BERNSTEIN
SOCIETE GENERALE GROUP

MAY 2026

ARTIFICIAL INTELLIGENCE: INSIDE THE WAR FOR AI DATA CENTER CONNECTIVITY

From PCB to copper to optical links, connectivity is the backbone of AI performance.

Connectivity is the new bottleneck and battleground. AI data centers are shifting from compute-bound to connectivity-bound architectures, elevating the strategic importance of electrical and optical interconnects across both scale-up and scale-out networks.

AI connectivity lifts board, substrate, and materials content. Higher speeds and densities are driving structural growth in multi-layer printed circuit board, HDI, and ABF substrates, while tight supply of critical materials reinforces near-term earnings momentum despite competitive pressures beyond 2026.

CPO's promise is clear and will start from scale-out. Co-packaged optics (CPO) deliver superior power efficiency and signal integrity, yet face cost, yield, and serviceability hurdles. We see scale-out deployments from 2H26, with copper remaining dominant in scale-up over the next three years.

Value creation spans a complex, evolving supply chain. The transition toward advanced connectivity reshapes economics across optical engines, lasers, fiber array units, PCBs, and testing equipment, favoring technology leaders with manufacturing depth and ecosystem control.

Published: 08 May 2026 20:05 UTC | Revised 08 May 2026 20:05 UTC

OVERVIEW

AI data centers are no longer constrained by compute power alone; data movement within and between racks has become the primary bottleneck. Connectivity is now a key performance and investment frontier across both electrical and optical domains. NVIDIA's roadmap confirms that copper and optical interconnects will coexist, each optimized for different distance and power needs. For investors, the focus is not which technology wins, but how each captures value across the supply chain and over what timeline. This *Whitebook* outlines the key technology trends, economics, and supplier positioning shaping scale-up and scale-out connectivity.

Chapters titled "Co-Packaged Optics: Inside the War for Data Center Connectivity" and "Mapping the CPO Value Chain" of this *Whitebook* explore the strategic value, adoption schedule, and supply chain of co-packaged optics (CPO). CPO integrates optical engines (OEs) closer to switch/AI chips, improving power efficiency and signal integrity. But, adoption faces challenges in serviceability, manufacturing yield, testing complexity, and supplier concentration. We expect small-volume CPO deployment in scale-out networks starting 2H26, with copper remaining mainstream for scale-up over 2026-27. Also, CPO cost is likely ~10% over pluggable optics after system margins. CPO manufacturing and testing are challenging, favoring technology leaders. We believe TSMC's COmpact Universal Photonic Engine (COUPE) will become the main platform. The transition to CPO will likely benefit OE, laser, fiber array unit, and equipment vendors (such as Chroma). For scale-up solutions, copper will likely remain the mainstream over 2026-28, and co-packaged copper (CPC) is likely to extend its lifecycle, benefiting players such as Luxshare.

Chapters titled "AI PCB Primer: Technology and Supply Chain" and "AI PCB Primer: Players, Positioning, and Profitability" of this *Whitebook* explore the advancement of printed circuit board (PCB) and integrated circuit (IC) substrate sectors amid the AI wave. Within AI server racks, connectivity advances also hinge on boards and substrates that sustain higher signal speeds and densities, lifting content value for multi-layer PCB, high-density interconnect (HDI) boards, and Ajinomoto Build-up Film (ABF) substrates. Critical materials such as T-glass, ABF film, high-end copper foil, and copper-clad laminates (CCLs) are shaping supply-chain dynamics. As of 2025, AI server demand has already fueled 50-100% revenue growth for several PCB and materials suppliers. Margin expansion may moderate from late 2026 amid intensifying competition, capacity additions, and rising depreciation. Nevertheless, new product cycles (e.g., midplanes and backplanes), tight T-glass supply, and growing Application-specific Integrated Circuit (ASIC) demand should sustain earnings momentum for technology leaders.

Alex Wang	alex.wang@bernsteinsg.com	+852 2123 2613
Stacy A. Rasgon, Ph.D.	stacy.rasgon@bernsteinsg.com	+1 213 559 5917
Mark Li	mark.li@bernsteinsg.com	+852 2123 2645
David Dai, CFA	david.dai@bernsteinsg.com	+852 2918 5704
Euan McLeish	euan.mcleish@bernsteinsg.com	+81 3 5962 9611
Shirley Yang	shirley.yang@bernsteinsg.com	+852 2123 2660
Ethan Xu	ethan.xu@bernsteinsg.com	+852 2123 2634

May 9, 2026

TABLE OF CONTENTS

CO-PACKAGED OPTICS: INSIDE THE WAR FOR DATA CENTER CONNECTIVITY

COPPER AND OPTICAL TECHNOLOGIES EVOLVE TOGETHER FOR AI DATA CENTER CONNECTIVITY

Why does CPO matter? By integrating optimal engines (OEs) with accelerated processing units (XPUs) or switch chips, CPO offers superior power efficiency, signal integrity, and network resiliency compared to pluggable optical transceivers. NVIDIA targets 3.5x more power efficiency and 10x better network resiliency compared to pluggable solutions, with initial deployment guided for late 2026 with three partners. Broadcom is already on its third-generation CPO switch and Lumentum (not covered) has confirmed several hundred million dollars of CPO product orders for shipment in 2027.

CPO adoption faces meaningful obstacles, including manufacturing yield, test complexity, fiber coupling, and cloud service provider (CSP) concerns around serviceability and vendor concentration. NVIDIA and Broadcom are expected to use TSMC's COUPE technology to bond electronic and photonic dies, with FAUs and micro-optics precisely aligned on top. Compared with pluggable optical transceivers, this approach improves performance, power efficiency, bandwidth density, and reliability. However, failures may require swapping an entire switch, increasing downtime and operational impact. The CPO market is also likely to be far more concentrated than today's pluggable optics ecosystem, reducing CSP bargaining power.

Mass adoption of CPO will likely take time, starting with scale-out. NVIDIA's "CPO" architecture is closer to near-package optics (NPO), which uses detachable OEs to improve serviceability. Interim solutions such as linear pluggable optics (LPO) can cut power consumption by roughly two-thirds versus traditional pluggables, providing a bridge while CPO matures. We expect CPO and NPO to begin shipping in small volumes in 2H26 for scale-out applications. By contrast, we remain more conservative on scale-up adoption, as the industry needs additional time to validate CPO reliability under real-world conditions before integrating it into high-value XPUs. That said, we expect increased visibility on CPO technology progress and customer adoption plans over the next 12 months.

We expect CSPs to favor pluggable optical transceivers over CPO in the next two to three years. While investors worry about the long-term role of pluggable vendors, new technologies in the AI data center, such as high-voltage direct current (HVDC) power, PCB backplanes, Chip-on-Wafer-on-Panel (CoWoP), and micro-channel lid, have consistently taken time to mature. **This is because reliability and supply-chain robustness matter as much as the theoretical benefits of new designs.** For scale-up solution, copper will likely remain the mainstream over 2026-27, and CPC is likely to extend its lifecycle, benefiting players such as Luxshare (Outperform).

COPPER AND OPTICS ARE THE MAINSTREAM CONNECTION SOLUTIONS IN AIDC

In recent years, AI advancement has closely followed the scaling law: increasing compute power, training data, and model parameters generally leads to better performance. Consequently, AI infrastructure has expanded rapidly through two distinct models — scale-up and scale-out. **Scale-up** refers to adding resources within a single system, such as installing more AI accelerators within a rack or node to boost performance. **Scale-out**, by contrast, expands capacity by connecting more racks to form larger clusters in an artificial intelligence data center (AIDC).

This exponential growth in computing power has driven unprecedented demand for interconnect bandwidth. Scale-up architectures, in particular, require much higher bandwidth than scale-out systems because large-language-model (LLM) training depends on parallel computing methods such as tensor and expert parallelism — both of which involve intensive data exchange within tightly coupled scale-up pods.

Currently, copper and optics represent the two primary connection methods in AI data centers. Copper interconnects are primarily utilized for intra-rack connectivity in scale-up scenarios. On the other hand, optical links deliver high bandwidth (≥224Gbps) and long reach (≥10 meters) with low attenuation and terabit scalability. Therefore, they serve as the mainstream solution for inter-rack links in scale-out architectures. In **NVIDIA's GB300**, for example, communication between the Superchip (GPU + CPU) and the switch chip relies on copper links, while the ConnectX network interface card (NIC) on the compute tray connects to 800Gb/s optical transceivers. These transceivers convert electrical signals into optical signals and link the system to either Quantum-X InfiniBand or Spectrum-X Ethernet fabrics for scale-out connectivity.

Over 2000-25, switch capacity grew at 40% a year with no sign of deceleration amid the rapid evolution of AI accelerators and the surge in data movement. Serializer-Deserializer (SerDes) drives short-reach electrical traces or cables into an optical module that then converts electrical signals to optical signals for longer-distance links. SerDes per-lane I/O bandwidth has increased much more slowly than overall switch throughput, forcing vendors to add more lanes per switch. As a result, the number of optical links and modules grows even faster than switch shipments themselves. Annual shipments of 800G modules hit 10 million by 2025, marking a new record, compared to a decade for 100G modules. With 1.6T transceivers now in mass production, Coherent (COHR, not covered) and LightCounting expect it will take only four years to reach 10 million units sales, creating meaningful revenue upside for optical component suppliers by 2026 (Exhibit 1).

Driven by surging shipment volume and average selling price (ASP) tailwind thanks to the migration to high-speed modules, the optical transceiver market is experiencing rapid growth. LightCounting estimates the sales of optical transceivers and related products exceeded $23Bn in 2025, up 50% from 2024. Currently, the scale-out network market is primarily dominated by two competing technologies: InfiniBand (known for superior performance and led by NVIDIA's Mellanox) and Ethernet (offering an open ecosystem and significant cost advantages, led by Broadcom). LightCounting estimated Ethernet optical transceivers made up ~74% of the total in 2025, reaching $17Bn and up 60% YoY. Additionally, LightCounting forecasts that the Ethernet optical transceiver market will grow at a 59% CAGR in 2024-26, then decelerate to 15% in 2026-30 (Exhibit 2).

EXHIBIT 1: **Optical interconnects product mix will likely shift toward higher rates**

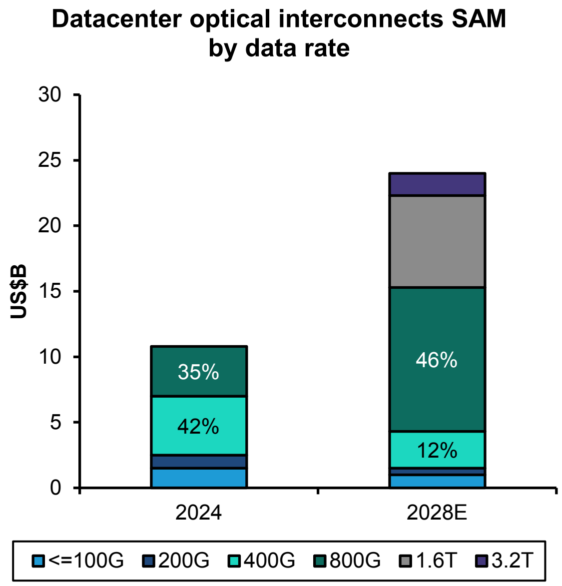

Note: Forecast from Coherent and LightCounting; SAM refers to Coherent's Service Addressable Market.
Source: Coherent, LightCounting, Bernstein analysis

EXHIBIT 2: **Ethernet optical transceiver market is projected to grow at a 59% CAGR over 2024-26, then decelerate to a 15% CAGR over 2026-30**

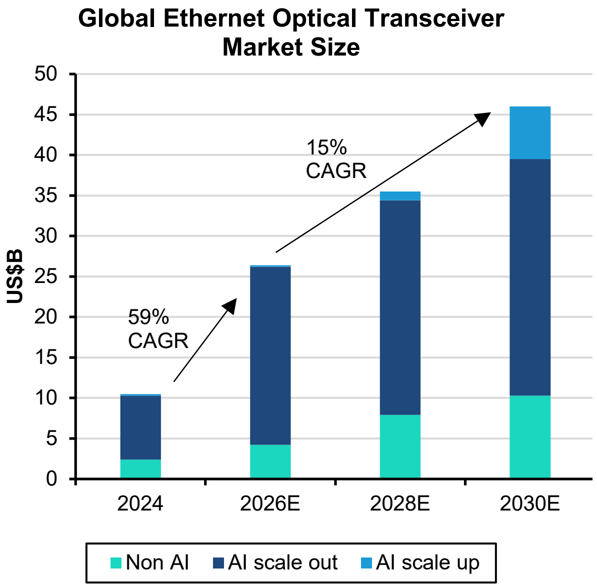

Source: LightCounting data and estimates, Bernstein analysis

WHY DO WE NEED CPO?

For pluggable optical transceivers, the transceiver cage mounted on the front panel of a switch or compute tray is typically positioned 15-30cm away from XPUs or switch chips. Therefore, a digital signal processor (DSP) is needed to recover and condition the signal. However, this performance comes at the cost of higher power usage. NVIDIA estimates that a 1.6Tbps module may consume ~30 watts, with the DSP accounting for more than half of that.[1] Traditional pluggable solutions rely on labor-intensive assembly and expensive components such as DSPs, retimers, and electro-absorption modulated laser (EML) devices, which drive up bill of materials (BOM) costs and constrain scalability.

Additionally, optical modules typically have a hard failure rate of ~100FIT (failure in time), which equates to about nine failures a year per 10,000 optical modules. Soft failures caused by dust contamination or connector and interface sensitivity will likely create further uptime and reliability issues that require manual intervention. These limitations in traditional pluggable optics are pushing the industry's shift toward next-generation connectivity solutions for AI data centers.

To boost power and thermal efficiency, the industry is moving toward CPO solutions. Under this new packaging technique, optical engines are integrated alongside chips on the same substrate, which enhances signal integrity and thereby eliminates DSPs, enabling data transmission via lower-power SerDes instead. NVIDIA claims that its CPO-based switches deliver 3.5x more power efficiency, 63x greater signal integrity, 10x better network resiliency at scale, and 1.3x faster deployment, compared to traditional optical transceivers (Link[2]). Broadcom

---

[1] A New Era in Data Center Networking with NVIDIA Silicon Photonics-based Network Switching | NVIDIA Technical Blog.

[2] https://nvidianews.nvidia.com/news/nvidia-spectrum-x-co-packaged-optics-networking-switches-ai-factories

reports that its CPO switch delivers 40% lower optics cost per bit (Link[3]), likely helped by reduced component count and the replacement of costly EML lasers with continuous-wave (CW) lasers, and a more streamlined semiconductor-based design and manufacturing process. Meta's testing results[4] of Broadcom's Bailly CPO switch suggests 500W+ power saving in a 51.2T switch system (Exhibit 3 and Exhibit 4).

EXHIBIT 3: **Traditional switch with pluggable optical transceivers**

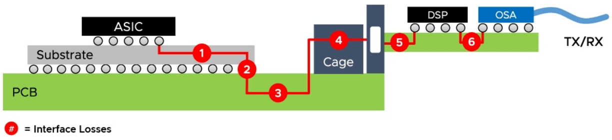

Source: Broadcom

EXHIBIT 4: **CPO switch**

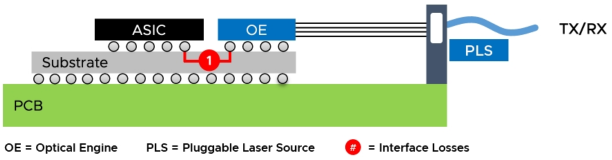

Source: Broadcom

CPO IN SCALE-OUT AND SCALE-UP

Based on the industry roadmap (Exhibit 5), CPO adoption will start in scale-out (switch), and move to scale-up connectivity (XPU) over the long term.

CPO adoption in scale-out network will start late 2026, but long-term TAM will likely hinge on its energy and cost advantages, system reliability, and how quickly pluggable optics continue to improve. According to the IEA, scale-out networking accounts for only about 5% of total cluster power, while servers consume roughly 75%. This means that even large power-efficiency gains from CPO switches would translate into only modest reductions in total data center power use (Exhibit 6).

Operational reliability is another concern for CSPs. Today, data center operators can hot-swap failed pluggable transceivers without affecting other ports. In contrast, a highly integrated CPO package — where silicon chips and optical engines are packaged together — cannot be easily separated for servicing, raising maintenance and downtime risks.

On cost, while CPO-based switches may offer lower BOM through simplified design and cheaper components, they also shift more value upstream to semiconductor suppliers, creating a more consolidated supply chain. The profitability requirements of semiconductor suppliers and transceiver companies are also very different. As a result, early-stage CPO ASPs are likely to exceed those of conventional pluggable solutions. Meanwhile, pluggable optics continue to advance, with 3.2T-class modules expected to arrive in the near future.

For scale-up, copper interconnects are expected to reach performance limits earlier than pluggable optical transceivers, rendering scale-up a more addressable market for CPO. Limited transmit range and bandwidth-upgrade impediments of copper links hinder the further expansion of AI clusters, which we elaborate on in the following section. In contrast, CPO could mitigate these bottlenecks by extending connectivity ranges and offering more bandwidth, thereby enabling the realization of super-large-scale clusters comprising over 1,000 XPUs.

---

[3] https://www.broadcom.com/info/optics/cpo

[4] Evaluation of CPO Performance and Pluggable Optics Health for Reliable AI Infrastructure.

