AI networking – scale-across & scale-in

Core players expanding business frontiers

Scale-across network's development may lead to stronger demand for optical network solutions; scale-in demand rising in the Agentic AI era

As AI model training size exceeds trillions of parameters and inference demands grow exponentially, the expansion of a single AI cluster or data center faces limitations such as power, cooling, and land, at scale-up and scale-out level. It has given rise to distributed training that connects multiple clusters or data centers as a unified computing pool to schedule, which needs to be supported by scale-across network. A scale-across environment needs high-speed optical networking such as coherent optical modules, hollow-core optical fiber, and DCI/WAN/metro routers and switches, forming a cross-site backbone. Compared to traditional DCI, scale-across network has higher networking requirements including higher bandwidth, lower latency and packet loss rates, which have to be supported by advanced technologies such as high-speed coherent optical transceivers (800G/1.6T), OCS, and hollow-core optical fibers. We think the development of scale-across network will accelerate demand for these technologies, and we expect YOFC (6869 HK, Buy) to benefit from commercialization of hollow-core optical fibers, while Innolight (300308 CH, Buy) and Accelink (002281 CH, Buy) should benefit from increasing demand for high-speed coherent optical transceivers. In addition, NVIDIA (NVDA US, Not rated) has introduced a concept of scale-in, which refers to offloading and accelerating traffic such as user access, data handling, storage read/write, and secure isolation through DPUs and Ethernet at the north-south boundary of the AI factory. Scale-in connects AI factory and enterprise data to AI compute, and we think it will become important in the Agentic AI era..

Major technologies used in scale-across network: high-speed coherent optical modules, OCS and hollow-core optical fibers

For optical modules, IM-DD optical modules used in scale-out is less than 2 km, while coherent optical modules can be extended to 80-120 km, making them more suitable for long-distance scale-across scenarios. Moreover, coherent-lite transceiver was first proposed by Google (GOOG US, Not rated) in 2021, which uses a simplified DSP instead . We expectcoherent-lite transceivers with more powerful DSPs to be used in the scale-across network. In addition, as we mentioned in Global AI Trend Tracker – Optical circuit switch at center of AI network, OCS as an all-optical switch provides the flexible networking over long distances and high bandwidth, which could become a crucial component in scale-across network. For optical fibers, as hollow-core fibers use air, inert gases, or vacuum instead of quartz as transmission media, it can generate ultra-low-loss transmission latency compared to traditional optical fibers, which is more suitable for scale-across network with low-latency requirement.

Scale-across market TAM and competitive landscape analysis.

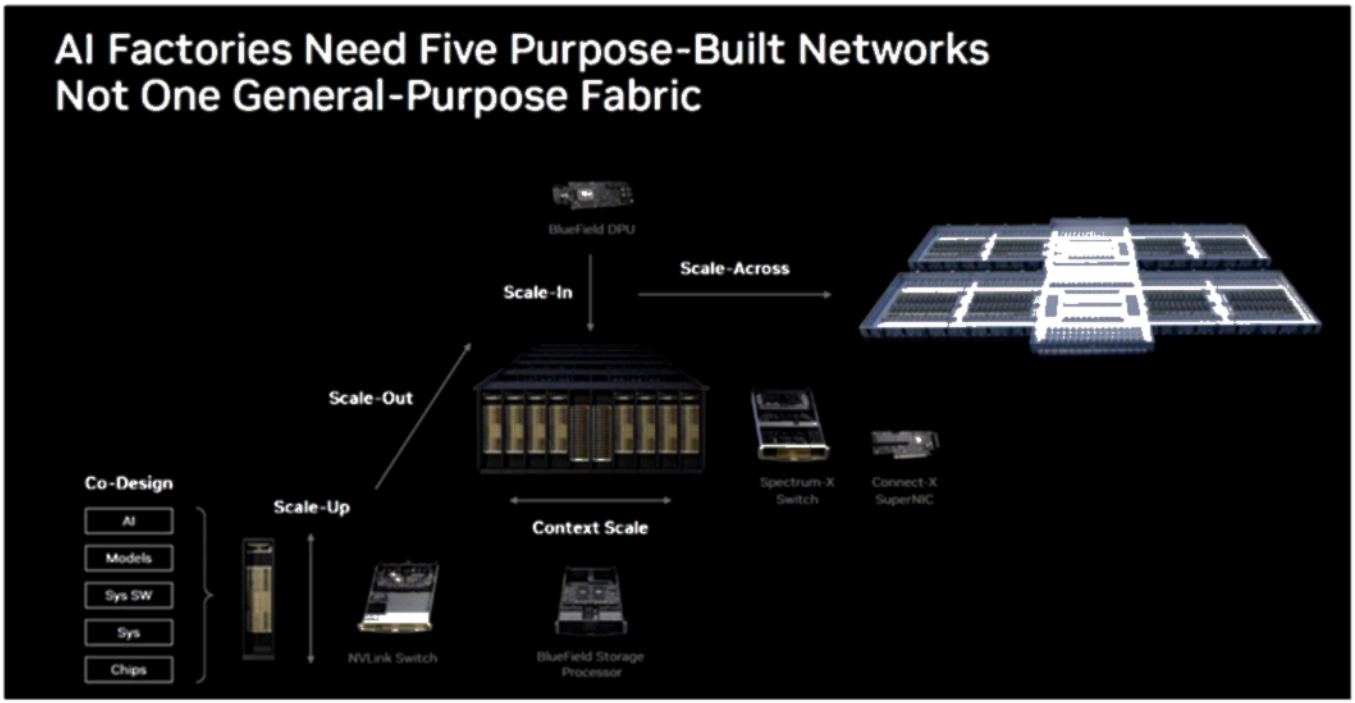
Ciena (CIEN US, Not rated) management expects scale-across market to reach USD8-10bn by 2029E at its 2QFY26 earnings call. We think major players in sectors such as coherent optical transceivers, OCS, and optical fibers could benefit from this incremental market. Marvell (MRVL US, Not rated), Acacia (a subsidiary of Cisco (CSCO US, Not rated), Ciena and Nokia (NOK US, Not rated) remain the major players in 400G / 800G ZR/ZR+ modules. The core suppliers of hollow-core optical fiber include YOFC (6869 HK, Buy) , Hengtong Optoelectronics (600487 CH, Not rated), Lightera (Unlisted, a subsidiary of Furkawa (5801 JP, Not rated)), and Relativity Networks (Unlisted).. Google, Coherent (COHR US, Not rated) and Lumentum (LITE US, Not rated) currently are the main players in the OCS market.

