage** (OW, covered by Daniel Yen) is set to benefit from higher PCIe interconnect content in larger, distributed supernodes. As systems scale across racks, denser and longer PCIe links among CPUs, xPUs, switches, NICs and peripherals should support demand for PCIe retimers and PCIe switch chips. We also expect China AI GPU vendors, such as **Hygon** (OW, covered by Daisy Dai) (along with its ecosystem partner Sugon), **Cambricon** and **Iluvatar** (both OW, covered by Charlie Chan), to introduce or leverage scale-up technology to improve their AI server rack performance.

(Continued below.)

VORGAN STANLEY ASIA LIMITED+

Ethan Jia

Research Associate

Ethan.Jia@morganstanley.com

+852 2943 1187

VORGAN STANLEY TAIWAN LIMITED+

Charlie Chan

Equity Analyst

Charlie.Chan@morganstanley.com

+886 2 2712 1215

Daniel Yen, CFA

Equity Analyst

Daniel.Yen@morganstanley.com

+886 2 2712 0863

VORGAN STANLEY ASIA LIMITED+

Daisy Dai, CFA

Equity Analyst

Daisy.Dai@morganstanley.com

+852 2815 7310

Henry Zhao

Research Associate

Henry.Zhao@morganstanley.com

+852 2234 7731

VORGAN STANLEY TAIWAN LIMITED+

Tiffany Yeh

Equity Analyst

Tiffany.Yeh@morganstanley.com

+886 2 2711 3037

Lucas Wang

Research Associate

Lucas.Wang@morganstanley.com

+886 2 2712 0863



GREATER CHINA TECHNOLOGY SEMICONDUCTORS

Asia Pacific

Industry View

Attractive

Morgan Stanley does and seeks to do business with companies covered in Morgan Stanley Research. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of Morgan Stanley Research. Investors should consider Morgan Stanley 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.

China AI Interconnect - Scale-up


Global AI server scale-up technology trend

What is scale-up networking?

Scale-up networking refers to high-speed communication between GPUs/accelerators within the same server or rack. Interconnecting accelerators allow them to function as one large supercomputer, accessing the same memory and processing the same workload. This differs from scale-out networking, which connects multiple systems across a broader data center infrastructure.


Montage's PCIe 6.x/CXL3.x x16 AEC has completed interoperability testing, and targets supernode and inter-rack deployments. Chinese systems deploying more xPUs per unit of compute could further increase PCIe content. Beyond standardized PCIe connectivity, multi-rack scale-up should also accelerate demand for domestic optical interconnect.

Exhibit 1: Scale-up interface data rate per lane



Source: Company data, Morgan Stanley Research

Exhibit 2: Scale-up bandwidth per accelerator



Source: Company data, Morgan Stanley Research

Exhibit 3: US AI networking scale-up technologies

Technology	NVLink5	UALink 1.0	SUE	PCIe Gen 6	ICI	NervuonLink
Purpose	GPU-GPU; CPU-GPU	xPU-xPU	xPU-xPU	Primarily CPU-XPU; xPU-xPU	TPU interconnect	Trainium interconnect
Bandwidth	100 GB/s	100 GB/s	100 GB/s	64 GT/s	50-100 c B/s	not standardized
Latency	10-20 nm	<100 nm	-	-	varies	varies
Ecosystem	NVIDIA (proprietary)	Open standard	Open standard	Open standard	Google (proprietary)	Amazon (proprietary)
Scalability	576 GPUs	1,024 xPUs	1,024 xPUs	varies	varies	64 Trainium