EXHIBIT 5: **Broadcom's roadmap indicates that copper still consumes less power versus traditional CPO, while CPO may exhibit better power efficiency versus copper and optical solutions in 2028-29**

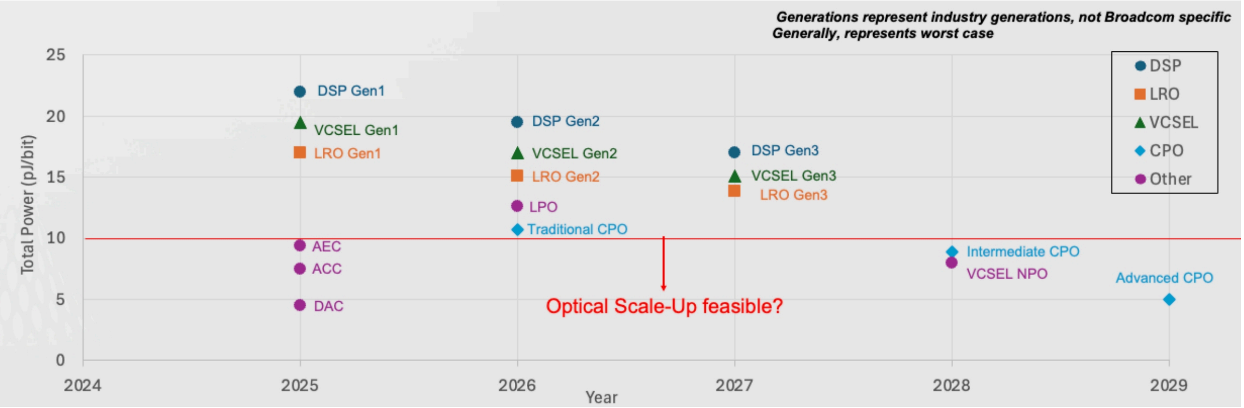

Source: Broadcom

EXHIBIT 6: **Network only makes up a small portion of total power consumption of hyperscale data centers**

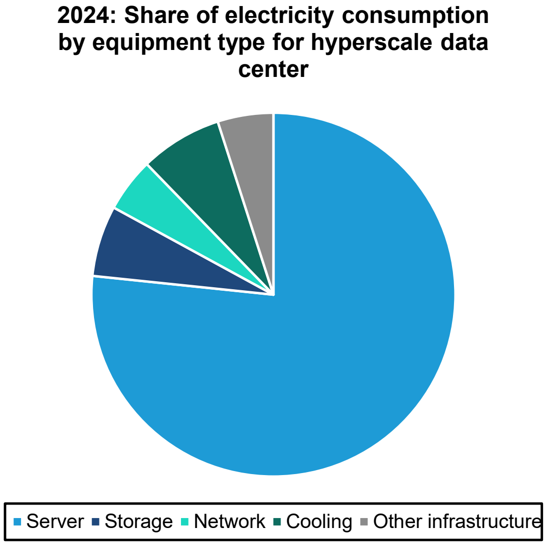

Source: IEA, Bernstein analysis

INSIDE NVIDIA AND BROADCOM'S RACE FOR THE CPO ERA

NVIDIA and Broadcom are the forerunners in the CPO industry. At the GPU Technology Conference (GTC) 2025, NVIDIA exhibited three distinct CPO-based switches, including Quantum X800-Q3450 (Exhibit 7 and Exhibit 8), Spectrum 6810, and Spectrum 6800. In addition, Broadcom, as one of the first companies to offer a CPO-enabled system, unveiled its first-generation CPO switch Humboldt in 2022, second-generation product Bailey in 2024, and expects its next-generation offering Davisson to reach commercial readiness in 2027. Broadcom shipped its early CPO systems to Tencent in small volumes, while its latest products have undergone testing at Meta.

We summarize the key metrics of NVIDIA's and Broadcom's CPO-based switches in Exhibit 9. (1) NVIDIA's CPO switch features a detachable optical engine, so technically it's closer to a NPO than to a true CPO. This approach improves serviceability by allowing easier repairs in the case of malfunction. In contrast, Broadcom's Bailey integrates its optical engine directly onto the substrate, making it non-detachable. (2) It's also noteworthy that Broadcom's Humbolt and Bailey (launched in 2022-24) are based on fan-out wafer-level packaging (FOWLP) from Siliconware Precision Industries Co., Ltd. (SPIL) (not covered), while the company will likely migrate to TSMC's CPO platform for future generations.

EXHIBIT 7: **NVIDIA Quantum-X800 InfiniBand CPO switch**

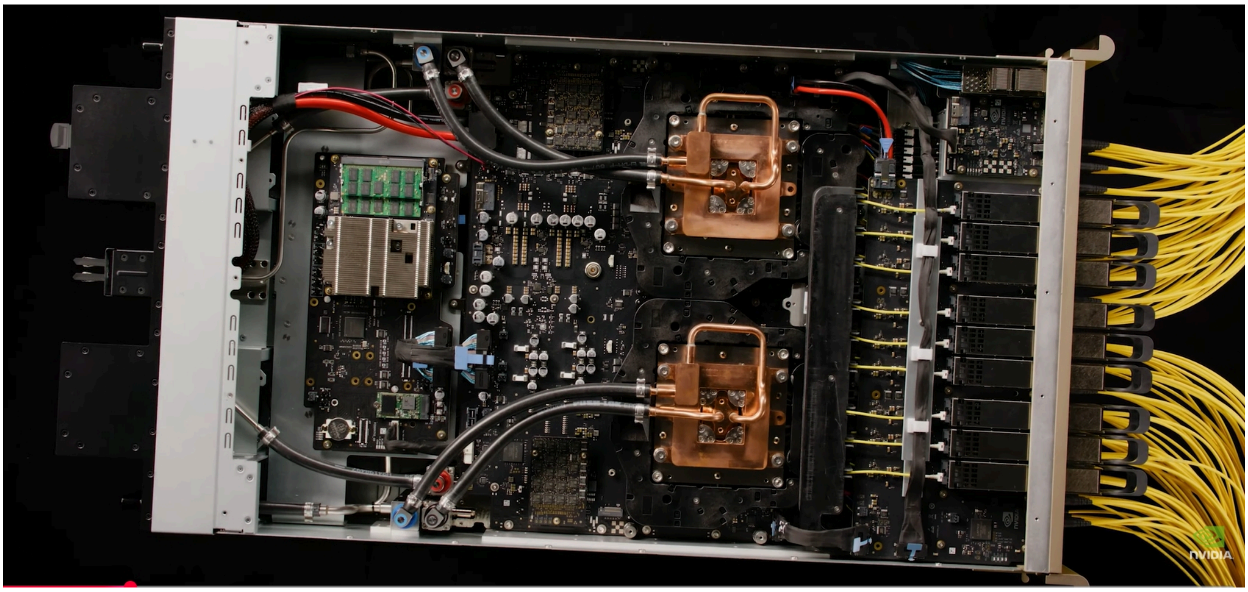

Source: NVIDIA

EXHIBIT 8: **The NVIDIA Quantum-X800 InfiniBand CPO switch comprises six optical subassemblies, with each subassembly housing three optical engines**

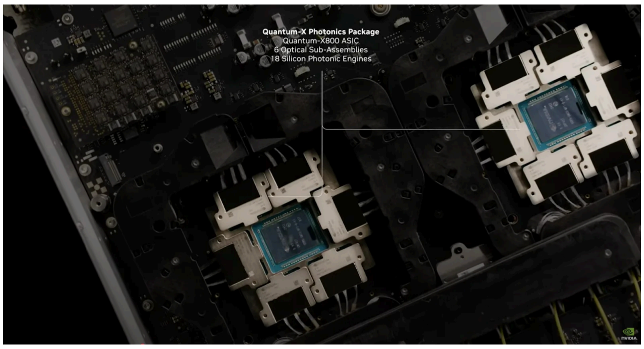

Source: NVIDIA

EXHIBIT 9: **NVIDIA versus Broadcom: CPO switch key metrics comparison**

	Nvidia			Broadcom		
	Quantum X800-Q3450	Spectrum 6810	Spectrum 6800	TH4-Humboldt DR	TH5-Bailly FR4	TH6 Davisson
Pilot Production date	2H 2025	2H 2026	2H 2026	2022	2024	2026
Networking standard	InfiniBand	Ethernet	Ethernet	Ethernet	Ethernet	Ethernet
Switch ASIC	Quantum-3	Spectrum-6	Spectrum-6	Tomahawk 4	Tomahawk 5	Tomahawk 6
Throughput per package	28.8Tps	102.4 Tps	102.4 Tps	25.6 Tbps	51.2 Tbps	102.4 Tbps
Number of switch package	4	1	4	1	1	1
Switch Aggregate Bandwidth	115.2 Tbps	102.4 Tpbs	409.6 Tbps	25.6 Tbps	51.2 Tbps	102.4 Tbps
SerDes speed (Gb/s unidi)	200G	200G	200G	100G	100G	200G
Optical connectivity	DR optics	DR optics	DR optics	DR optics	FR optics	DR4 optics

Source: NVIDIA, Broadcom, Bernstein analysis

MANUFACTURING STEPS OF CPO

TSMC's CPO technology as an example

There are many approaches to manufacture CPO, and we summarize TSMC's approach in this section as with the adoption by NVIDIA, Broadcom, etc., it is likely to be the mainstream approach for many years to come.

Exhibit 10 shows the overall advanced packaging platform that TSMC offers to AI and high-performance compute (HPC) applications, and TSMC's CPO technology is part of it. At the heart of TSMC's CPO technology, COUPE is what TSMC uses to manufacture an optical engine (OE). An OE is where optical and electrical signals convert into each other. TSMC first produces an EIC (electronic IC) on a wafer (6nm process now) and a PIC (photonic IC) on another (65nm silicon-on-insulator (SOI) process now). The EIC wafer is then cut into dies and tested. Then a known good die (KGD) of EIC will be flipped, and then face-to-face bonded to the PIC wafer using bumpless hybrid bonding (marketed as SoIC-X by TSMC). The PIC wafer is then cut into dies too, so that each PIC die will have connections with an EIC die on the top (Exhibit 11). The connections between them allow electrical and optical signals to couple and convert into each other, and the use of hybrid bonding shortens the connection distance. Additionally, TSMC puts silicon-based micro lenses on the top of the OE and a metal reflector at the bottom (Exhibit 12). These, together with the short distance enabled by hybrid bonding, help increase the conversion efficiency of the OE.

After the OE is ready, a Fiber Array Unit (FAU) (Exhibit 17) holding many optical fibers together, will be attached to the top of the OE, so that the light beams from the fibers after falling on a mirror will be reflected and shed onto the micro lenses precisely (Exhibit 13).

Though the output of the COUPE process described, in theory, can be also adopted in DSP pluggable and other configurations (Exhibit 14), CPO, enabled by TSMC's chip-on-wafer-on-substrate (CoWoS) technology, is clearly the primary target, so that the OE and switch chips or XPUs are co-packaged to become CPO. TSMC further breaks down the integration of COUPE with CoWoS into two phases. The first is to package the OE on the substrate of the CoWoS structure. This mostly likely will be outsourced to ASE (3711 TT, not covered) as TSMC has been outsourcing most of the "oS" part of the CoWoS production to ASE, and we expect the practice to continue. In the second phase, the OE will be attached on an interposer. For this phase, TSMC will most likely be handling the attachment as the current practice is of TSMC handling all of the "CoW" part of CoWoS internally.

Also note, COUPE can be employed regardless of the use of CPO. For example, before the integration of COUPE and CoWoS is production-ready, OEs made of COUPE can also be packaged in a pluggable format and then placed on a PCB. This is called NPO or on-board optics (OBO), and is an intermediate step before CPO is ready. COUPE's transition to the CPO format can be further divided into putting OE on substrate first, likely in 2026, and then on interposer later as TSMC gradually overcomes the technology complexity and brings OEs closer to processor.

In terms of coupling technology, COUPE supports both grating coupling (GC) and edge coupling (EC), but TSMC and NVIDIA currently focus on grating coupling.

The roadmaps (Exhibit 15 and Exhibit 16) presented by TSMC in August 2025 showed that its CPO will be first adopted by switch chips with OE on substrate, and then by XPU with OE on interposer. Another roadmap (Exhibit 14) presented in September 2025 revealed the possibility of having CPC before CPO is adopted by switch chips.

Competitively, Intel also has silicon photonic and CPO technologies (webpage[5]). GlobalFoundries offered CPO services to Ayar Labs earlier, but lost out to TSMC in the third-generation of Ayar Labs' products. ASE also packaged Broadcom's earlier CPO products. However, considering the need to integrate with other advanced packaging technologies (e.g., CoWoS) and advanced front-end node (e.g., 6nm for EIC and above), and the scale and economy jointly established by TSMC and NVIDIA, we believe TSMC's CPO offering is more competitive and will solidify TSMC's position as the default manufacturer of AI processors and networking chips.

EXHIBIT 10: **TSMC's advanced packaging technologies offered for HPC/AI applications, with COUPE being part of it and shown on the right**

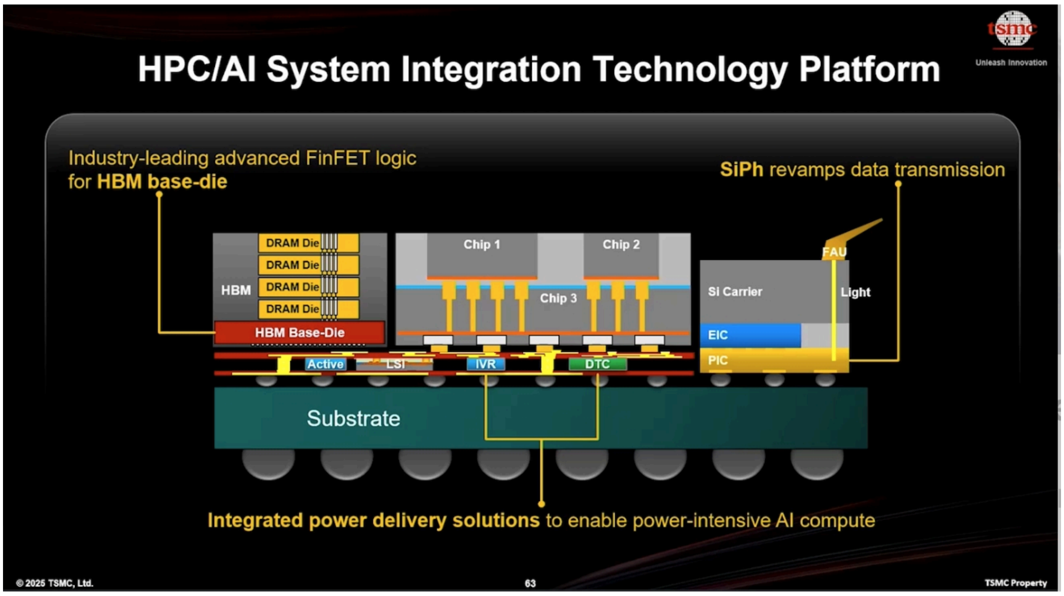

HPC/AI System Integration Technology Platform

Industry-leading advanced FinFET logic for **HBM base-die**

SiPh revamps data transmission

Integrated power delivery solutions to enable power-intensive AI compute

Source: TSMC

---

[5] https://www.intel.com/content/www/us/en/products/details/network-io/silicon-photonics.html

EXHIBIT 11: **In COUPE, TSMC creates an optical engine (OE) by bonding EIC (electronic IC) and PIC (photonic IC) together through die-to-wafer face-to-face hybrid bonding**

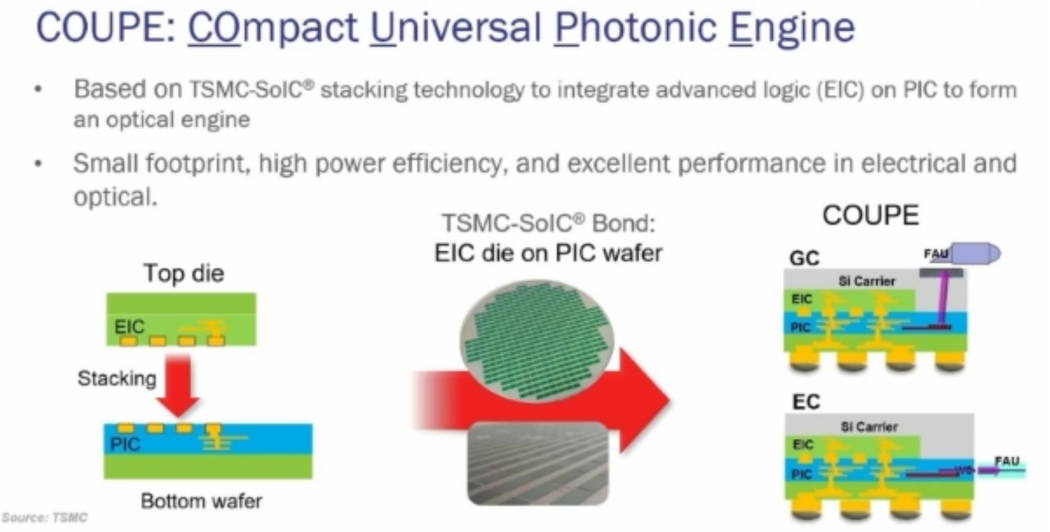

COUPE: COmpact Universal Photonic Engine

- Based on TSMC-SoIC® stacking technology to integrate advanced logic (EIC) on PIC to form an optical engine
- Small footprint, high power efficiency, and excellent performance in electrical and optical.

Source: TSMC

EXHIBIT 12: **In COUPE, TSMC puts silicon-based micro lenses on the top of the OE and the metal reflector at the bottom to increase conversion efficiency**

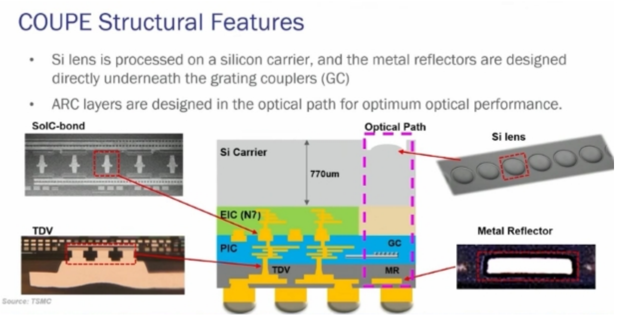

COUPE Structural Features

- Si lens is processed on a silicon carrier, and the metal reflectors are designed directly underneath the grating couplers (GC)
- ARC layers are designed in the optical path for optimum optical performance.