AI networking – scale across network explained

In the past three years, the rise of generative AI market and large language models (LLMs) has led to an unprecedented demand for large-scale, high-speed and low-latency AI networks. As scale-up/out network (within a single AI data center or cluster) cannot fulfill the ever increasing communication demand for large AI clusters, due to limitations such as power, cooling, and land, a scale-across network was proposed by NVIDIA in 2025. Scale-up network connects multiple GPUs into a single compute node, and it focuses on within-rack interconnection. Scale -out connects more cabinets within a single data center and it expands overall computing power by increasing the number of nodes, while scale-across interconnects multiple data centers scattered across different regions into a larger-scale AI factory, enabling them to collaboratively run a single AI task.

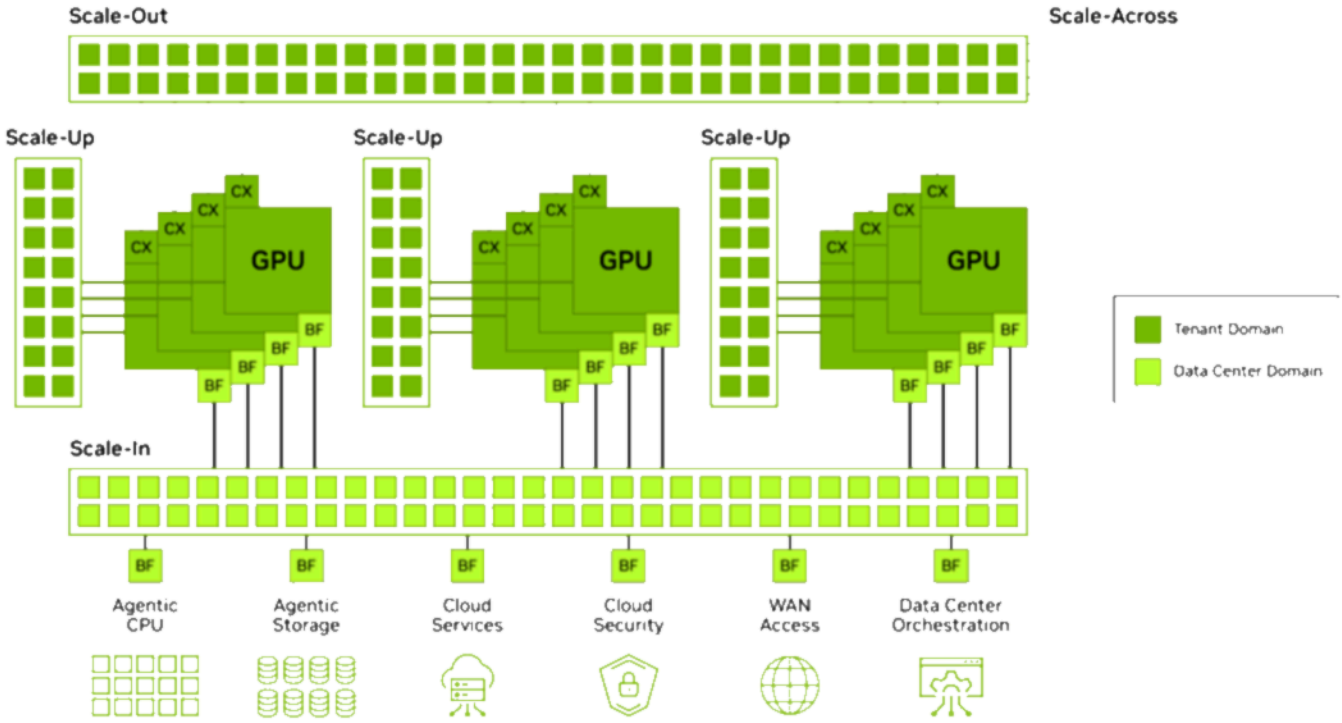
Following scale-across, NVIDIA has introduced a new concept-scale-in as one of the five pillars for the AI networking infrastructure. Scale-In refers to offloading and accelerating traffic such as user access, data handling, storage read/write, and secure isolation through DPUs and Ethernet at the north-south boundary of a AI factory. The five pillars of AI networking are scale-in, context memory, scale-up, scale-out, and scale across. We think scale-in network may become increasingly important, due to: 1) the increasing scale of AI agents has increased the complexity for data transfer, storage and multi-rounds of communications; 2) expensive CPUs handle a lot of storage, security, and other miscellaneous tasks, which require extensive communication workloads; and 3) highly sensitive data in finance, healthcare, and other sectors needs encryption and isolation before entering the GPU / CPU clusters, so moving the security boundary forward to the DPU is more controllable than relying on CPUs within the host. Scale-in connects AI factory and enterprise data to AI compute, and is powered by NVIDIA BlueField-4 and NVIDIA DOCA, and connected via NVIDIA Spectrum-X Ethernet..