Source: Company data, Morgan Stanley Research

Exhibit 4: AI supernode scale-up solutions

Scale-up tech summary	NVIDIA	AMD	Ascend	Huawei	Moore Threads	Sugon	MetaX	Altair	Enflame	Enflame	Bion	Bion
Supernode / accelerator name	GA100 NVL72 (144 die)	AMD Holm / MI 455A	N/A	Atlas 950 SuperPoD	MTT C256 / MT M10000 R100	scale640	Hygon H800 / C800	Parisc A100 / Zhenyu M800	Turrisi G64-G	NPO optical prototype	Light Pattern L78 / B4 155A	ML200 NPO supernode (announced)
Scale-up solution	NVLink 5 + NVSwitch	Link 1.0 / Ulink	Verapex V-Series	UnifiedBus / B	MTLink 2.0	scale640	MetasLink and MetasLink-E	AIlink system; M800 implements ICN + ICN Switch 3.0	G64-LINK	Undisclosed	Undisclosed	Bank 2.0
Technology foundation	Proprietary	Ethernet PHY-based, memory-remains, Ulink fabric	PCIe Gen5	Proprietary	Proprietary	Open multi-vendor system architecture	Proprietary	Supports proprietary protocol and Ulink	Proprietary	Undisclosed	SiPh OCS	Proprietary
Single node speed	Physical line rate 90A, 100GB/s per NVLink 5 link	200Gb/s per lane, 212.5GT/s raw signaling rate	Undisclosed	800.25Gb/s PAM4 under USB2.0 specification	112Gb/s per lane	Undisclosed	Undisclosed	112G/224G SerDes, supported at A100 platform level	Undisclosed	Undisclosed	Undisclosed	Undisclosed
Total bandwidth per accelerator	1.8TB/s	3.6TB/s	Undisclosed	2TB/s	800GB/s	Undisclosed	Undisclosed	800GB/s	Undisclosed	Undisclosed	Undisclosed	Undisclosed
Total accelerator units in supernode	72	72	Up to 80 per switch (overhead solution)	1,024 NPUs demonstrated / 8,192 NPUs announced	128 GPUs in one rack, 256 GPUs across two racks	640	64 GPUs in 5000 rack, 128 GPUs through 6800 MetasLink-E	128 per rack, 64 per chassis of full-bandwidth ICN domain	64	512-plus	128	1,000
Physical interface	Passive copper, 100Gb/s full-duplex edge / backplane	Electrical/copper	Electrical/copper	Hybrid orthogonal optical, electrical interconnect intra-rack, optical interconnect inter-rack	All-copper Cable Tray	Electrical/copper	Cables OCS electrical interconnect	For optical fiber, orthogonal electrical, full interconnect	Cables OCS electrical interconnect	NPO optical	SiPh optical circuit switching, supported for Lightelligence	NPO optical
Copper cable	Yes	Yes	Yes	No (published)	Yes	Undisclosed	No (published)	Yes	No (published)	No	No	No
CPO	No for scale-up (Yes for scale-out)	No	No	Undisclosed	No	Undisclosed	No	No	No	No	No	No
NPO	No	No	No	Undisclosed	No	Undisclosed	No	No	No	Yes	No	Yes

Source: Company data; Morgan Stanley Research

Exhibit 5: Scale-up vs. scale-out networking



Source: Broadcom 2025 GCP Summit Keynote; Morgan Stanley Research

China AI server scale-up development


In China, electrical interconnect remains the preferred option within racks, while optical becomes more important as scale-up domains extend across racks. Short electrical links offer lower cost and simpler integration within boards, trays and single-rack systems. Moore Threads' C256 uses an all-copper cable tray design, while the orthogonal zero-cable architectures by Enflame and MetaX also remain electrical. Optical interconnection becomes more attractive across racks, where copper reach, signal loss and power consumption become constraints. Huawei's Atlas 950 SuperPoD extends UnifiedBus across racks within the same SuperPoD using optical interconnect.

The transition toward optical connectivity may create demand for domestic switching and optical engine solutions designed for AI systems. At WAIC 2026, Lightelligence (not covered) showcased its new NPO/CPO switch solutions (Exhibit 8, Exhibit 9), which pair the switch ASIC from a leading domestic vendor, which we believe is Centec Communications (not covered), with Lightelligence's optical engine, including the SiPh PIC and selected EICs such as the TIA. The NPO/CPO solutions shorten the high-speed

electrical path between the switch ASIC and optical engine, improving bandwidth density and power efficiency, while NPO retains greater serviceability than CPO. The same architecture can serve both scale-up and scale-out networks, with scale-up becoming increasingly relevant as tightly coupled accelerator domains extend across racks.


More broadly, scale-up is emerging as the principal area of architectural differentiation, while scale-out remains the backbone of large-scale clusters. System solutions generally use a specialized fabric to connect accelerators within a supernode, and Ethernet-based fabrics such as RoCE or UBoE, or Infiniband to connect multiple supernodes across a cluster. NVIDIA combines NVLink with InfiniBand or Spectrum-X Ethernet; AMD Helios pairs UALink with Ethernet scale-out; Huawei uses UnifiedBus within the SuperPoD, and UBoE or RoCE for scale-out between SuperPoDs.

Exhibit 6: Relative strengths of AI industries in the US and China - optical networking almost on par



Source: Morgan Stanley Research

Exhibit 7: Domestic chips have lower TCO and comparable per token cost (AI LLM inference) vs. NVIDIA's processors for China




Source: Company data; Morgan Stanley Research estimates

Optical interconnect is progressing towards scale-up system deployment


Enclave and Lightelligence previously demonstrated an xPU-CPO prototype that places the optical engine beside the accelerator, enabling direct optical output from the xPU and reducing the electrical distance between the compute silicon and optics. At WAIC 2026, Enclave displayed an NPO-equipped GPU-server solution targeting scale-up domains of 512 accelerators or more, developed in partnership with Lightelligence ([Exhibit 10](#)). Biren also unveiled a next-gen BR2xx-based NPO architecture designed to support a scale-up domain of up to 1,024 GPUs ([Exhibit 11](#)). Also, LightSphere X (jointly developed by Lightelligence, Biren, ZTE, and other partners) incorporates Lightelligence's SiPh interconnect and dOCS technology, enabling the accelerator domain to extend across racks and allowing the topology to be reconfigured according to workload requirements ([Exhibit 12](#)). These projects show domestic optical technology being integrated at both the accelerator optical-I/O layer and the distributed optical-switching layer, to support larger, multi-rack scale-up domains.