Source: TSMC

EXHIBIT 13: **By aligning FAU precisely with the micro lens on the top of the COUPE structure, optical signals transmitted by fibers can be coupled into PIC and then EIC**

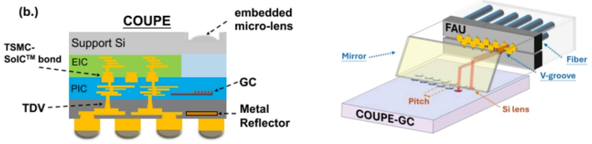

Source: TSMC

EXHIBIT 14: **TSMC roadmap shown in September 2025 revealed the possbility of CPC as an intermediate step before CPO adoption by switch chips**

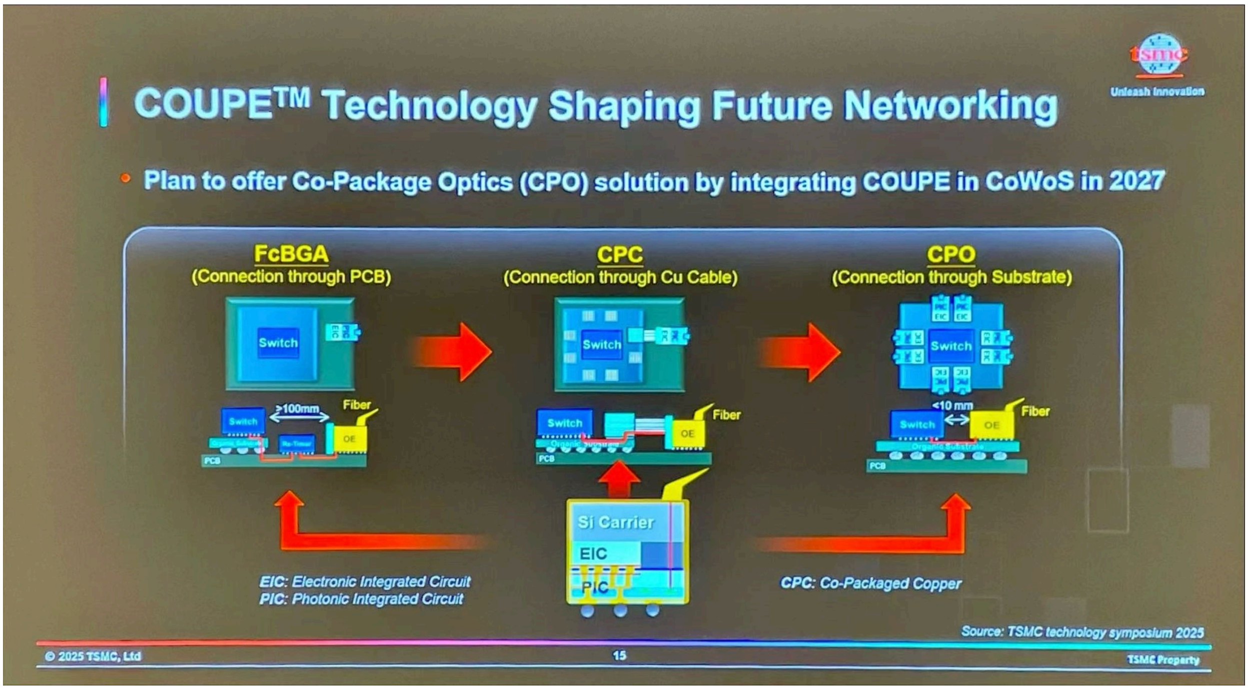

COUPE™ Technology Shaping Future Networking

- Plan to offer Co-Package Optics (CPO) solution by integrating COUPE in CoWoS in 2027

EIC: Electronic Integrated Circuit
PIC: Photonic Integrated Circuit

CPC: Co-Packaged Copper

Source: TSMC technology symposium 2025

© 2025 TSMC, Ltd 15 TSMC Property

Source: TSMC

EXHIBIT 15: **Roadmap provided by TSMC in August 2025 showed CPO will likely be first adopted by switch chip with OE on substrate, and later by XPU with OE on interposer**

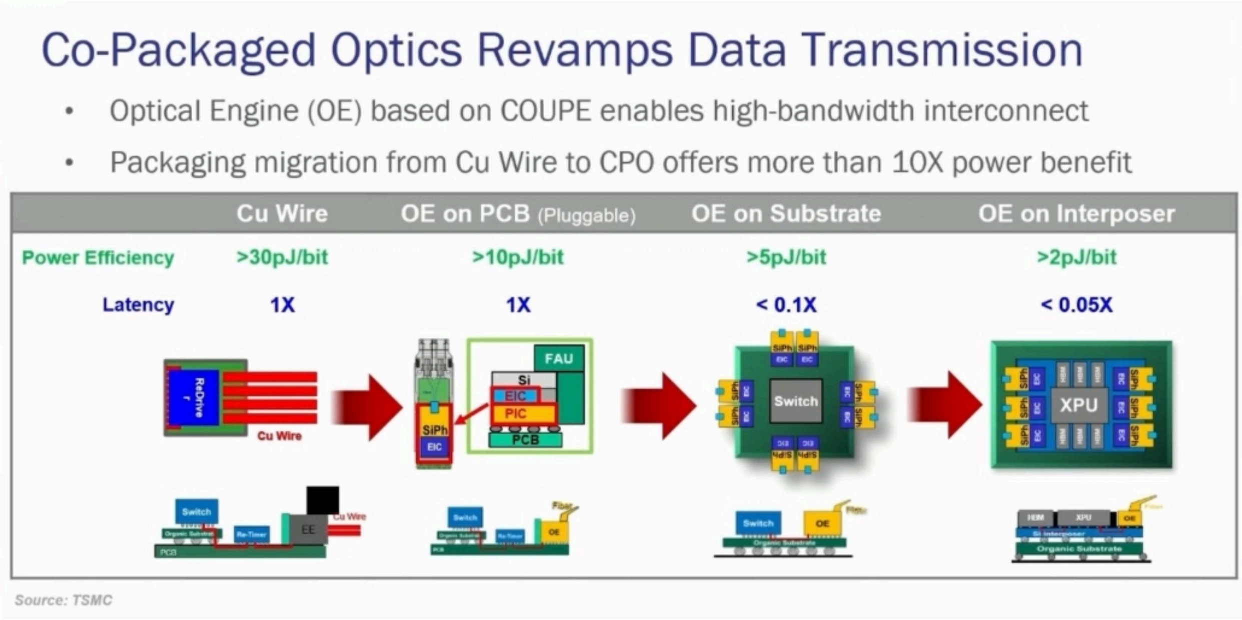

Co-Packaged Optics Revamps Data Transmission

- Optical Engine (OE) based on COUPE enables high-bandwidth interconnect
- Packaging migration from Cu Wire to CPO offers more than 10X power benefit

	Cu Wire	OE on PCB (Pluggable)	OE on Substrate	OE on Interposer
Power Efficiency	>30pJ/bit	>10pJ/bit	>5pJ/bit	>2pJ/bit
Latency	1X	1X	< 0.1X	< 0.05X

Source: TSMC

Source: TSMC

EXHIBIT 16: **TSMC also plans to support higher transmission speeds and higher number of fibers in the future**

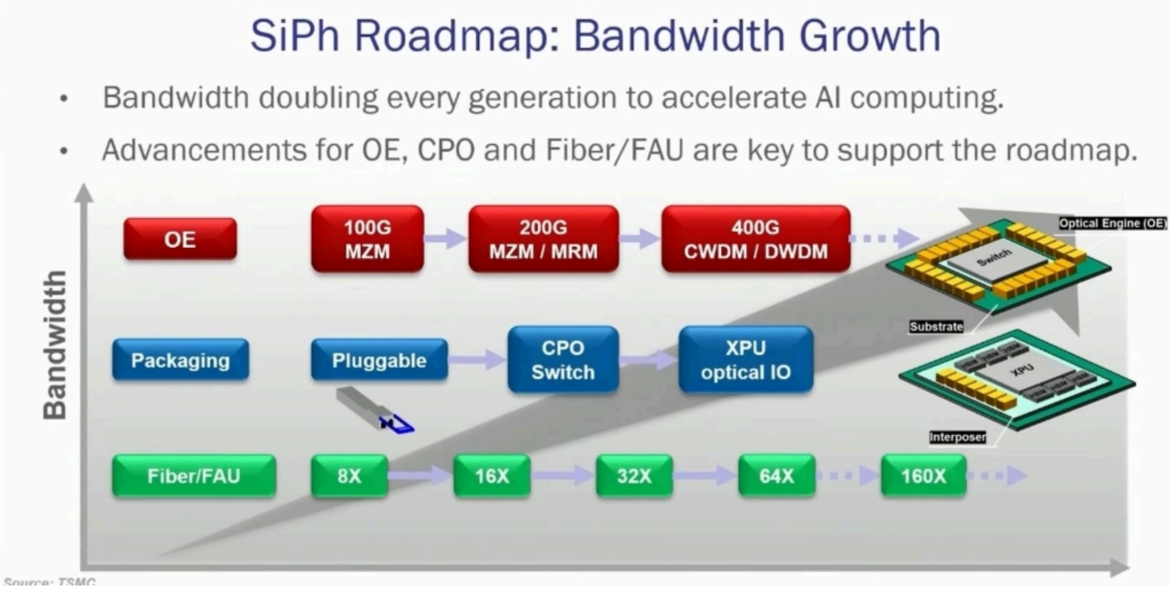

SiPh Roadmap: Bandwidth Growth

- Bandwidth doubling every generation to accelerate AI computing.
- Advancements for OE, CPO and Fiber/FAU are key to support the roadmap.

OE → 100G MZM → 200G MZM / MRM → 400G CWDM / DWDM

Packaging → Pluggable → CPO Switch → XPU optical IO

Fiber/FAU → 8X → 16X → 32X → 64X → 160X

Source: TSMC

WHY MASS ADOPTION OF CPO WILL TAKE TIME

When will we hit the bottleneck in current architecture?

The large-scale adoption of CPO remains uncertain and will likely occur only when existing interconnect technologies come close to their performance limits. For pluggable optical transceivers, 1.6T modules are in mass production in 2026, and the industry expects 3.2T modules to enter volume production in 2028. Certainly there are challenges for the shifts, as it requires a 2x jump in per-lane rate, channel count, power density, and packaging complexity at the physical limits of today's electronics, optics, and cooling.

For copper links in AI data centers, industry roadmaps — such as those from the OIF (Optical Internetworking Forum), IEEE 802.3 Ethernet Task Force, and the Ethernet Technology Consortium (ETC) — increasingly treat **448G** electrical lanes as the practical upper speed node, especially for short-reach, in-rack connections, with meaningful volume expected only after the late-2020s as ecosystems mature. This "ceiling" is driven by the trade-off between data rate and channel reach — as bandwidth increases, the effective distance over which copper can maintain signal integrity declines exponentially.

Challenges of CPO adoption

The adoption of CPO remains encumbered by a number of challenges at the current stage, including technological impediments and customer adoption propensity.

From a technical perspective, major hurdles include: (1) sophisticated packaging and manufacturing workflows, coupled with yield issues and testing difficulties; (2) photonic device manufacturing difficulties, including efficient waveguide-to-fiber coupling and wavelength stability under thermal fluctuations; (3) shortages of key components, such as lasers and indium phosphide (InP) substrates; (4) thermal management challenges arising from compact packaging and elevated power density, e.g., modulators' thermal sensitivity; and (5) the lack of industry-wide standardization frameworks.

Challenges in FAU

An FAU (Exhibit 17) holds optical fibers at a precisely defined pitch to enable efficient coupling with optical waveguides on PICs, making it one of the key CPO components. FAU manufacturing requires extremely tight control over fiber spacing to achieve low insertion loss and high packing density, demanding micron-level alignment accuracy. However, FAU assembly still relies heavily on skilled manual labor, which limits yield and scalability.

After an FAU is ready, it has to be connected to an OE, so that lights transmitted by fibers can be coupled and converted into electrical signals. There are multiple coupling approaches; Exhibit 18 and Exhibit 19 summarize two of them, EC and GC, sometimes called surface or vertical coupling. Each has trade-offs in efficiency, alignment tolerance, and manufacturability. EC supports a wider range of optical wavelengths and has lower coupling loss than GC, but its main drawbacks are: (1) tighter alignment requirement and, hence, complex manufacturing and mechanical and thermal reliability challenges, and (2) lower fiber density (single row of fibers versus multiple rows in GC). Broadcom's earlier CPO products adopted EC, as the company prioritizes performance over other considerations, but we believe NVIDIA favors GC. TSMC's CPO in theory supports both GC and EC, but we believe TSMC prefers GC too and will offer GC first. With that, we believe Broadcom also will likely adopt GC in its newer CPO products.

Shortages of laser and InP substrates

In today's datacom pluggable market, Vertical-Cavity Surface-Emitting Lasers (VCSELs) are widely used for short-reach multi-mode 400G and 800G links (≤100 meters), while Electro-Absorption Modulated Lasers (EMLs) and silicon photonics with CW lasers are the main light-source technologies for longer-reach single-mode 400G and 800G, and evolving 1.6T modules. An EML integrates a Distributed Feedback (DFP)-CW laser and Electro-Absorption Modulators (EAMs), so the light is generated and modulated inside the optical transceiver modules (Exhibit 20 to Exhibit 22).

In CPO, by contrast, modulation moves onto the PIC at the chip level (e.g., via micro-ring modulators (MRMs), and light generation is provided by an external CW laser. This lets the high-power CW source sit away from hot CPUs and ASICs, improves thermal management, and allows failed lasers to be replaced without touching the OE package. A single high-power CW laser can feed many PIC channels through splitters, reducing total system power. However, CPO requires ultra-high-power CW diodes, with each device delivering hundreds of milliwatts (mW), which in turn demands specialized epitaxy and careful thermal design.

InP substrates are critical for both CW and EML lasers. The market is currently tight, reflecting surging demand from optical transceivers and China's stricter controls on indium compounds. Leading InP substrate suppliers such as **Sumitomo Electric Industries** (5802 JT) and **AXT Inc** (AXTI US, both not covered) are expanding capacity. Meanwhile, laser makers including **Lumentum** (Lite US) and **Coherent** (COHR US, both not covered) are adding InP laser chip capacity. Coherent has roughly tripled its InP capacity over FY2024-FY25, and its new 6-inch line is now ramping, which should materially lift internal output in 2026. Moving from 4-inch to 6-inch wafers can roughly quadruple die output per wafer and cut unit cost by about half, according to Coherent management. Lumentum announced in late-2024 that it planned to increase its capacity by >40%, and management commented in February 2026 earnings that it is more than halfway through that capacity expansion (Exhibit 23 and Exhibit 24).

On March 2, 2026, NVIDIA announced $2Bn in strategic investments each in Lumentum and Coherent, together with multi-year, non-exclusive partnerships that include large purchase commitments and future access or capacity rights for laser and optical networking products.[6, 7]

Bottlenecks in modulators

Modulators convert electrical data into optical signals. The three mainstream technologies are Mach–Zehnder modulators (MZMs), MRMs, and electro-absorption modulators (EAMs), with their pros and cons summarized in Exhibit 25. EAM is an InP-based intensity modulator commonly used in EML optical transceivers, while MZMs are widely adopted in silicon photonics (SiPh) pluggable transceivers.

In CPO, **MRMs** have drawn strong interest because of their very small footprint, low power consumption, and high bandwidth density, which is why NVIDIA has chosen an MRM-based approach. MRMs sit on the PIC die and imprint the XPU or ASIC electrical data onto an optical carrier from an external CW laser; their resonant behavior also makes them naturally well-suited for dense wavelength-division multiplexing (WDM), with each ring selecting and modulating a specific wavelength in a minimal area. The trade-off is that MRMs are highly sensitive to temperature and fabrication variation, so NVIDIA is working with **TSMC** to address manufacturing tolerances and to implement robust control schemes (Exhibit 26).

EXHIBIT 17: **Structure of fiber array units**

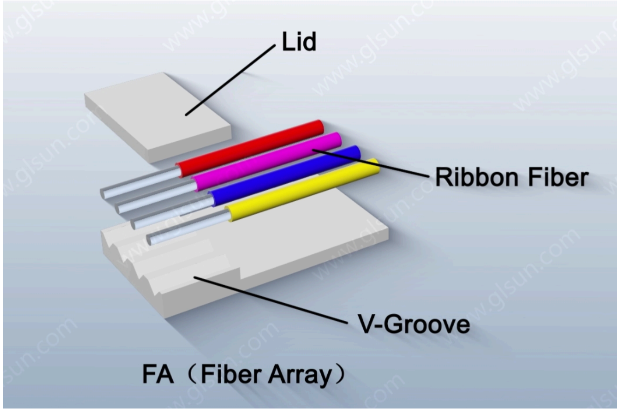

Source: GLsun Science, Tech Group

EXHIBIT 18: **Edge coupling (EC) and grating coupling (GC) are two common approaches to couple lights from fibers into OEs**

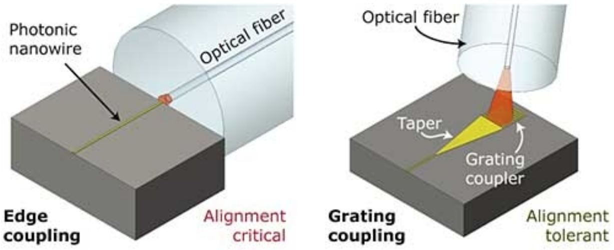

Source: KTH Royal Institute of Technology

---

[6] NVIDIA and Coherent Announce Strategic Partnership to Develop Optics Technology to Scale Next-Generation Data Center Architecture | NVIDIA Newsroom.

[7] NVIDIA Announces Strategic Partnership with Lumentum to Develop State-of-the-Art Optics Technology | NVIDIA Newsroom.