Fig. 1: New definition of AI factory networking



Source: NVIDIA. Nomura Research

Fig. 2: Five-pillar of AI networking architecture



Source: NVIDIA, Nomura research

Scale-up requires extremely high bandwidth and low latency connections, with copper interconnects usually used for GPU link, while scale-out requires high bandwidth density switching architectures and low-latency communication, with the use of optical interconnects. Scale-across needs wide-area/campus optical interconnects, requiring ultra-high bandwidth and extremely low transmission latency jitter over long distances, as well as high reliability due to high maintenance costs. The key products in scale-across network include hollow-core fiber and coherent optical transceivers / equipment. A typical scale-across environment consists of multiple data centers are interconnected through high-speed optical links such as 800G/1.6T optical transceivers, hollow-core optical fiber, and DCI/WAN/metro routers and switches, forming a cross-site backbone.

Fig. 3: Scale-up vs. Scale-out vs. Scale-across vs. Scale-in

Category	Scale-up	Scale-out	Scale-across	Scale-in
Primary scenario	A single node (server/GPU box)	A cluster inside one DC	Multiple DCs as a compute pool	Interaction between AI cluster and external data
Scaling method	NVLink, NVSwitch, PCIe, UALink	Ethernet/Infiniband	Optical networking	NVIDIA BlueField-4, Spectrum-X Ethernet
Performance	Highest bandwidth and lowest latency	Aggregate performance across nodes; subject to network overhead	High capacity DCI links optimized for resilience and throughput; latency higher than intra-DC	NVIDIA BlueField-4 and Spectrum-X Ethernet Scale-In storage path delivers up to 1.45x the throughput of off-the-shelf Ethernet
Scalability	Limited by server chassis space, power, cooling and switch ASIC capacity	Vritually unlimited, constrained by network fabric design	Limited by DCI cost and optical reach	-

Source: Nomura research

Fig. 4: Primary technology comparisons in scale-up/scale-out/scale-across

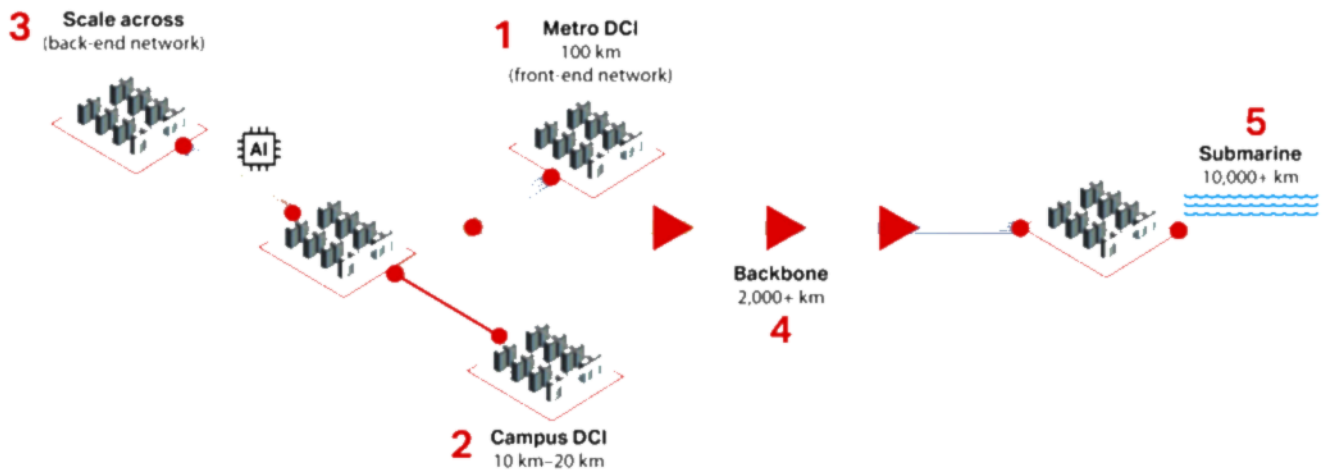
Level	Reach	Primary technology	Latency
Scale-up	<4m typical, up to ~10m	NVLink, UALink, Infinity Fabric, DAC/AEC copper, CPO	Sub-microsecond
Scale-out	100m to 2km	Infiniband NDR/XDR, Ethernet, Ultra Ethernet, IM-DD optics	Single-digit microseconds
Scale-across	>500m, up to 20km for synchronous training, up to 100km+ for async	Coherent pluggables module, DWDM, OCS, hollow-core fiber	Tens of microseconds to milliseconds

Source: Nomura research

Traditional AI clusters mainly rely on scale-up and scale-out networking, but as model training size exceeds trillions of parameters and inference demands grow exponentially, this architecture is beginning to face major limitations of bottlenecks in data center power and heating, and increasing network topology complexity. When AI clusters break through the physical boundaries of a single data center campus, scale-across becomes the inevitable choice, as scale-across can connect AI computing resources distributed across multiple data centers through high-speed optical networks to form a unified computing pool for supporting larger-scale AI workloads, in our view.

Traditional DCI is mainly used for data synchronization, backup, and content distribution across campus/metropolitan areas, with low bandwidth requirements (100G-400G), relying on DWDM and coherent optical transceiver technologies, emphasizing reliability and long-distance transmission. In 2025, NVIDIA CEO Jensen Huang proposed scale-across network to support distributed training, setting much higher requirements for bandwidth, latency, and packet loss rates, as compared to traditional DCI. According to Cisco, scale-across requires about 14 times the bandwidth required to connect 1 million XPUs compared to WAN/DCI networks.