Exhibit 8: 51.2T CPO switch module (displayed by Lightelligence)



Source: Company data, Morgan Stanley Research

Exhibit 9: NPO switch module (displayed by Lightelligence)




Source: Company data, Morgan Stanley Research

Exhibit 10: Enflake's NPO optical interconnect solution for scale-up




Source: Company data

Exhibit 12: LightSphere X (Lightelligence)



Source: Company data, Morgan Stanley Research


Exhibit 11: Biren's NPO optical interconnect solution for scale-up



Source: Company data


US scale-up networking ecosystem and TAM forecast

Exhibit 13: Chinese CSPs' capex will be a key demand driver for Chinese AI GPUs



Source: Company data, Morgan Stanley Research

Exhibit 14: US semi : Scale-up TAM estimate (\$mn)



Source: Company data, Morgan Stanley Research (F) estimates

This report references U.S. Executive Order 14032 and/or entities or securities that are designated thereunder. U.S. persons may be prohibited from buying certain securities of entities named in this report. Readers are solely responsible for ensuring that their investment activities are carried out in compliance with applicable laws.

This report references export controls and/or entities that may be subject to export control restrictions. Readers are solely responsible for ensuring that their investment or trade activities are carried out in compliance with applicable laws.

This report references U.S. Executive Order 14105 and/or entities that may be in scope of such order. U.S. persons may be prohibited from engaging in certain transactions or otherwise require certain other transactions be notified to the U.S. Department of Treasury. Readers are solely responsible for ensuring that their investment or trade activities are carried out in compliance with applicable laws.

Valuation Methodology and Risks

Montage Technology Co Ltd (6809.HK)

Base case, residual income model. Key assumptions:

- 8.4% CoE (1.2 beta, 3.0% risk-free rate, and 4.5% risk premium)
- 30% payout ratio
- 19.3% medium-term growth rate
- 4% terminal growth rate

These assumptions reflect cloud capex growth, DRAM interface technology migration, and Montage's strong position in China's datacenter semi localization. We then apply an exchange rate of 1.15 HKD:1 RMB, assuming no H-share discount vs the A-share.

Risks to Upside

- Faster-than-expected phase-out of US peers
- Faster-than-expected spec migration

Risks to Downside

- Weaker-than-expected cloud demand
- Slower-than-expected DRAM interface technology migration
- Delay in new product launches

Montage Technology Co Ltd (688008.SS)

Base case, residual income model. Key assumptions:

- 8.4% CoE (1.2 beta, 3.0% risk-free rate, and 4.5% risk premium)
- 30% payout ratio
- 19.3% medium-term growth rate
- 4% terminal growth rate

We believe these assumptions are justified, given the cloud capex growth, DRAM interface technology migration, and Montage's strong position in China's datacenter semi localization.

Risks to Upside

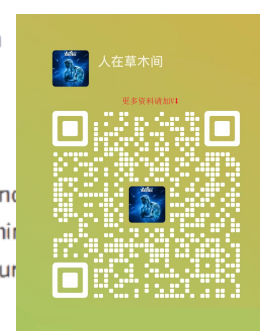
- Faster-than-expected phase-out of US peers
- Faster-than-expected spec migration

Risks to Downside

- Weaker-than-expected cloud demand
- Slower-than-expected DRAM interface technology migration
- Delay in new product launches

Hygon Information Technology Co., Ltd. (688041.SS)

We assume an 8.0% cost of equity (beta 1.05, risk-free rate 2.0% and a risk premium of 6.0%), a payout ratio of 50%, a medium-term growth rate of 25.0%, and a terminal growth rate of 4.0%, all of which are in line with other China AI chip companies under our coverage.



Risks to Upside

- Hygon further differentiates itself from local CPU and GPU competitors and gains significant market share in China's CPU and GPU markets
- China AI demand is stronger than expected
- Faster ramp-up of leading node capacity

Risks to Downside

- Intensified pricing competition among local GPU companies
- China AI demand is weaker than expected
- Slower yield improvement and capacity buildout at local leading node foundries

Cambricon Technology Corporation (688256.SS)

Key valuation assumptions underpinning our model include: an 8.4% cost of equity (derived from a beta of 1.06, risk-free rate of 2.0%, and equity risk premium of 6.0%), a long-term payout ratio of 40% (was 57%), a medium-term growth rate of 16%, and a terminal growth rate of 6.0%.

Risks to Upside

- Stronger-than-expected AI demand
- CSP order ramp-up
- Accelerating localization

Risks to Downside

- Capacity and yield constraints
- Customer concentration risk
- Slower technology iteration

Iluvatar CoreX Semiconductor Co., Ltd. (9903.HK)

Key valuation assumptions include:

- An 8.3% cost of equity (derived from a beta of 1.05, risk-free rate of 2.0%, and equity risk premium of 6.0%)
- A long-term payout ratio of 34% (was 33%)
- A medium-term growth rate of 16%
- A perpetual terminal growth rate of 6%

Risks to Upside

- Stronger-than-expected CSP orders
- Faster CUDA replacement with Iluvatar's software
- Expansion of overseas and domestic capacity

Risks to Downside

- Order ramp below expectations
- Escalation of sanctions
- Intensifying competition