EXHIBIT 19: **Each of two mainstream coupling methods face multiple unresolved technical challenges**

	Edge coupling	Grating coupling
Fiber density	Only allow one row of fibers, limited density	Allow multiple row of fibers, more compact
Manufacturing difficulty	Complex, require additional cleaving and polishing process to create the coupling facet	Simple, and compatible with wafer-level testing
Coupling loss	Low	Relatively high
Operating bandwidth	Broad	Relatively narrow
Polarization sensitivity	Low	Relatively high
Representative players	GlobalFoundries; Broadcom (earlier CPO products)	Nvidia, TSMC

Source: Ansys, Tsinghua University, Bernstein analysis

EXHIBIT 20: **Schematic view of an EML laser using butt-joint technology**

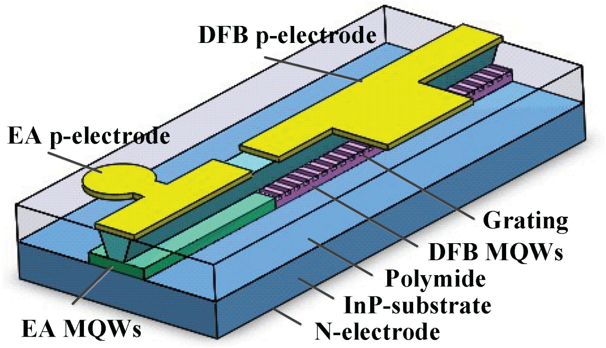

Source: Chinese Academy of Sciences, University of Chinese Academy of Sciences, Beijing Key Laboratory of Low Dimensional Semiconductor Materials and Devices

EXHIBIT 21: **EML lasers (DFB-based) use a different substrate and emission mechanism compared with VCSEL lasers**

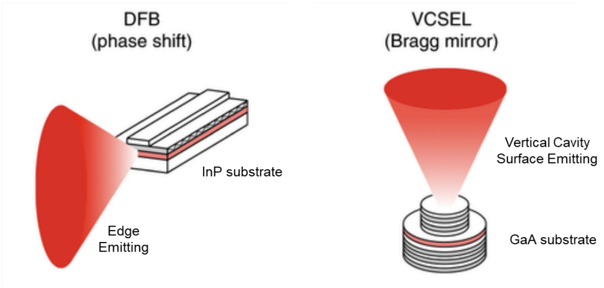

Source: Chinese Academy of Sciences, University of the Chinese Academy of Sciences, Songshan Lake Materials Laboratory

EXHIBIT 22: **100G and 200G/lane lasers for 800G and 1.6T**

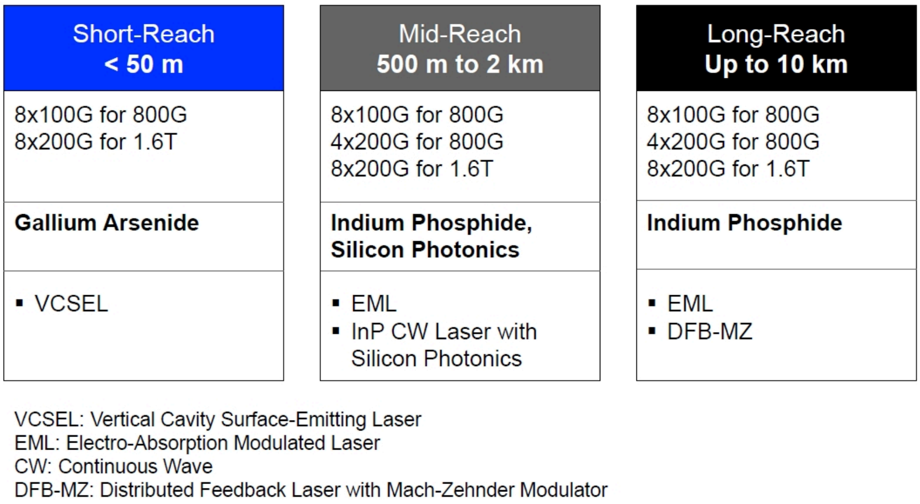

Short-Reach **< 50 m**	Mid-Reach **500 m to 2 km**	Long-Reach **Up to 10 km**
8x100G for 800G 8x200G for 1.6T	8x100G for 800G 4x200G for 800G 8x200G for 1.6T	8x100G for 800G 4x200G for 800G 8x200G for 1.6T
Gallium Arsenide	**Indium Phosphide, Silicon Photonics**	**Indium Phosphide**
▪ VCSEL	▪ EML ▪ InP CW Laser with Silicon Photonics	▪ EML ▪ DFB-MZ

VCSEL: Vertical Cavity Surface-Emitting Laser
EML: Electro-Absorption Modulated Laser
CW: Continuous Wave
DFB-MZ: Distributed Feedback Laser with Mach-Zehnder Modulator

Source: Coherent

EXHIBIT 23: **Sumitomo Electric targets to more than double its InP substrate capacity from 2023 to 2028**

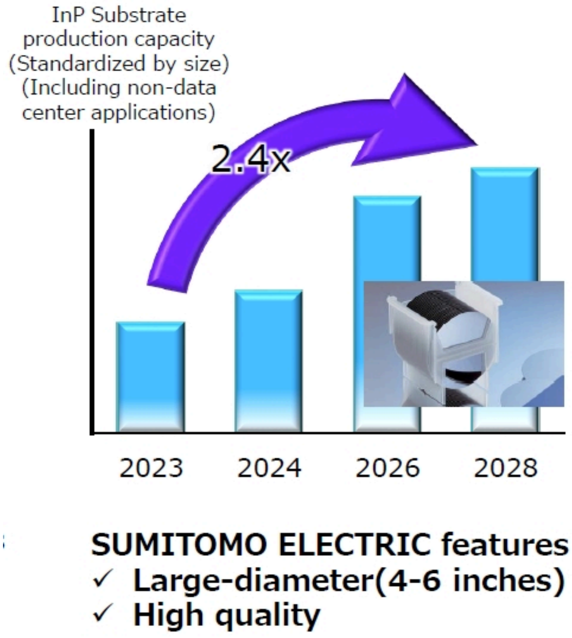

Source: Sumitomo Electric

EXHIBIT 24: **Coherent has tripled its InP capacity in FY2024-FY25, and will add more in 2026**

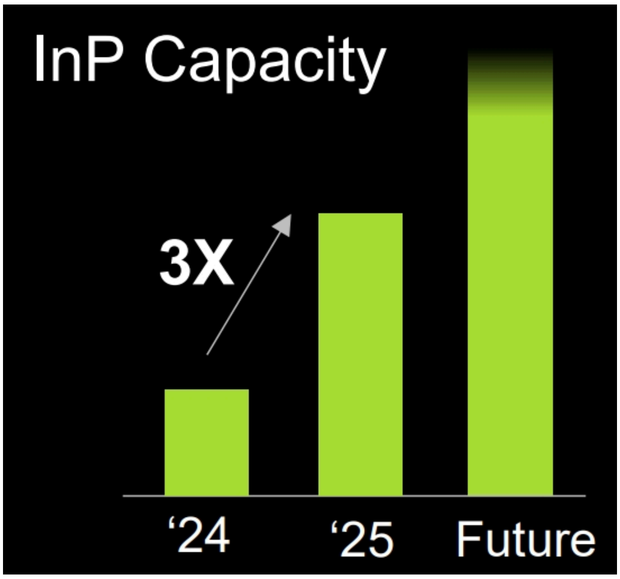

Source: Coherent

EXHIBIT 25: **Key differences of three major modulators**

Modulators	MZM	EAM	MRM
Size (μm^2)	~12,000	~250	100-200
Thermal stability	Yes (>50°C)	Yes (>50°C)	No (<1°C)
Power consumption	High	Medium	Low
Optical bandwidth	Optically broadband	30nm	Narrow band
WDM suitability	External Multiplexer/Demultiplexer needed	Medium Requires per-λ source + Multiplexer	Native WDM Each ring = 1 λ channel High density cascading

Source: Celestial AI, Bernstein analysis

EXHIBIT 26: **NVIDIA partners with TSMC to tackle manufacturing and control challenges for MRM**

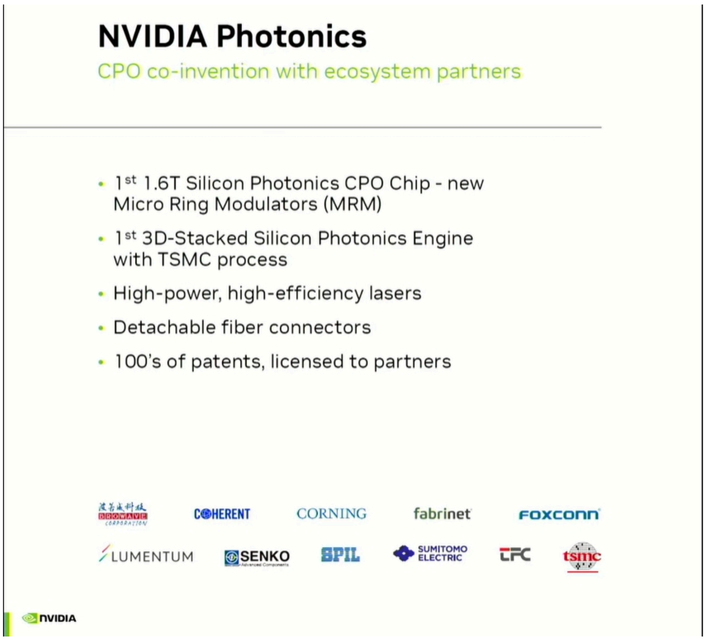

NVIDIA Photonics

CPO co-invention with ecosystem partners

- 1st 1.6T Silicon Photonics CPO Chip - new Micro Ring Modulators (MRM)
- 1st 3D-Stacked Silicon Photonics Engine with TSMC process
- High-power, high-efficiency lasers
- Detachable fiber connectors
- 100's of patents, licensed to partners

Note: TSMC and NVIDIA are covered, all other companies mentioned are not covered.
Source: NVIDIA

CPO SHIPMENT IN 2H26-1H28 WILL BE FOR SCALE-OUT ONLY

Despite the challenges discussed earlier in this chapter, CPO developers are positive about mass adoption. (1) NVIDIA's roadmap outlined shipments of its Quantum-X InfiniBand CPO switch in 2H25 and deliveries of its Spectrum-X Ethernet CPO switches in 2H26. In its March 2026 CPO webinar, NVIDIA further noted that deployments will begin in 2026, with CoreWeave (covered), Lambda (private), and the Texas Advanced Computing Center among the first adopters. (2) Broadcom announced the shipment of its Tomahawk 6 switch series in June 2025, while Alpha Networks (3380 TT, not covered) launched1.6T liquid-cooled switches equipped with the Broadcom TH6 chip in October 2025.

Supply chain companies, however, express mixed views. (1) Laser developer Lumentum (Lite US, not covered) confirmed receiving several hundred million dollars of additional laser orders for CPO scale-out applications, with shipments scheduled for 2027. Management expects its first scale-up CPO shipments by late 2027. (2) FAU provider FOCI (3363 TT, not covered) indicated that mass shipments of CPO switches may begin in 2027-28 (Link[8]). (3) Innolight (optical transceiver provider, 300308 CH, not covered) stated that none of its CSP customers plan to deploy CPO switches for scale-out network in a large scale in 2026-27.

We believe CSPs will continue favoring pluggable optical transceivers over CPO in 2026-28 because: (1) In CPO architectures, photonic components are embedded inside the switch. Any failure typically requires replacing the entire switch or returning it for factory repair — resulting in longer downtime and significantly higher operational impact. In contrast, pluggable modules can be swapped on-site by standard IT staff with minimal service disruption. While the September 2025 test results from Meta indicate the BaillyCPO annual link failure rate (ALFR) is 5x better than the pluggable solutions, the testing period is still limited and is in a lab

---

[8] https://finance.technews.tw/2025/11/24/foci-3363/

environment. (2) Compared to the optical transceiver market, we expect the CPO market to be highly concentrated, which diminishes CSPs' bargaining power.

We believe that CPO shipment in 2H26 will be in small volumes for scale-out, primarily aimed at gathering market feedback and verifying supply chain readiness. The deployment in scale-up is likely to be in 2H28, as the industry needs more time to test CPO reliability under real-world conditions before integrating it into high-value XPUs. LightCounting projects the large-scale shipment of CPOs to commence in 2028 (Exhibit 27).

EXHIBIT 27: **LightCounting projects the large-scale shipment of CPOs to commence in 2028**

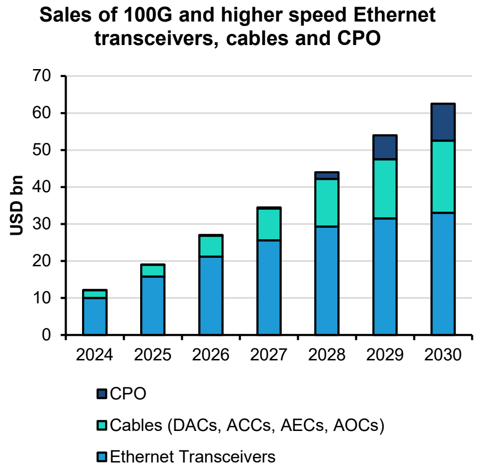

Source: LightCounting data and estimates (2026-30), Bernstein analysis

BRIDGING THE GAP: DESIGNING FOR THE PRE-CPO TRANSITION

As CPO is unlikely to see mass adoption in the near term, switch and optics vendors are rolling out transitional solutions, with LPOs and NPOs emerging as the most practical.

LPO is an evolution of traditional pluggable modules that removes DSPs and relies mainly on linear components, such as transimpedance amplifiers (TIAs) and drivers to handle signal loss. This reduces power, latency, and BOM cost while reusing the existing pluggable ecosystem, so 800G LPO is already in commercial production. However, because the optics still sit at the faceplate, farther from the XPU/ASIC than NPO or CPO, electrical traces are longer, which increases loss and link power.

NPO is architecturally closer to CPO; the OE is mounted on a separate substrate/PCB and plugs into a socket near the ASIC, making it detachable for easier maintenance. **NVIDIA's current "CPO"-branded solutions effectively follow this NPO model**, as its OEs are removable rather than permanently co-packaged.

Where do traditional optical transceiver companies fit in different products? Traditional optical transceiver companies fit naturally into LPO, as it builds on the same module form factors, optics, and supply chain. Vendors such as Innolight and Eoptolink (300502 CH, both not covered) have already shipped 800G LPO products. By contrast, their **NPO** offerings remain in development, because NPO demands semiconductor-class packaging and close integration with ASICs, which differ significantly from conventional pluggable-module assembly. Innolight has been involved in NPO R&D (focus on packaging) and aims to finalize product design in 2026.

CPO is still in the early stage of development (e.g., ELS and FAU related products) (Exhibit 28 to Exhibit 31).

EXHIBIT 28: **Difference in designs of DSP Pluggable, LPO, NPO, and CPO**

DSP Pluggable

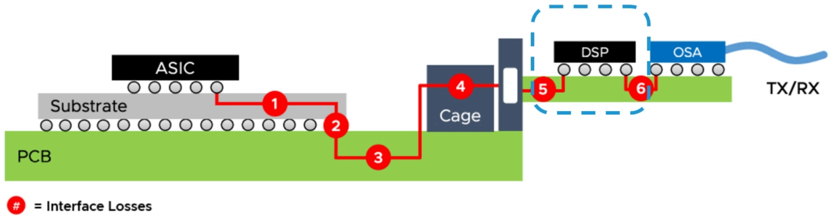

LPO/LRO Pluggable

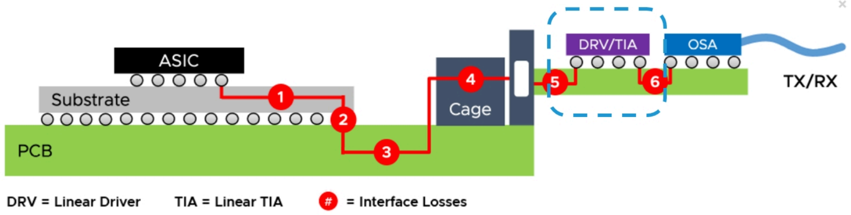

NPO

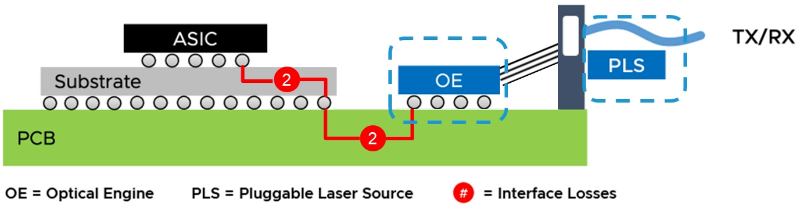

CPO

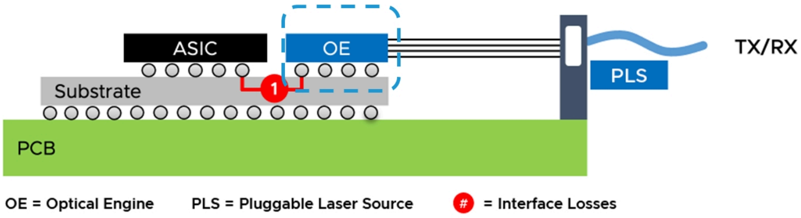

Source: Broadcom

EXHIBIT 29: **Comparison of DSP Pluggable, LPO, NPO, and CPO**

	DSP Pluggable	LPO	NPO	CPO
Design	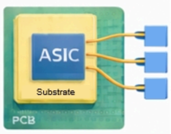	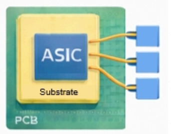	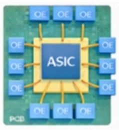	
Serviceability	Hot swappable	Hot swappable	Packaged but detachable	Packaged
Optical module location	Attached to Switch PCB	Attached to Switch PCB	Switch PCB	Substrate
Signal processing	DSP	TIA and drivers, instead of DSP	TIA and drivers, instead of DSP	TIA and drivers, instead of DSP
Laser type	Internal	Internal	Mostly external	Mostly external
SiPh integration	Not necessarily SiPh	Not necessarily SiPh	SiPh integration	SiPh integration
Energy efficiency	c. 20-22 pJ/bit	c. 16-18 pJ/bit	c. 7-8 pJ/bit	c. 5 pJ/bit
Reliability	Low	Low	High	The highest

Source: Broadcom, Bernstein analysis

EXHIBIT 30: **1.6T LPO and NPO achieve lower power consumption than retimed optics and offer longer reach than copper-based solutions**

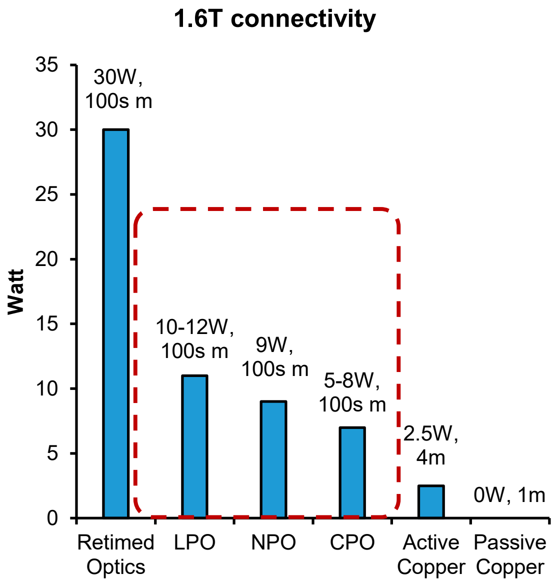

Source: Nubis Communication, Bernstein analysis

EXHIBIT 31: **LightCounting projects more LPO shipment than CPO by 2030**

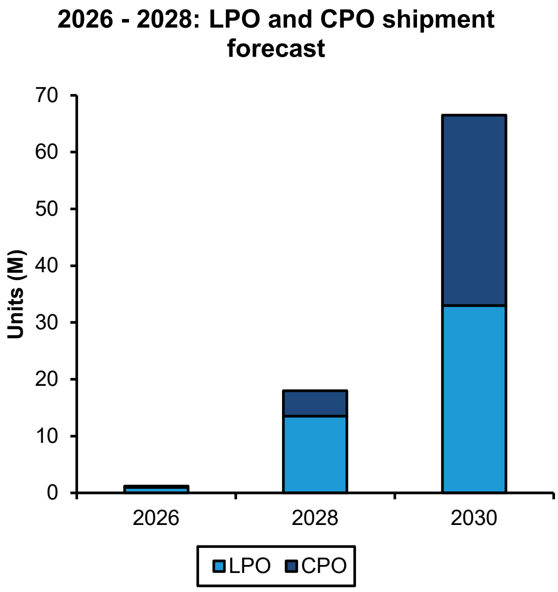

Source: LightCounting data and estimates, Bernstein analysis

WHAT ABOUT COPPER CONNECTIVITY IN AI DATA CENTERS?