Fig. 5: Data center connectivity—from 10 km to 10,000+ km



Source: Ciena, Nomura research

Scale-across market TAM analysis

Major technologies used in scale-across

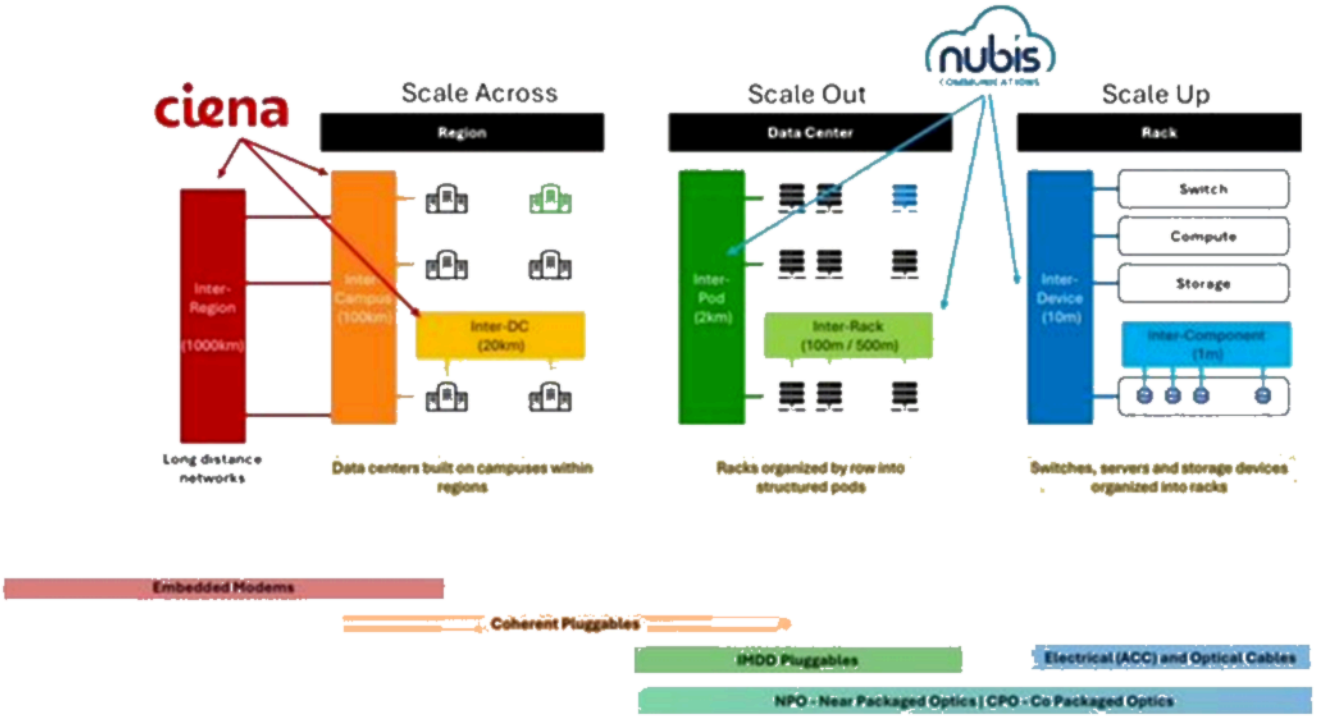
To achieve high-capacity, high-reliability, and low-latency scale-across transmission, system architecture places high demand on underlying hardware such as switches, optical fibers, and optical transceivers. As a result, we expect demand for technical solutions such as multi-rail line system, hollow-core fiber, and coherent optical transceiver to accelerate.

As the AI training environment is becoming geographically distributed, the capacity required to interconnect the scale-across network can reach tens of petabits per second across hundreds of fiber pairs, which would require many additional line amplification sites , along with large increases in space and power that are not viable in traditional approach. Multi-rail describes a new generation of photonic line systems designed to support multiple fiber pairs, with each rail functioning as a high-capacity optical highway with dedicated amplification, monitoring, and control. Multi-rail presents multiple benefits including maximized capacity per rack, reduced need for additional line systems and sites, improved space and power efficiency, and simplified operations at scale. At the scale-across level, multi-rail combines hundreds of optical fibers in parallel, aggregates ultra-high bandwidth, which may greatly increase deployment density in scale-across network, in our view. According to Ciena (CIEN US, Not rated), this technology will increase the number of single-rack fiber pairs by 32 times, further boosting demand for high-speed coherent optical transceivers and high-density optical amplifiers.

In terms of optical fibers, the delay caused by the inherent refractive index of glass in traditional solid-core single-mode fibers limits the synchronization of microsecond-level parameters in distributed data centers. Hollow-core fibers use air, inert gases, or vacuum instead of quartz as transmission media, resulting in a lower refractive index than solid-core glass. According to FiberHome (600498 CH, Not rated), its new generation of ultra-low-loss hollow-core fiber transmission latency is 31% lower than traditional solid-core fiber. Hollow-core optical fibers offer advantages of low latency, low loss, and high capacity, making it an ideal medium for long-distance and low-latency transmission.

Currently, the mainstream optical transceivers in AIDC are using the IM-DD (Intensity Modulation-Direct Detection) modulation technology, which only modulates light intensity at the transmitter. On the other hand, the coherent optical transceivers use IQ (In-Phase Quadrature) modulation at the transmitter to simultaneously control the amplitude and phase of the light and the receiver introduces a local oscillator laser, to greatly increasing information carrying capacity. The local oscillator at the receiving end of the coherent optical transceiver interferes with the signal light to demodulate, overcoming signal damage such as dispersion and attenuation caused by long-distance transmission. Therefore, the transmission distance of IM-DD optical transceivers is less than 2 km, while coherent optical transceivers can be extended to 80-120 km, making them more suitable for long-distance scale-across scenarios. The core components of coherent optical transceivers include narrow linewidth lasers, local oscillator lasers, modulators, mixers, and high-performance coherent DSP chips. The concept of coherent-lite was first proposed by Google in 2021, which utilizes the three dimensions of light like coherent optical transceivers, but uses a simplified DSP instead, which is cheaper and has lower power consumption than long-range coherent modules. Coherent-lite offers transmission distances ranging from 2–20 km, covering interconnection between data centers and campuses. We expect this technology to be applied in different network sections such as the scale-out (spine layer), scale-across, and DCI scenarios.

Fig. 6: Optical interconnect technologies in different layer of AI networking



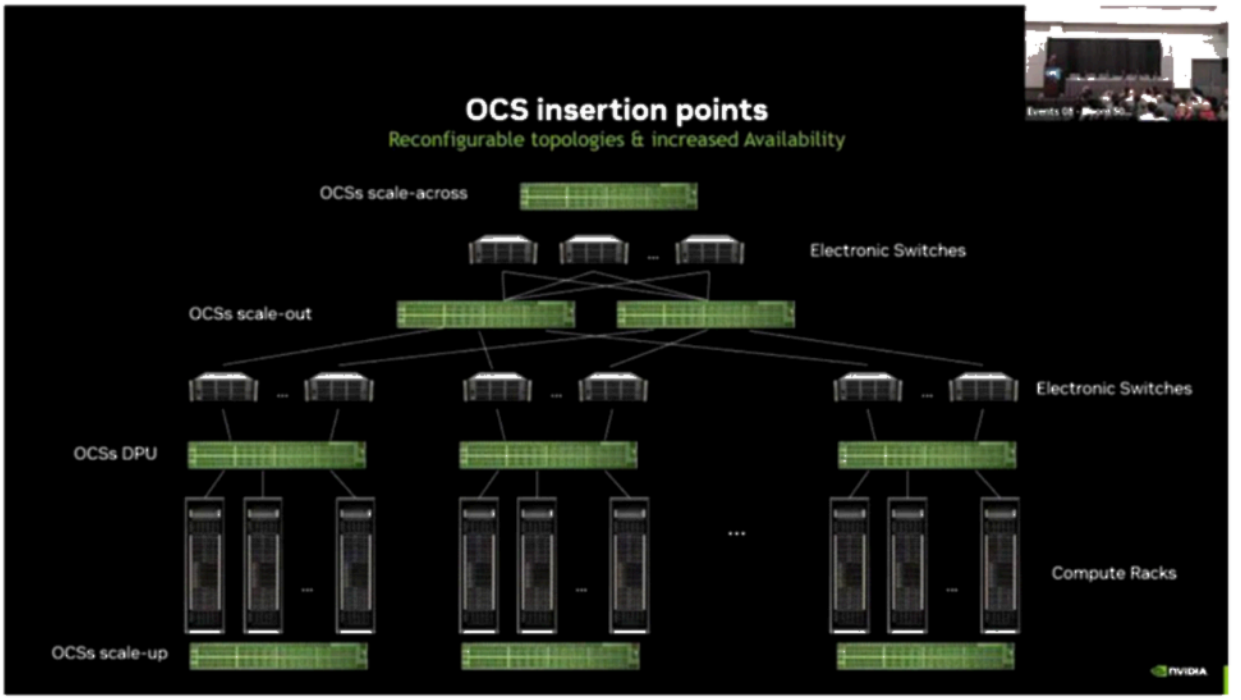
Source: Ciena, Nomura research

Furthermore, the use of coherent lite technology together with high bandwidth of OCS may fulfill the needs of ultra-large-scale AI clusters at the scale-across level. Through a simplified coherent detection scheme, coherent-lite can recover the link budget margin consumed by OCS, allowing the OCS architecture to continue scaling into the 2.4T/3.2T module era. Additionally, Lumentum (LITE US, Not rated) and Marvell jointly showcased the Lumentum R300 OCS platform and Marvell's Aquila 1.6T coherent-lite DSP, 1.6T PAM4 DSP, and COLORZ 800 ZR/ZR+ DCI modules at OFC 2026, demonstrating the feasibility of OCS+ coherent-lite.

Besides Google, NVIDIA has made it clear that OCS can be used for all-optical switching at the scale-across layer, replacing traditional electrical switching to support the efficient interconnection of 100,000-card-level AI clusters. For scale-across, NVIDIA has also introduced Spectrum-XGS Ethernet, which is not just a single hardware product, but a complete solution based on the collaboration of software algorithms and hardware. Spectrum-XGS hardware configuration primarily includes:

- Host network cards and accelerator cards: ConnectX-8 and BlueField-3 SuperNICs. The ConnectX-8 supports 800 Gb/s and, when combined with 1.6T optical modules, ensures high-bandwidth transmission.
- Long-distance optical transceivers for cross-data center connectivity: To interconnect geographically dispersed data centers, coherent transceivers can be used.