Optical versus copper connectivity in data centers: Copper interconnection has always been valued for its low power consumption, strong cost-effectiveness, and high reliability. As bandwidth has scaled, however, the practical reach of copper has shrunk, eroding its share in high-speed data center links over the past decades (Exhibit 32). At 224Gbps bandwidth, passive direct-attach cables (DACs) are typically limited to about 2 meters, while active electrical cables (AECs) can extend copper reach to roughly 7 meters, keeping copper mainly in very short-reach and in-rack applications. In contrast, optical links provide extended reach (≥10 meters) with low signal attenuation and terabit-scale scalability, making optics the de facto choice for data-center interconnect (DCI), intra-datacenter, and rack-to-rack connectivity.

CPC technology is likely to extend copper's lifecycle in AIDC. As bandwidth continues to rise and CPO gains attention as a scale-up solution, there has been a recurring debate over whether copper will be phased out. Compared to traditional copper solutions, CPC integrates high-speed copper cable connectors directly into ASIC or switch packages and bypasses the PCB traces. It avoids much of the board-level loss and makes even 448 Gbps-class copper links technically feasible. Compared with CPO, CPC preserves compatibility with existing copper or connector ecosystems. It offers a meaningful cost advantage and faces fewer integration risks, potentially giving copper a longer runway in next-generation AI infrastructure (Exhibit 33).

Luxshare (Outperform) is a new supplier for GB300 copper connectivity, competing with Amphenol (APH US, not covered). Luxshare is likely involved in the Vera Rubin platform (see China Next Winners: Tech Hardware - Luxshare's AI play and China's dominance in optical transceiver & HDI).

EXHIBIT 32: **Optics vs. copper connectivity application trend**

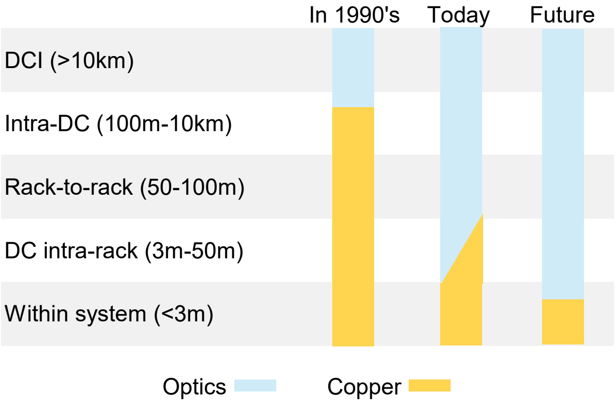

Source: Yole, Bernstein analysis

EXHIBIT 33: **LightCounting projects that copper cables will maintain around half of the 1.6T market share in 2029**

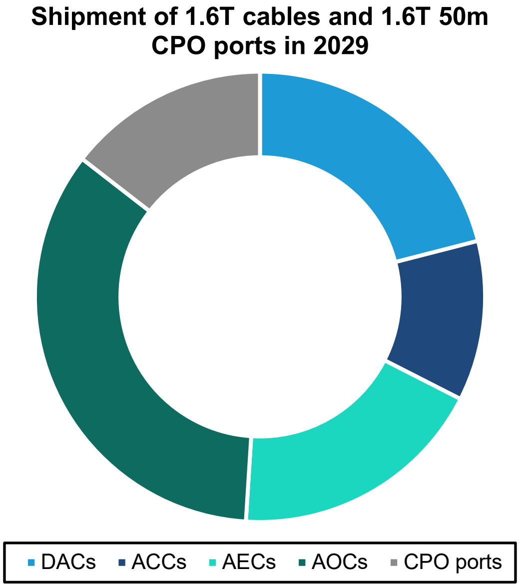

Source: LightCounting estimates, Bernstein analysis

MAPPING THE CPO VALUE CHAIN

INDUSTRY OVERVIEW

We model CPO solution costs at least 10% higher than optical transceivers. By integrating OEs with XPUs or switch chips, CPO offers superior power efficiency, signal integrity, and network resiliency compared to pluggable optical transceivers. We estimate that OE accounts for ~45% of the CPO's BOM. While the BOM of CPO OE + ELS is lower than that of 1.6T transceivers, once we include the margin captured by switch vendors (e.g., NVIDIA), the total upfront cost is at least 10% higher for customers.

We believe TSMC's COUPE will become the mainstream CPO platform. Although GlobalFoundries entered silicon photonics earlier, key customers such as Broadcom and NVIDIA are shifting to TSMC's COUPE, which bonds electronic and photonic IC with hybrid bonding, and attaches micro lens and FAU precisely on the top. This allows EIC and PIC to optimize their node selections separately and, more importantly, working with TSMC allows customers to integrate COUPE with CoWoS so that the adoption can advance from NPO to CPO, starting gradually in 2026. We believe the direct revenue contribution of CPO and silicon photonics to TSMC is small, but they will strengthen TSMC's front-end wafer and CoWoS business.

CPO manufacturing and testing are challenging, favoring suppliers with technology leadership. NVIDIA's Quantum switch chip integrates18 OEs/chip, rising to over 30 OEs in Spectrum, and each FAU carries tens of fibers with extremely tight fiber spacing to achieve low insertion loss and high packing density. We expect **TFC** and **Senko** (both not covered) to supply FAUs in 2026, with **FOCI** (not covered) likely catching up as OE moves from PCB to substrate. CW lasers in CPO need to deliver hundreds of mW with strong thermal stability, which makes **Lumentum** (not covered) stand out. Finally, OE-FAU coupling, and testing at wafer, die, and package levels create new opportunities for equipment vendors including **Chroma ATE (Outperform), All Ring, Teredyne,** and **ficonTec** (all three not covered).

CPO startups are also playing an important role in pushing the ecosystem forward with a focus on scale-up solutions. They target more disruptive architectures, ranging from in-package optical interconnect (e.g., Ayar Labs (private), Celestial AI (acquired by Marvell, not covered), and Lightmatter (acquired by Cactus, not covered)) to integrated CPO platform (e.g., Nubis, acquired by Ciena (not covered)). Their technologies could enable higher bandwidth density, lower latency, and tighter XPU integration over the longer term, but most remain in the early pilot stage with limited production visibility.

What about Google's Optical Circuit Switch (OCS)? CPO is not the only path hyperscalers are exploring to solve the bandwidth and power challenges. Since TPU v4, Google has adopted OCS, which uses Micro-Electro-Mechanical Systems (MEMS) mirrors to steer light from input to output, avoiding the optics-to-electronics conversions of traditional switches. As more CSPs show interest in merchant OCS solutions from vendors such as Lumentum and Coherent, broader adoption still faces hurdles, including the need for full stack redesign (software and tooling) and a relatively immature ecosystem.

CPO CONTENT ANALYSIS – WE ESTIMATE THAT CPO IS 10% MORE EXPENSIVE THAN PLUGGABLE SOLUTIONS

At GTC 2025, NVIDIA exhibited three distinct CPO-based switches, including Quantum X800-Q3450, Spectrum 6810, and Spectrum 6800. The Quantum CPO switch[1] features four ASICs, with each ASIC surrounded by 18 OEs (three OEs per sub-assembly and six sub-assemblies). The switch is powered by 18 external light sources (ELS), each with eight CW lasers. We estimate overall CPO switch costs of ~$70k (excluding profit of switch solution). **Compared to the ASP of 1.6T optical transceiver, the OE and laser in CPO ASP is at least 10% more expensive.** This includes gross profit of optical transceiver suppliers (assuming 40% GM), and assuming 50% GM (or higher) at CPO switch vendor (NVIDIA). The upcoming Spectrum 6810 will upgrade bandwidth/OE to 3.2Tb/s and the total bandwidth is similar to Quantum. A more powerful switch — Spectrum 6800 — is also available, housing four Spectrum-6 ASICs and reaching a total bandwidth of 409.6 TB/s[2] (Exhibit 1 to Exhibit 8).

EXHIBIT 1: **NVIDIA's CPO-based switches**

	Nvidia	
	Quantum X800-Q3450	**Spectrum 6810**
Pilot Production date	2H 2025	2H 2026
Networking standard	InfiniBand	Ethernet
Switch ASIC	Quantum-3	Spectrum-6
Throughput per package (Tb/s)	28.8	102.4
Number of switch package	4	1
Switch Aggregate Bandwidth (Tb/s)	115.2	102.4
SerDes speed (Gb/s unidi)	200G	200G
Optical connectivity	DR optics	DR optics
Bandwidth per optical engine (OE; (Tb/s)	1.6	3.2
Number of OEs	72	36
External Light Sources (ELS)	18	16

Source: NVIDIA, Bernstein analysis

EXHIBIT 2: **OE makes up ~45% of a Quantum X800 switch**

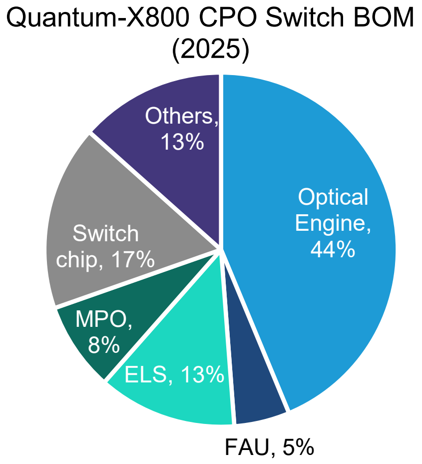

Source: SemiAnalysis, Bernstein analysis and estimates (all data)

---

[1] How Industry Collaboration Fosters NVIDIA Co-Packaged Optics | NVIDIA Technical Blog.

[2] Scaling Power-Efficient AI Factories with NVIDIA Spectrum-X Ethernet Photonics | NVIDIA Technical Blog.

EXHIBIT 3: **Compared to the ASP of 1.6T optical transceivers, OE and laser in CPO ASP is 10% higher**

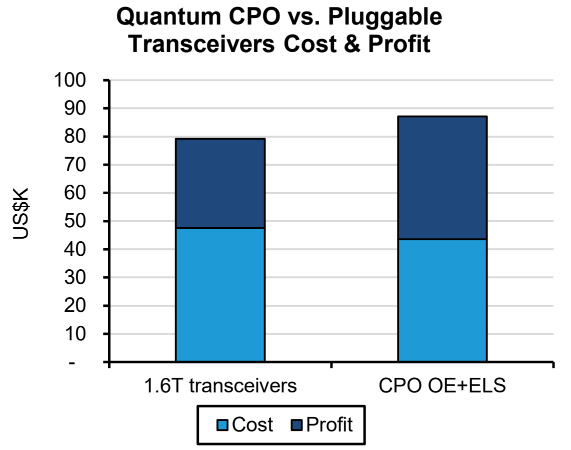

Source: Bernstein analysis and estimates (all data)

EXHIBIT 4: **The NVIDIA Quantum X CPO switch is equipped with six optical sub-assemblies (OSA), and each OSA houses three OEs**

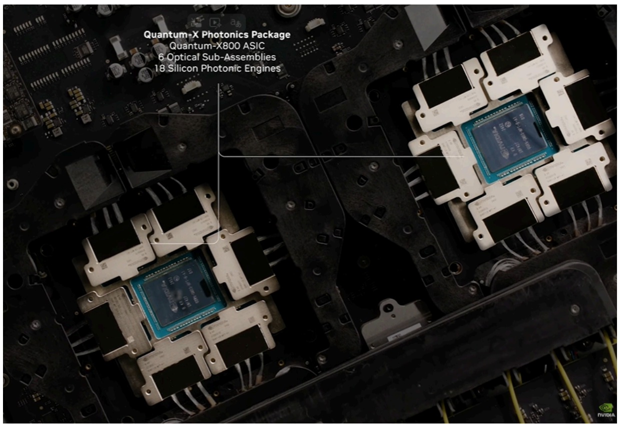

Source: NVIDIA

EXHIBIT 5: **Each switch chip includes 36 laser inputs from 18 ELS units, along with 288 data links**

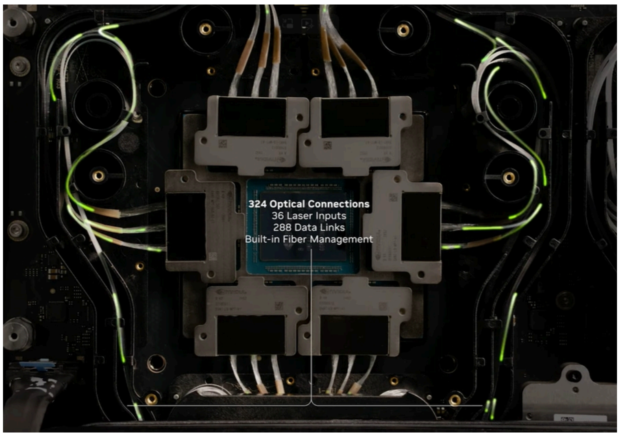

Source: NVIDIA

EXHIBIT 6: **Each OE is connected to 18 fibers through an FAU**

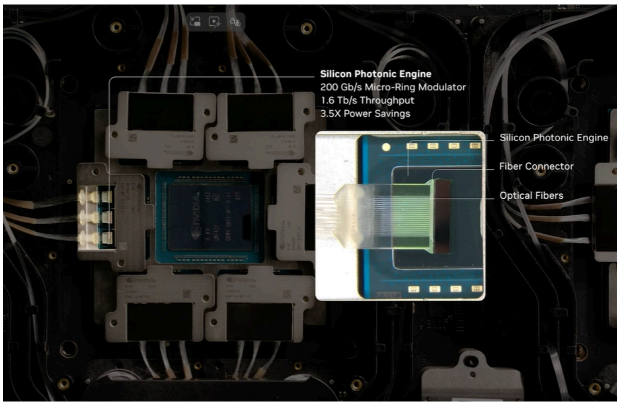

Source: NVIDIA

EXHIBIT 7: **Within each OE, the EIC and PIC are 3D-stacked**

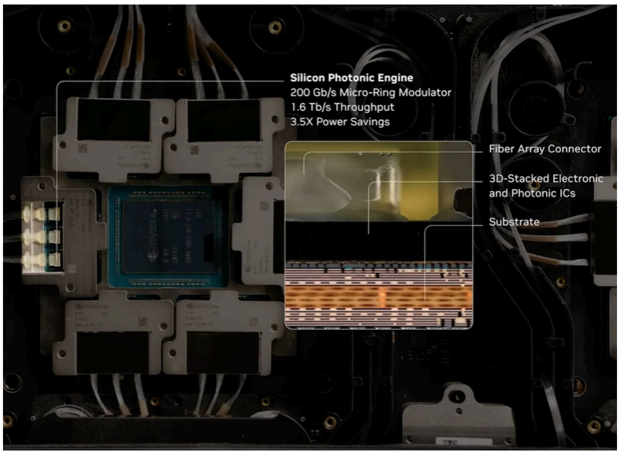

Source: NVIDIA

EXHIBIT 8: **Spectrum ASIC features over 30 OEs**

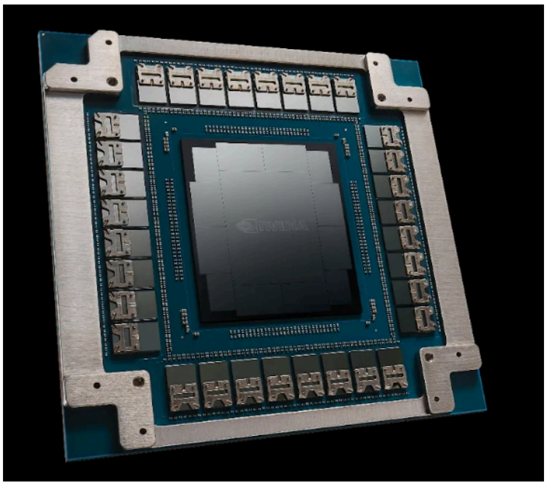

Source: NVIDIA

FOUNDRY FOR CPO PLATFORMS

TSMC (Outperform): Being the leader in XPU manufacturing, TSMC is also taking a leading role in working with many customers and partners to create a CPO solution and integrate the solution with its front-end wafer and advanced packaging services. The most notable offering in TSMC's CPO platform is its COUPE. Detailed production steps of COUPE can be found in the chapter titled "Co-Packaged Optics: Inside the War for Data Center Connectivity" (Exhibit 9).