Fig. 7: Topology for OCS deployment designed by NVIDIA

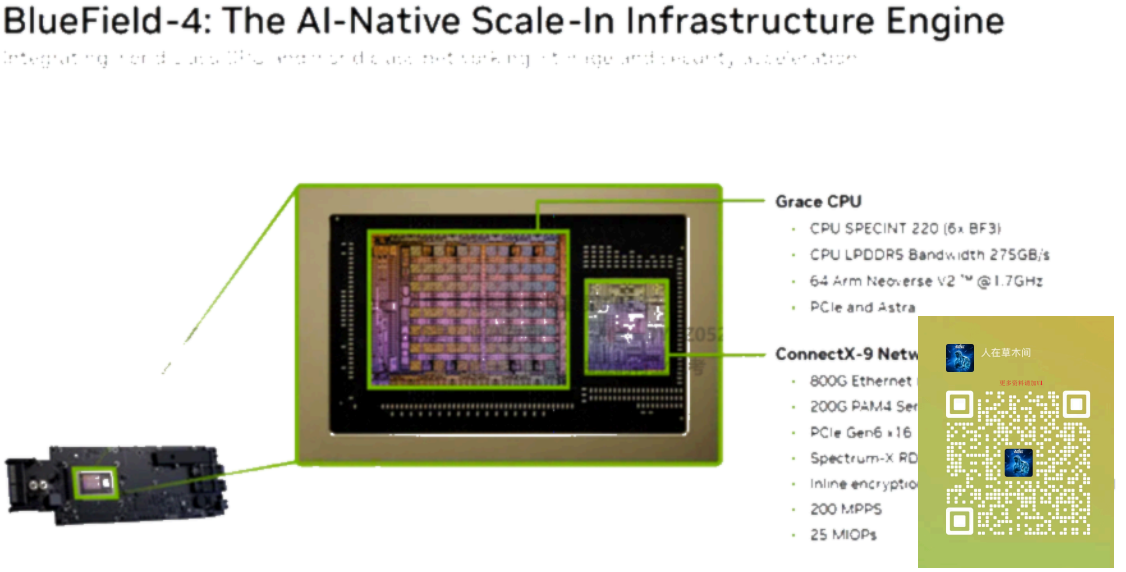


Source: NVIDIA, Nomura research

To realize the vision of "AI Gigafactory," Microsoft (MSFT US, Not rated) has designed a Fairwater series of data centers, deploying NVIDIA's GB200/GB300 NVL72 rack-scale systems, scalable to hundreds of thousands of Blackwell GPUs, and connecting two intercontinental data centers 1,000km apart via Microsoft's specially designed AI wide area network (WAN), deploying 120,000 miles of dedicated fiber optic cables.

In scale-in network, the core hardware is the BlueField-4 (BF-4) DPU, which provides hardware-level Zero Trust isolation and multi-tenant storage offloading, delivering 18X of AI Factory service bandwidth and 2X of storage acceleration performance, freeing host CPUs and GPUs from chores like network scheduling, security review, and virtualization management. BF-4 is defined by NVIDIA as an "AI-Native Scale-In Infrastructure Engine", integrating Grace CPU and ConnectX-9 (CX-9) at the packaging level. BF-4 is designed specifically for Agentic AI workloads, serving as the third major computing pillar, independent of CPUs and GPUs, specifically taking over data center networks, storage, and security governance.

Fig. 8: Components of BlueField-4



Source: NVIDIA, Nomura research

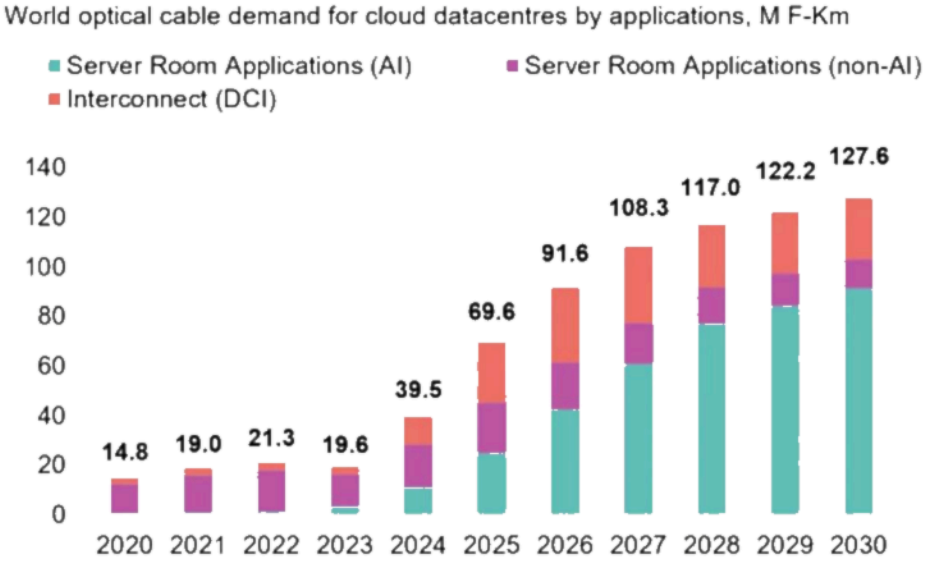
Scale-across TAM and value chain analysis

The total addressable market (TAM) of AI-WAN, which is inclusive of network, DCI and scale-across fabrics, is projected to increase from USD7bn to USD14bn, representing ~ 30% CAGR for 2026-2029E, according to Ciena. Ciena management also expects the scale across market to reach USD8-10bn by 2029E. Ciena noted in its 2026 investor forum that its scale-across TAM for a particular 80km span (requires 400+ fiber pairs) was about USD500mn, or USD1,000 per GPU on a 500k GPU cluster deployment (not include amplifiers). For a 1,000km span, which requires inline amplifiers, the TAM doubles to USD1bn, and could further jump to USD2bn for a 2,000km span, which reflects a shift from ZR/ZR+ pluggables to high-performance modems at the end points.

For coherent optical transceivers, Acacia (a Cisco subsidiary) pointed out that coherent pluggable solutions will play an important role in supporting AI scale-across networks,. with 800G ZR/ZR+ expected to become the main product, and the 1.6T ZR/ZR+ products are expected to become key support for next-generation cross-domain interconnection. Cisco management noted that Acacia had its strongest quarter with more than USD1bn in orders in the March-26 quarter and another USD1bn+ quarter in the June-26 quarter, driven mainly by 800G ZR/ZR+. Cignal AI estimates that the 800G ZR/ZR+ market size will have a CAGR of 145% from 2025 to 2029E.

According to the UK Commodity Research Agency (CRU), demand for fiber optic cables for AI workload likely grew 77% y-y in 2025E, far outpacing non-AI applications, and is projected to have a five-year CAGR of 26% during 2025-29E. Driven by demand from North American CSPs, hollow-core fiber begun commercial use in 2025. Currently, the core suppliers of hollow-core optical fiber include YOFC, Hengtong Optoelectronics (600487 CH, Not rated), Lightera (Unlisted, a subsidiary of Furukawa Electric (5801 JP, Not rated)), and Relativity Networks (Unlisted).

Fig. 9: Fiber optic cables for AI workload market size forecast



Source: CRU, LightCounting, Nomura research

For he OCS market, the demand was dominated by Google's MEMS OCS in 2025, with an overall market size of about USD400mn, according to Cignal AI's forecasts, and the OCS market size will exceed USD2.5bn in 2029E, exhibiting a CAGR of ~58% over four years. Meanwhile, Coherent (COHR US, Not rated) raised its 2030E global OCS market size guidance from USD2bn to USD4bn at the OFC 2026 investor meeting.

China scale across market vs global scale across market

The overseas scale-across market, especially in North America, is driven by the power and land constraints of individual data centers and the computing power bottlenecks of individual data centers. Such a situation is driving overseas CSPs to accelerate the deployment of AI infrastructure for scale-across networking, such as hollow-core fiber, coherent optical modules and OCS. In China, although electricity supply is relatively sufficient, it also faces physical

limitations in power and heat dissipation at a single data center as demand for LLM training & inference grows exponentially, thereby leading to the demand for the scale-across architecture.

Leading players in China's AIDC market have already started to work on the scale-across network technologies, products, and solutions. For example, China Mobile (941 HK, Buy) released its original scale-across technology—GSE-DCI (Globally Scheduled Ethernet Data Center Interconnect) at MWC 2026, and publicly showcased for the first time a 115.2T scale- across router. This device can achieve training computing efficiency over 98% per node over 100km. In 2026, ZTE (763 HK, Neutral) also released a new generation of 800G intelligent computing switches and the scale-across computing interconnection solution, equipped with independently developed large-capacity switching chip that optimize end-to-end transmission latency by about 33%. Meanwhile, the number of supporting networking equipment, optical modules, and optical fibers has been reduced by about 40%, and the solution has completed

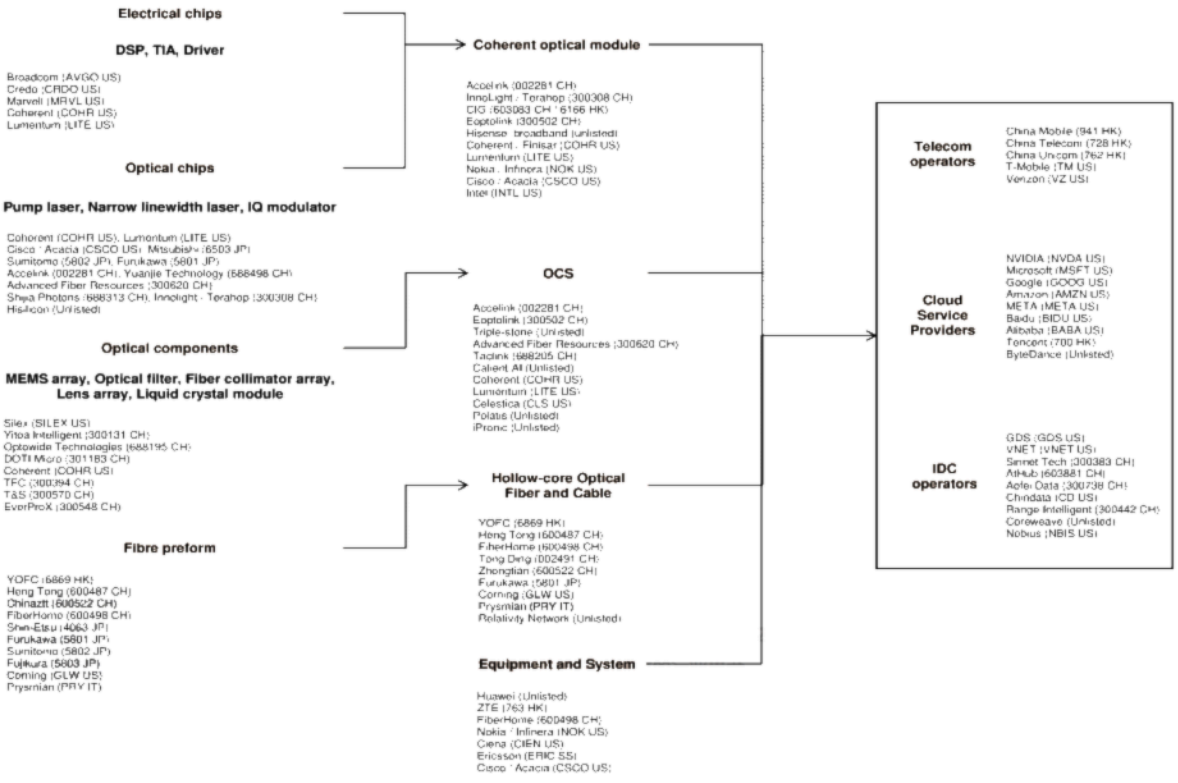
commercial verification for telecom operators on the network. We think that scale-across. business currently accounts for a relatively small share in China, compared to the global market, and the demand for large-scale applications has not yet reached a turning point.

Scale-across market competition landscape

The core of the scale-across industry chain is based on DCI technology, which is undergoing a restructuring due to the higher performance demands of AI scenarios. Its upstream mainly consists of optoelectronic chips, optical components, new types of fiber optics, and optical transceivers, while the midstream consists of networking and equipment manufacturers, including switches, routers, optical network system. Downstream are CSPs, IDC service providers, and telecom operators, who build cross-regional computing networks.

In terms of competitive landscape of pluggable coherent transceivers, Marvell and Acacia remain the leaders in 400ZRx with movement likely as Ciena and Nokia are likely to begin shipments of 800ZRx soon. According to Cignal AI 3Q25 optical component report, Marvell took the lead in coherent 400G+ market, followed by Acacia, Ciena, Nokia and Huawei. For optical fiber, Corning (GLW US, Not rated) held the highest share of the global fiber optic cable market in 2025, accounting for about 19.52%, followed by YOFC (6869 HK, Buy), with a market share of about 14.57%, and Zhongtian Technology (600522 CH, Not rated)) ranked third with a market share of about 11.43%, according to CRU.