Though TSMC is leading many companies to jointly turn CPO from concept to commercialization, we believe for TSMC the direct revenue contribution from CPO will be limited in the near future. CPO, however, will likely be a complementary technology that strengthens TSMC's position in front-end wafer and advanced packaging services, and our TSMC model (available on request) reflects this.

GlobalFoundries (GFS US, not covered): GlobalFoundries brands its silicon photonic offering as "Fotonix." The company started work on silicon photonic technology before TSMC, originally through the acquisition of IBM Microelectronics in 2015, followed by the development of 90nm and in 4Q21 by the production of 45nm. Similar to TSMC, GlobalFoundries offers its silicon photonic technologies with front-end production on 300mm wafers and in integration with packaging, but GlobalFoundries differentiates itself from TSMC with its ability to manufacture EIC and PIC on a single monolithic die.

The earlier start than TSMC attracted many customers, e.g., Broadcom, Cisco, Marvell, NVIDIA, Ayar Labs, Lightmatter, PhiQuantum, Ranovus, and Xanadu, to work with GlobalFoundries, but we believe many of these customers, notably NVIDIA, Broadcom, and Ayar Labs, have shifted their recent products to TSMC. In terms of ecosystems, GlobalFoundries has partners including Corning and Senko for FAU and fiber; Amkor and ASE on packaging (e.g., for Lightmatter); and Cadence and Synposys for Electronic Design Automation (EDA). In late 2025, the company also announced the acquisition of a small silicon photonics foundry Advanced Micron Foundry (AMF), which has a 200mm-fab in Singapore, and of InfiniLink, a Cairo-based startup for high-speed design capabilities for silicon photonics and transceivers. We note that MediaTek invested in InfiniLink in early 2025.

GlobalFoundries guided for the acquisition of AMF to bring over $75Mn in revenue in 2026. Though the company also expects its silicon photonic revenue to exceed $1Bn by 2030, we note that while older products may stay at GlobalFoundries, many of its customers are shifting their new products to TSMC. We attribute this shift mainly to the transition from NPO/OBO to CPO necessitating CoWoS, in which TSMC commands a clear leadership. We also note that NVIDIA picking 6nm for its EIC suggests a leading-edge logic node may become necessary for EIC. A slide (Exhibit 10) from TSMC also highlights that the node disparity between EIC (on leading node) and PIC (on mature node) will favor a heterogenous approach (making them separately on their respective optimal nodes and bonding them together afterward) over GlobalFoundries' monolithic approach. In light of these considerations, though GlobalFoundries is competitive now, its technology offering not going below 12nm may be a disadvantage going forward.

Tower Semiconductor (TSEM US/TSEM IT, not covered): Tower also had an early start in silicon photonics. In early 2024, the company claimed to have over 50 active silicon photonic customers, including Innolight, Coherent, and Marvell. In February 2026, Tower announced a plan to team up with NVIDIA for the PIC used in "1.6T data center optical modules design for NVIDIA networking protocols." Though Tower announced the availability of silicon photonic process and wafer bonding technology on 300mm wafer, we believe Tower's primary strength and customer traction is on 200mm, with its 0.18um process. The lack of advanced logic nodes and advanced packaging equivalent to CoWoS will likely keep Tower's traction mostly in OE, especially in PIC, but not much in CPO.

Tower registered $228Mn silicon photonics revenue in 2025, more than double the $106Mn in 2024. In 4Q25 alone, revenue was $95Mn. The company plans to expand silicon photonics capacity over 5x versus its 4Q25 shipment, and claimed more than 70% of the new capacity already reserved or being reserved through 2028.

Intel has silicon photonic and CPO technologies (webpage) too, but we find little customer traction, either through its own products or through foundry service.

Overall, GlobalFoundries and Tower have a head-start in silicon photonics, but TSMC leads in advanced logic and packaging, and the necessary integration with them will likely make TSMC a strong supplier in EIC and CPO (Exhibit 11). TSMC is strong in manufacturing electrical devices with silicon, but optical devices are often made with other materials such as silicon on insulator (SOI) and silicon germanium (SiGe). GlobalFoundries and Tower are more experienced in these materials and their earlier start also positions them well to benefit from the rise of silicon photonics, particularly through the existing products that customers have engaged with them. Going forward, it appears EIC will advance to finer nodes and that should make TSMC improve its position in EIC, and possibly in PIC too as TSMC leverages the interface between EIC and PIC. More importantly, CPO needs to employ advanced packaging CoWoS to co-package XPU/switch chip and EOs together. Hence, we expect TSMC to become a major supplier of CPO in the next few years. That said, we also expect this contribution is relatively small versus TSMC's existing revenue size and, hence, choose to model it implicitly through better wafer and advanced packaging outlook.

EXHIBIT 9: **Compared with micro-bump electrical interconnects, TSMC's COUPE lowers power consumption by 40% at the same speed, or reaches 170% speed gain with the same power**

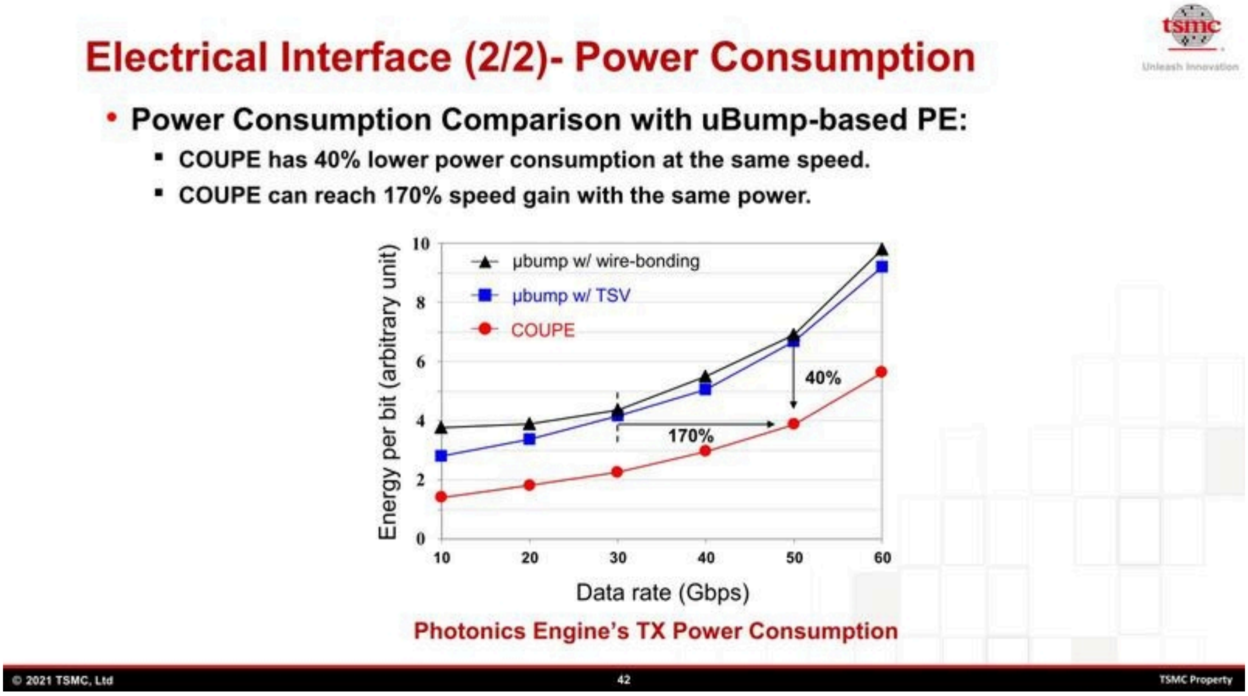

Electrical Interface (2/2)- Power Consumption

- **Power Consumption Comparison with uBump-based PE:**
  - COUPE has 40% lower power consumption at the same speed.
  - COUPE can reach 170% speed gain with the same power.

Photonics Engine's TX Power Consumption

Source: TSMC

EXHIBIT 10: **TSMC highlighted that node disparity between EIC and PIC favors heterogenous integration, i.e., making EIC and PIC separately on their optimal nodes, respectively, and bonding them together afterward**

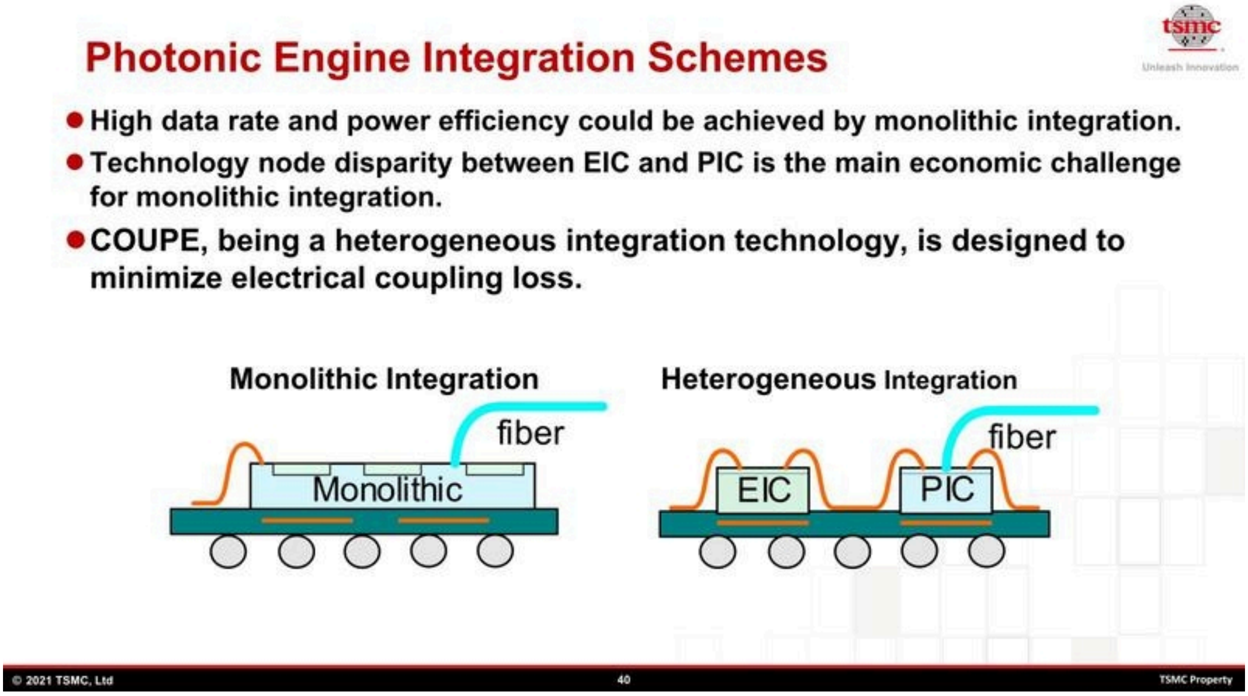

Photonic Engine Integration Schemes

- High data rate and power efficiency could be achieved by monolithic integration.
- Technology node disparity between EIC and PIC is the main economic challenge for monolithic integration.
- COUPE, being a heterogeneous integration technology, is designed to minimize electrical coupling loss.

Source: TSMC

EXHIBIT 11: **Comparison of TSMC, GlobalFoundries, and Tower in CPO and silicon photonics**

	Strength	Weakness
TSMC	> Adv. packaging tech necessary to achieve CPO > Finer logic node for better EIC > Many customer engagements, esp. with leading ones	> Late start in silicon photonics
GlobalFoundries	> 45nm 12" process to make EIC & PIC monolithically > Many customers on existing products	> No CoWoS or fine logic node
Tower	> Mostly 0.18um 8" process > A major 1.6T supplier for NVIDIA networking protocols	> No CoWoS or fine logic node > Limited 12" capabilities

Source: Company reports, Bernstein analysis

SUPPLY CHAIN IMPACT – NEW OPPORTUNITIES IN LASERS, FAUS, AND TESTERS

Compared with the pluggable transceiver ecosystem, CPO's highly integrated architecture shifts a larger portion of value creation toward semiconductor manufacturers. As a result, semiconductor suppliers and key optical component or equipment vendors are likely positioned as the primary beneficiaries of CPO adoption. As CPO depends heavily on semiconductor processes, the supply chain is led by NVIDIA (Outperform), Broadcom (Outperform), TSMC (Outperform), and outsourced semiconductor assembly and test (OSAT) companies. Component and photonic equipment makers such as Lumentum, TFC, and Senko, and Chroma ATE (Outperform) could see increasing demand as CPO solutions ramp (Exhibit 12 to Exhibit 14).

For traditional optical transceiver companies such as Innolight, their role in a CPO or NPO world is likely to be structurally reduced relative to their position in pluggable modules. Their NPO offerings are still under development and appear focused on OEs and ELS packaging.

Exhibit 15 illustrates revenue and profitability profiles across key CPO-related names; many of the supply chain companies have experienced a sharp share price appreciation, particularly since the start of 2026 (Exhibit 16).

EXHIBIT 12: **Overview of CPO supply chain and key players**

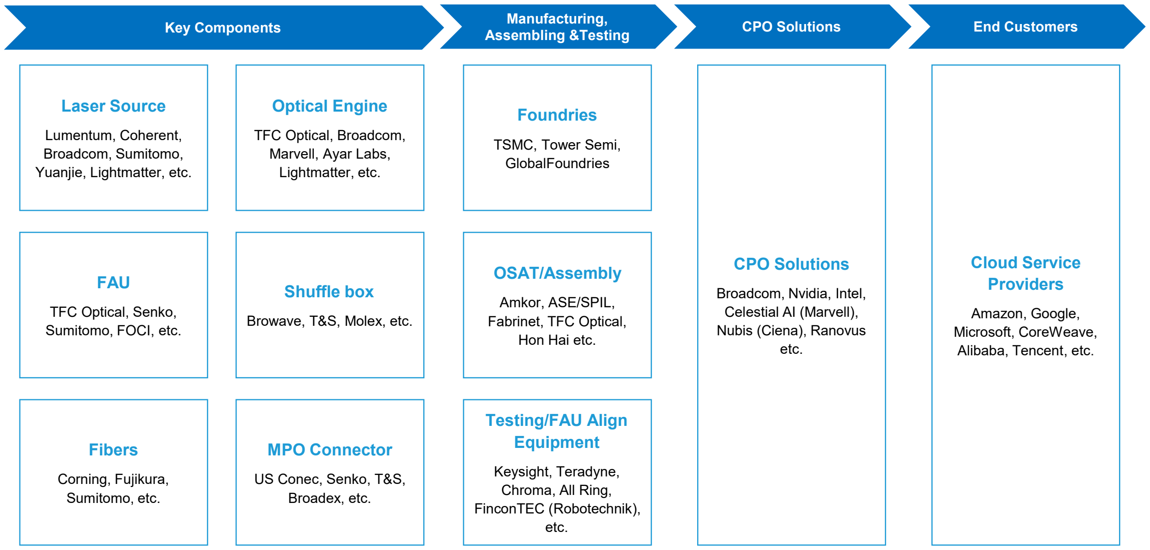

Note: Chroma, Broadcom, TSMC, NVIDIA, Intel, Amazon, Google, Microsoft, Alibaba, CoreWeave, and Tencent are covered, other companies are not covered.
Source: NVIDIA, FOCI, Bernstein analysis

EXHIBIT 13: **Overview of optical transceiver supply chain and key players**

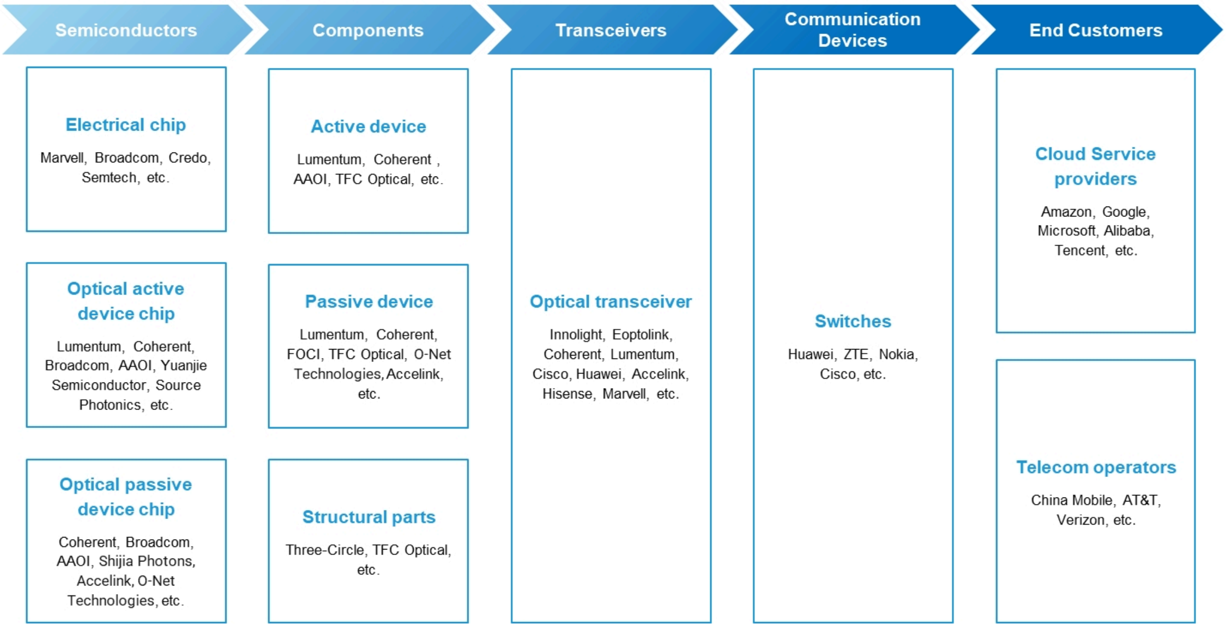

Note: Broadcom, Amazon, Google, Microsoft, Alibaba, Tencent, AT&T, Nokia, and Verizon are covered, other companies are not covered.