Fig. 10: Scale-across value chain



Source: Nomura research

Coherent DSP is the core component of coherent optical transceivers and determines the critical performance of the transmission system. Vertically integrated companies like

Ciena, Infinera (acquired by Nokia [NOK US, Not rated]), and Huawei (Unlisted) have developed their own DSPs for high-performance long-distance networks and generally for internal use, while Marvell (MRVL US, Not rated), Broadcom (AVGO US, Not rated), and Acacia (acquired by Cisco [CSCO US, Not rated]) dominate the coherent optical transceivers market with commercial DSPs. Acacia is a pioneer of pluggable coherent optical transceivers, pioneering the launch of an 800G ZR+ optical transceiver with interoperable PCS functionality, improving the signal-to-noise ratio for long-distance transmission. In 3QFY26, Acacia shipped over 750,000 units for 400G ports and over 40,000 units for 800G ports, according to company. Ciena is the first to enter the 1.6T era, launching the industry's first 1.6T coherent technology using a 3nm CMOS process. In China market, Accelink uses Marvell's DSP to launch a 1.6T coherent-lite optical transceiver, making it one of the first domestic 1.6T coherent optical transceiver.

Fig. 11: DSP and Coherent optical module product progress

Company name	Product	Characteristic
Marvell	1.6T ZR/ZR+ coherent module with 2nm coherent DSPs	- Support scenarios including campus (20km), metro (120km), and regional (1,000km) distances; - Products expected to begin sampling in the second half of 2026
Infinera/Nokia	1.2T coherent DSP	- 5nm process; - Can deeply integrated with photonic integrated circuits (PIC) to form an embedded coherent optical engine
	1.6T coherent optical module	- Optimized for the deployment of streamlined wavelength converters, supporting a wide range of land and submarine cable network applications
	3.2T coherent-lite optical module	- Optimized for low power consumption to meet the needs of short-distance campuses and enterprise applications; - The coherent optical solution family is expected to begin prototyping in mid-2027, with official commercial deployment since 2H2027
Ciena	WaveLogic 6 Nano 800G coherent optical module	- Support both C-band and L-band
Acacia (Cisco)	Delphi DSP	- Support metro and regional DCI upgrades with 800G I/O ports
Accelink	1.6T O-band Coherent-Lite module	- adopt Marvell Aquila Coherent-lite DSP, offering interconnection distances of up to 20km
Terahop (Innolight)	800G LR2 OSFP Coherent-lite module	- Uses the Marvell Aquila DSP for distributed campus data center interconnects
FiberHome	Coherent DSPs for 100G~800G	- It has achieved mass production of DSPs and 800G coherent optical module
ZTE	800G coherent DSP	- It has already been commercialized on a large scale
Taclink	EDFA (fiber amplifiers)	- Coherent signals become severely attenuated after long-distance transmission of 80-120 kilometers and must be amplified in the optical domain via EDFA
	400G/800G coherent optical modules	- 400G coherent modules have achieved small-batch shipments, while 800G coherent devices and modules are in the pre-research stage

Source: Company data, Nomura research

In terms of demand for optical fibers, a single GW in data center requires several million to 10mn fkm of optical fibers, according to YOFC. CPO uses polarization-maintaining fiber in scale-up network, and YOFC management estimates that demand for polarization preserving fibers may grow by 10-20x over the next 1-2 years and may achieve several million fkm demand by 2030F. Scale-out network uses single-mode G.657.A1 fiber, multi- mode OM4 and OM5 fibers, Meanwhile, scale-across network may use hollow-core fiber in the future, but the volume should be relatively small (several million fkm by next 2-3 years), in our view.

North American CSPs have started deployed hollow-core fiber for scale-across network
By 2026, AWS (AMZN US, Not rated) has successfully deployed hollow-core fiber to its 10 core data centers. Besides AWS, Microsoft (MSFT US, Not rated) acquired UK high- tech company Lumenisity (Unlisted) as early as 2022, which specializes in hollow -core fiber technology development, and plans to deploy 15,000km of hollow-core fiber across Azure's global network by the end of 2026, while Google (GOOG US, Not rated) and Meta (META US, Not rated) have also begun related technology testing. For Chinese optical fiber suppliers, YOFC has achieved mass production of hollow-core fiber, with the lowest attenuation coefficient for its developed products, and has already delivered hollow-core fibers for the three telecom operators. Hengtong Optoelectronics is investing in building an AI advanced fiber optic materials R&D center, expected to reach 5mn fkm after completion in February 2026.

Fig. 12: Progress of hollow core optical fiber

Company name	Progress of HCF
Corning	Announced cooperation with Microsoft and Heraeus to jointly promote large-scale production of hollow-core fibers in Sep. 2025
Prysmian	A long-term cooperation agreement has been reached with Relativity Network, to mass-produce hollow-core fiber optic cable product
YOFC	Independently developed world-leading hollow-core fiber technology, achieving continuous transmission of over 20km per fiber with a decay coefficient of 0.05dB/km. Since 2026, the public tender projects for telecom hollow-core optical fibers have reached a scale of tens of millions, and YOFC has won all of them.
Hengtong	The company has established an hollow-core fiber optic technology platform and is laying out the preparation of key raw materials and manufacturing technologies. The company is collaborating with companies such as China Mobile on demonstration project applications
FiberHome	The company developed single-bridge structure hollow-core optical fiber in 2024, and granted a new patent titled "An Hollow-Core Anti-Resonant Active Optical Fiber" in April 2026

Source: Company data, Nomura research