Source: Company reports, Bernstein analysis

EXHIBIT 14: **Electrical chips such as DSP account for a large proportion of pluggable transceivers' BOM cost**

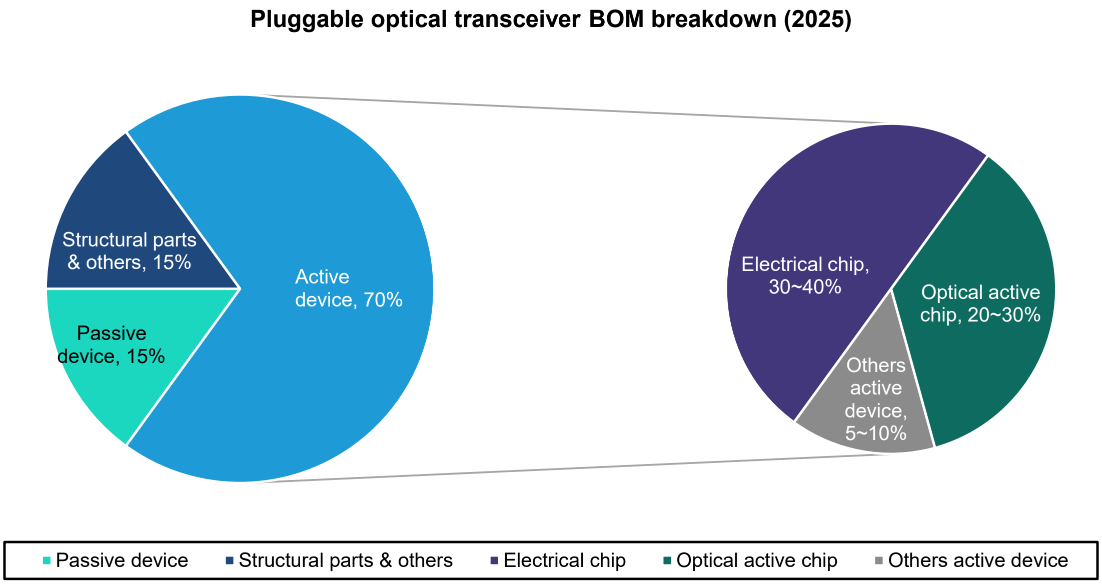

Source: Bernstein analysis and estimates (all data)

EXHIBIT 15: **Revenue and profitability profiles across key CPO-related names**

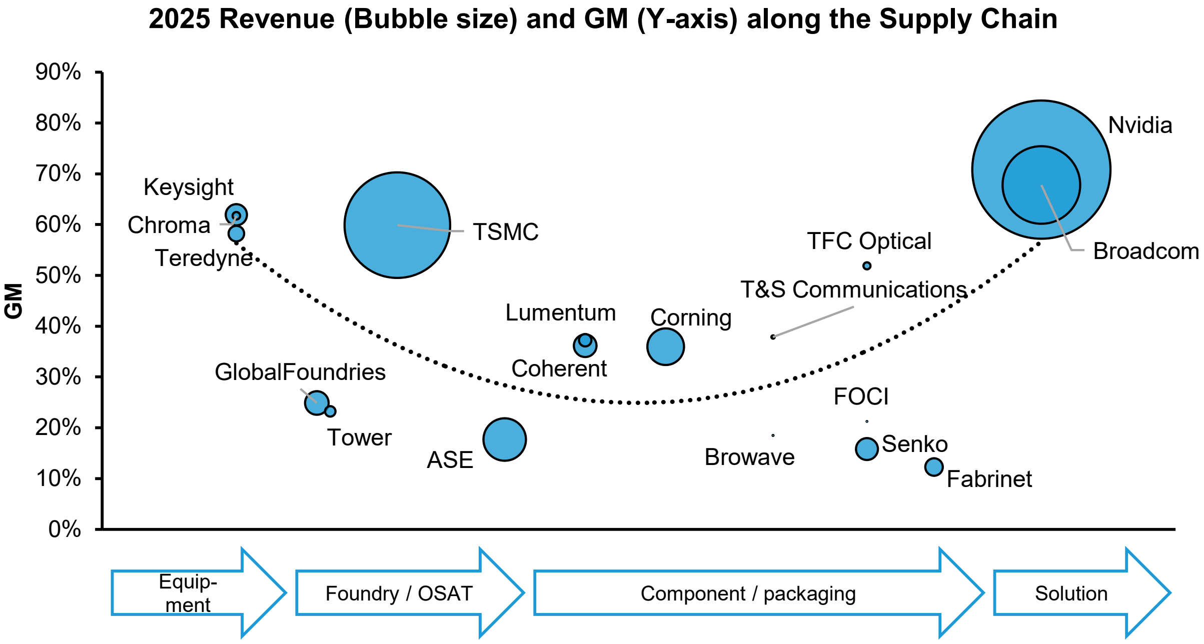

Note: Chroma, Broadcom, TSMC, and NVIDIA are covered, other companies are not covered.

Source: Bloomberg, Bernstein analysis

EXHIBIT 16: **Stock performance of CPO-related companies**

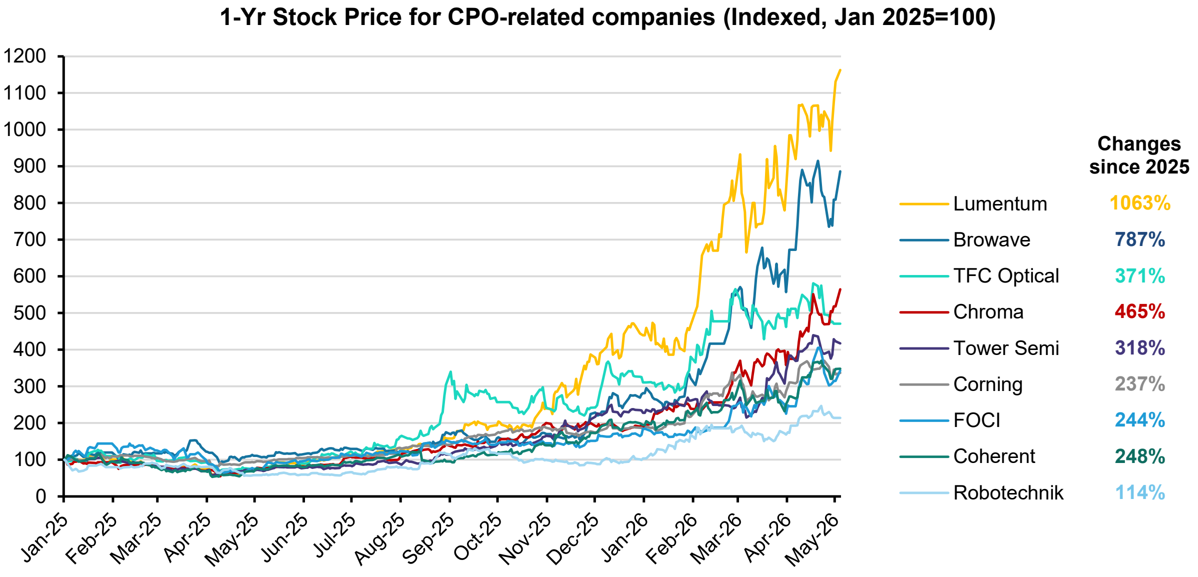

Note: Only Chroma is covered; market data as of May 4, 2026.

Source: Bloomberg and Bernstein analysis

COMPONENT SUPPLIERS

Lumentum (Lite US, not covered): Lumentum is a key supplier of EML lasers for 800G and 1.6T optical transceivers, and a pioneer in CW lasers for CPO, closely aligned with NVIDIA's roadmap (e.g., Spectrum-X). While CPO will likely replace EMLs with fewer CW lasers, the higher-value of CW lasers could keep total optical content per accelerator similar, according to management. More importantly, Lumentum's market share is likely higher in CPO-class CW/ELS than in the EML market, and moving from standalone laser chips to full ELS modules could meaningfully expand its addressable market. In its February 2026 earnings call, management said it has secured several hundred million dollars of CPO laser orders, with shipments beginning in 2027, and that it expects its first large-scale CPO "scale-up" shipments by late 2027 (Exhibit 17).

TFC Optical (300394 CH, not covered): TFC has maintained a three-year partnership with NVIDIA in the development of CPO technology and currently participates in three key segments: OEs, FAUs, and ELS modules. **(1) OEs:** TFC has long supplied OEs to NVIDIA for pluggable transceivers. TFC delivers the engines to Fabrinet (FN US, not covered) for packaging, after which the modules are supplied to Mellanox, a NVIDIA subsidiary. We expect TFC is also involved in NVIDIA's CPO/NPO OE packaging. **(2) FAUs:** TFC is likely supplying FAUs for NVIDIA's X800-Q3450 CPO switch. **(3) ELS:** TFC has also been involved in ELS packaging since 2H25, positioning it to benefit from future ELS demand in CPO.

Fabrinet (FN US, not covered): Fabrinet is a leading original design manufacturer (ODM) company providing advanced optical packaging and high-precision optical and electronic manufacturing services. Its products include optical modules, active optical cables, and lasers. With the long-standing partnership with NVIDIA, Fabrinet and Foxconn are "partners providing system-level CPO assembly and testing, as well as integrating switch-CPO assemblies and sub-assemblies into the switch system chassis."[3] Fabrinet management didn't disclose what products or solutions it provides for CPO, and our guess is that it's partnering with TFC and focuses on final assembly of components such as OEs and sub-assemblies around ASIC chips. The company has seen a small amount of CPO-related revenue and is also collaborating with two other customers in developing CPO products.

Senko (9069 JT, not covered): Senko Advanced Components, a global leader in FAUs and Multi Push On (MPO) connectivity, offers the **Metallic PIC Connector (MPC)** family as a detachable, high-density fiber-to-PIC interface for CPO and SiPh platforms. The MPC integrates a stamped metallic optical bench, freeform micro-mirror array, and precision-aligned fiber array to provide low-loss, repeatable optical coupling between fiber ribbons and PIC waveguides. Senko is recognized by NVIDIA as an official SiPh networking technology partner and has cooperated with Marvell on CPO development. Building on this MPC and FAU portfolio, Senko has also extended the concept to wafer level detachable interfaces in collaboration with GlobalFoundries, combining GF's etched-trench PIC platform with Senko's SEAT (Senko Elastic Averaging Technology) and MPC structures to enable precise, repeatable PIC-to-fiber alignment and high-throughput SiPh testing (Exhibit 18 and Exhibit 19).

FOCI (3363 TT, not covered): The company has been focusing on optical components such as fiber patch cords and fiber couplers for years, and successfully broke into the TSMC supply chain to supply FAU. The company has built a 1.6T LPO production line that is under product qualification. FOCI is building a mass production line for CPO FAU and expects product verification in 3Q26. Through its partnership with TSMC, we believe it's CPO/NPO FAU products will start shipment in 2027.

---

[3] A New Era in Data Center Networking with NVIDIA Silicon Photonics-based Network Switching | NVIDIA Technical Blog.

Corning (GLW US, not covered): Corning plays a pivotal role in CPO technology by developing the specialty fiber product tailored for CPO, as well as the configuration and deployment mechanisms needed to ensure higher reliability (Corning and Broadcom's CPO whitepaper). In May 2025, it announced plans for a collaboration with Broadcom to provide optical components for Broadcom's Bailly CPO system.

T&S (300570 CH, not covered): The company focuses on MPO and shuffle boxes sectors. It has developed patented automated systems for precision fiber alignment within shuffle boxes, creating barriers for competitors seeking alternative approaches. As a long-term partner of Corning, T&S contributes to the global AIDC supply chain by handling subcontracted production for Corning.

In addition to the companies discussed so far in this section, ABF substrate companies such as Unimicron (Outperform) should also benefit from the CPO, which leverages advanced packaging and requires an increasing area of substrate size. **Unimicron** aims to increase substrate package size from today's 90x90mm to 120x120mm body size by 2028. Companies that have strong optical know-how, including **Himax (HIMX US, not covered)** and **Largan (Market-Perform)**, could also benefit from CPO adoption depending on their product maturity. Largan confirmed that it is conducting dedicated in-house optic R&D targeting CPO applications, but underscored that mass production feasibility remains uncertain, given the technological complexity. For NVIDIA's upcoming NPO switch, socket suppliers such as **Lotes (3533 TT, not covered)** could also play a role. Lotes has been developing sockets for NVIDIA's compute trays, though that program has seen delays. The collaboration could naturally extend into future CPO front.

EXHIBIT 17: **Lumentum's external light source (ELS) for CPO**

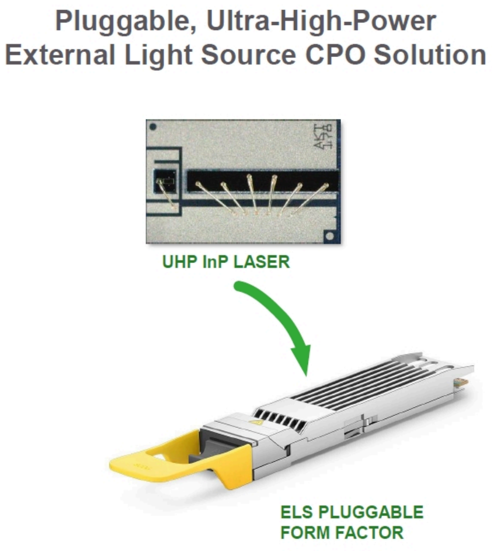

Source: Lumentum

EXHIBIT 18: **Senko's MPC use stamped metallic optical benches with intergrated micro mirror arrays to focus light between fibers and PIC**

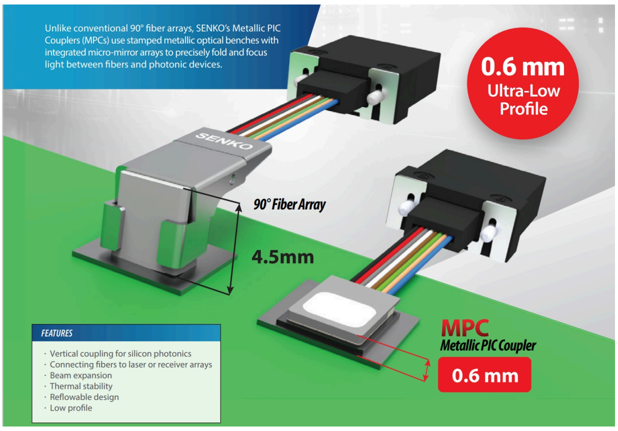

Source: Senko

EXHIBIT 19: **Senko's products for CPO**

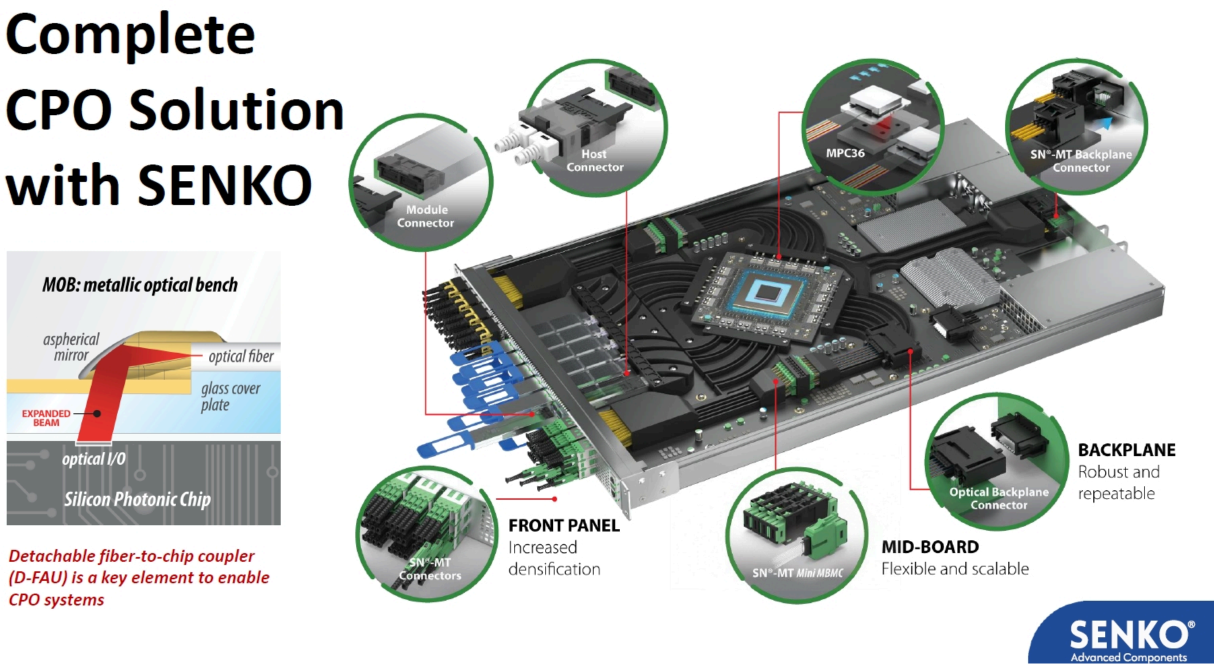

Source: Senko

EQUIPMENT SUPPLIERS

Chroma ATE (Outperform): Chroma ATE offers a growing portfolio of CPO-adjacent photonics test systems, including testers for PIC wafer and OE, laser diode burn-in, and reliability testers that target SiPh/PIC manufacturing. These systems combine automation, precise temperature control, and multi-channel optical and electrical measurement, and build on Chroma's long experience testing laser diodes and active optical components, which should translate well as CPO production scales. We expect its testers shipment to start in 2H26 with more revenue generation in 1H27 (Exhibit 20).

Teredyne (TER US, not covered): Teredyne UltraFLEXplus ATE exercises high-speed electrical I/O of the EIC. The company has partnered with ficonTEC to offer a double-sided wafer probe test cell, which integrates Teredyne's test platform and programming environment with ficonTEC's optical alignment and probing technology (Teradyne Announces Production System for Double-Sided Wafer Probe Test for Silicon Photonics). The solution helps foundry customers to test the "known good die" (KGD) before wafers are diced and packed into CPO devices.

ficonTEC (Robotechnik, 300757 CH, not covered): In March 2025, ficonTEC launched its 300mm double-sided elctro-optical wafer tester that enables the high-throughput test of silicon photonics. This tester is compatible with existing automated test equipment (ATE) at the software and hardware levels, enabling high data rate test capability on the top side, and precision optical six-axis active alignment probing below. ficonTEC is a Germany-based equipment developer, first acquired by a group of Chinese investors in 2020. In 2025, Robotechnik acquired a majority stake in the company for ~RMB1Bn, bringing the implied total acquisition cost for 100% of the shares to ~RMB1.2Bn. With ficonTEC's book value at roughly RMB200Mn and net profit around RMB40Mn, the transaction generated ~RMB1.1Bn in goodwill.

All Ring (6187 TT, not covered): We believe the company has been a key beneficiary of TSMC's CoWoS supply chain over the past few years. It has dedicated ~25% of its workforce to CPO-related projects, with key products in 2026 being automated fiber coupling equipment. With multiple customers under engagement, All Ring expects its CPO equipment revenue to start picking up in late-2026 to 2027. With multiple FAUs per switch and a low units per hour (UPH), All Ring expects strong demand for this equipment once CPO adoption begins.

MPI (6223 TT, not covered): The company provides a wafer-level mechanical and optical probing platform, covering insertion loss, optical eye diagrams, etc. Its 300mm wafer probes enable multi-axis fiber-array alignment and wide temperature control, and can be integrated with ATEs from third parties.

EXHIBIT 20: **Chroma ATE CPO testers**

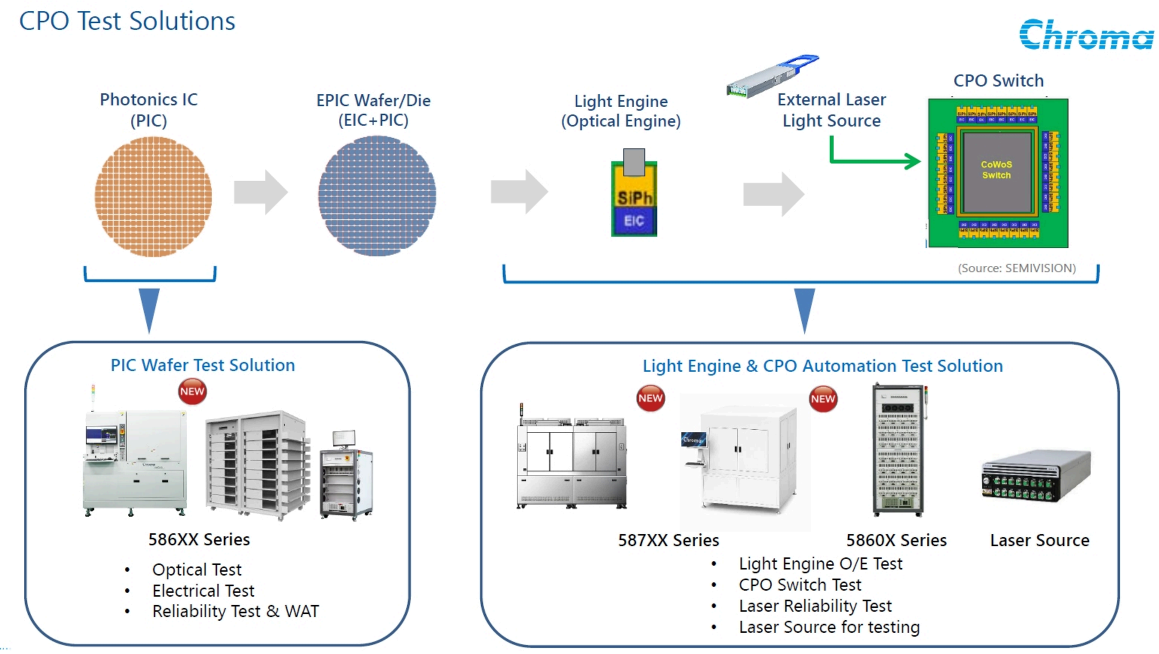

Source: Chroma ATE

STARTUPS IN CPO SUPPLY CHAIN

Celestial AI (acquired by Marvell (MRVL US, not covered): Celestial AI has developed an innovative technology, the **Optical Multichip Interconnect Bridge (OMIB)**, which is an advanced in-package optical interconnect architecture built on TSMC's CoWoS-L technology. OMIB places the interposer directly beneath the XPU/ASIC rather than locating optics at the board edge, as in traditional designs, and integrates photonics directly onto the embedded bridge. This approach bypasses shoreline limitations and enables faster, more efficient data movement. Celestial AI adopts electro-absorption modulators (EAMs) due to their relatively low thermal sensitivity and compact footprint, as the modulators are positioned close to high-power chips in the company's architecture. Marvell (MRVL US, not covered) announced the acquisition of Celestial AI in December 2025 and the deal was closed in February 2026. The deal is mainly to enhance Marvell's product solution for scale-up interconnect, and Marvell expects the Celestial revenue run rate to reach $500Mn by 4QFY2028 and $1Bn by 4QFY2029 (Exhibit 21 to Exhibit 24).

Nubis (private): Nubis provides OE chiplets designed to integrate directly with XPUs and ASICs. The company adopts MZM modulators due to their interoperability, reliability, and maturity. In addition, the company favors grating coupling and has emerged as an early mover in this approach. Nubis has partnered with Samtec to deliver optical modules compatible with Samtec's CPC connector. The company was acquired by Ciena (CIEN US, not covered) for $270Mn in 2025. Ciena targets to provide CPO/NPO solutions for inside and between racks by combining its high-speed SerDes with Nubis's OEs.

Ayar Labs (private): The company is a startup backed by many industry giants, including NVIDIA, AMD, TSMC, GlobalFoundries, and Intel. Earlier, it collaborated with GlobalFoundries for manufacturing, but has now pivoted to TSMC's COUPE platform, aiming to further enhance product performance. Its core products are the "Teraphy" optical chiplet and the "SuperNova" laser module. **A key differentiator is that the chiplet functions like a standards-compliant UCIe 2.0 retimer on the electrical side, making it compatible with a wide range of vendor packages**, including GPUs, CPUs, NPUs, switches, and memory devices. This positions Ayar Labs firmly within the open ecosystem rather than pursuing a proprietary solution. The company also demonstrated its reliability through long-run error tests, margin tests, and thermal stress tests at the 2025 Hot Chips conference (Exhibit 25 and Exhibit 26).

Lightmatter (private): The company is known for its 3D **optical interposer** solution "Passage," which is a thin silicon substrate that integrates multiple silicon chiplets, including waveguides, modulators, photodiodes, and couplers, within a single package located underneath or adjacent to XPUs/ASICs. The Passage M1000 interposer reaches sizes of up to 4,000 mm². The company adopts MRM modulators with active control loops to handle thermal challenges and to enable dense wavelength-division multiplexing (DWDM) on the interposer (Exhibit 27).

EXHIBIT 21: **Startups in CPO supply chain**

	Celestial AI	Nubis	Ayar Labs	Lightmatter
Core products	Photonic Fabric™	Vesta 100 1.6T NPX optical engine linear	TeraPHY optical engine chiplet	Passage™ M1000 3D Photonic Superchip
Power consumption (pJ/bit)	3.1	4.0	<5	4.6
Key differentiator	1. Optical Multichip Interconnect Bridge (OMIB) technology: places the interposer directly beneath the XPU/ASIC and integrates photonics directly onto the embedded bridge 2. Adopts EAM due to its low thermal sensitivity	1. Uses MZM modulators 2. Adopts grating coupling 3. Designed to be compatible with CPC-style detachable connectors	1. Chiplet exposes a standards-compliant UCIe 2.0-class electrical interface with integrated retiming, enabling compatibility with a wide range of host packages 2. Use MRM modulator or dense WDM and low-energy optical I/O	1. Employs a large multi-reticle photonic interposer that sits beneath multiple chips and can reach sizes of 4,000 mm², significantly larger than those used in Celestial AI's solutions 2. Uses MRM modulators
Commercialization	Marvell targets its revenue run rate of US$500 mn by FY2028 Q4 and US$1 bn by FY2029 Q4	N.A.	Two generations of TeraPHY silicon shipped; third generation in development and sampling with partners	Targets Passage platform commercialization in 2029
Founding round	Acquired by Marvell	Acquired by Ciena	D round	D round
Valuation	Acquisition: US$3.25bn upfront plus up to US$2.25bn earn-out linked to revenue milestones	Acquisition consideration US$270 mn	>US$1 bn	US$4.4 bn
Investors	Marvell	Ciena	GlobalFoundries, Intel Capital, Nvidia, AMD, TSMC, etc.	Google ventures, HPE, Lockheed Martin, T. Rowe Price, Fidelity, etc.

Source: Businesswire, VC news daily, company disclosures, Bernstein analysis

EXHIBIT 22: **OMIB places the interposer beneath the XPU and ASIC, and integrates photonics directly onto the embedded bridge**

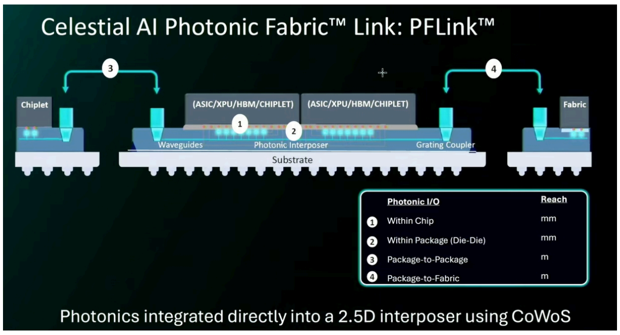

Celestial AI Photonic Fabric™ Link: PFLink™

	Photonic I/O	Reach
❶	Within Chip	mm
❷	Within Package (Die-Die)	mm
❸	Package-to-Package	m
❹	Package-to-Fabric	m

Photonics integrated directly into a 2.5D interposer using CoWoS

Source: Celestial AI

EXHIBIT 23: **Celestial AI positions optical IO in the center of XPUs and ASICs, bypassing shoreline limitations**

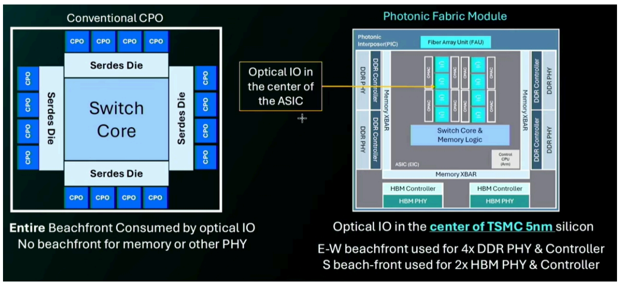

Entire Beachfront Consumed by optical IO
No beachfront for memory or other PHY

Optical IO in the **center of TSMC 5nm** silicon
E-W beachfront used for 4x DDR PHY & Controller
S beach-front used for 2x HBM PHY & Controller

Source: Celestial AI

EXHIBIT 24: **Celestial AI uses a photonic bridge architecture for interconnection**

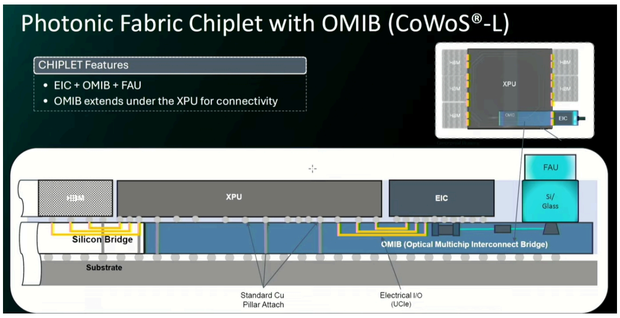

Photonic Fabric Chiplet with OMIB (CoWoS®-L)

CHIPLET Features
- EIC + OMIB + FAU
- OMIB extends under the XPU for connectivity

Source: Celestial AI

EXHIBIT 25: **Results of thermal-cycling link test demonstrated the reliability of Ayar Labs' products**

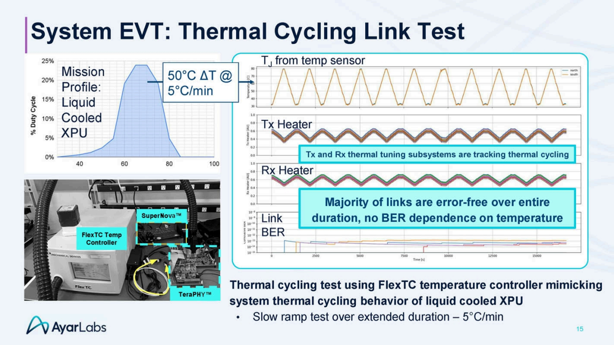

System EVT: Thermal Cycling Link Test

Tx and Rx thermal tuning subsystems are tracking thermal cycling

Majority of links are error-free over entire duration, no BER dependence on temperature

Thermal cycling test using FlexTC temperature controller mimicking system thermal cycling behavior of liquid cooled XPU
- Slow ramp test over extended duration – 5°C/min

Source: Ayar Lab

EXHIBIT 26: **Ayar Labs' core products are the "Teraphy" optical chiplet and the "SuperNova" laser module**

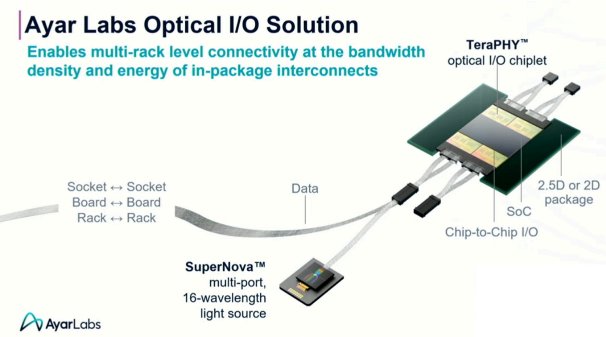

Source: Ayar Labs

EXHIBIT 27: **Lightmatter's 3D optical interposer Passage 1000 is a silicon substrate that integrates multiple silicon chiplets, including waveguides, modulators, photodiodes and couplers within a single package located underneath or adjacent to XPUs/ASICs**

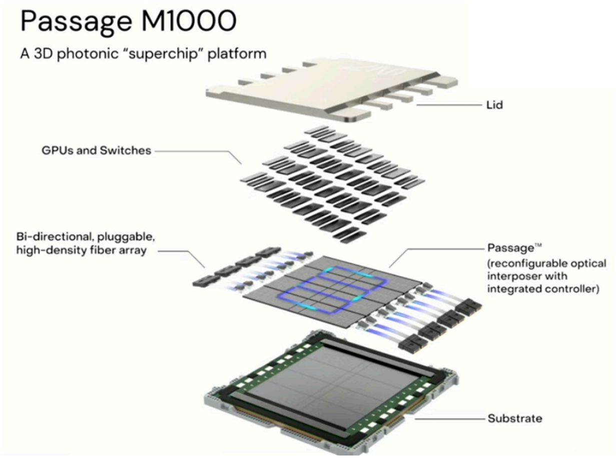

Source: Lightmatter

WHAT ABOUT GOOGLE OCS?

While this chapter has focused on CPOs that power the next-generation of connectivity in data centers, CPO is not the only path hyperscalers are exploring to solve the bandwidth, power, and topology challenges. Google has taken a different approach — OCS architecture. Its Apollo network replaced the traditional electronic spine switch with OCS, which steers light from input to output without converting signals between optics and electronics. This solution leverages 3D-MEMS mirrors to dynamically "aim" light beams between fibers, creating low-power, high-bandwidth optical paths that can be reconfigured in software to match changing traffic patterns and the needs of large AI data centers (Exhibit 28 to Exhibit 30).

By removing the optic-electronic-optic conversion in a traditional switch, OCS slashes power and cooling needs, and cuts long-term capex by eliminating multiple switch tiers, and it can serve multiple generations of link speeds (from 40G to 1.6T). This is because the traditional network needs a new generation of electronic switches when link speed changes, while OCS is agnostic to bit rate or protocol, as it just connects fiber A to fiber B.

In addition to Google, **Lumentum** is commercializing MEMS-based OCS such as the R300 (300x300 ports), with over $400Mn in backlog from three primary hyperscale customers, and OCS revenue is already starting to ramp. The product features insertion loss of 1.5dB and can support 100k GPU clusters (Exhibit 31). **Coherent**, by contrast, is pushing a liquid-crystal-based OCS product with more than 10 customer engagements.

Challenges for wider adoption of OCS include: (1) Its demand for a full stack redesign — operators must plan and schedule circuits rather than just send packets, so existing routing software, tooling, and skills don't transfer directly. (2) The optics and MEMS hardware need extremely tight alignment and low loss, which make large-port-count systems complex to manufacture, test, and maintain. (3) The 10-100ms reconfiguration means it works best for long, predictable flows (e.g., AI training) but not for short bursts of traffic. (4) The ecosystem is immature, which is why Google is pushing this technology to the Open Compute Project (OCP) and aims for defined open software APIs, control stacks, and minimum performance/

operational requirements so that multiple vendors' OCS boxes can interoperate and be managed like standard switches.

EXHIBIT 28: **Google's Apollo network replaced the traditional electronic spine switch with OCS**

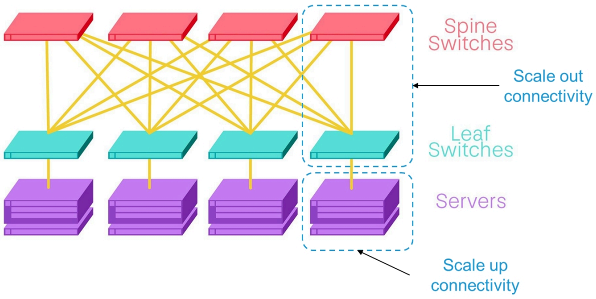

Source: Corning, Bernstein analysis

EXHIBIT 29: **Principle of OCS: Mirrors are tilted about axes to steer optical beams from one path (solid line) to another path (dashed line)**

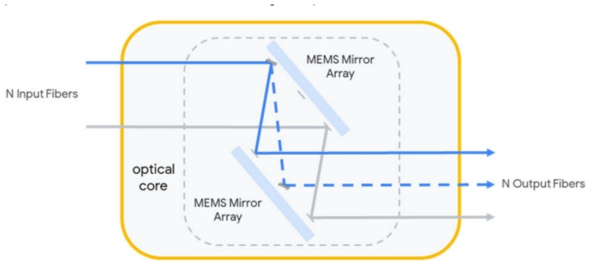

Source: Google

EXHIBIT 30: **Inside Google's OCS optical core: key components include fiber collimators, camera modules, MEMS, injection modules, dichroic splitter, and combiners**

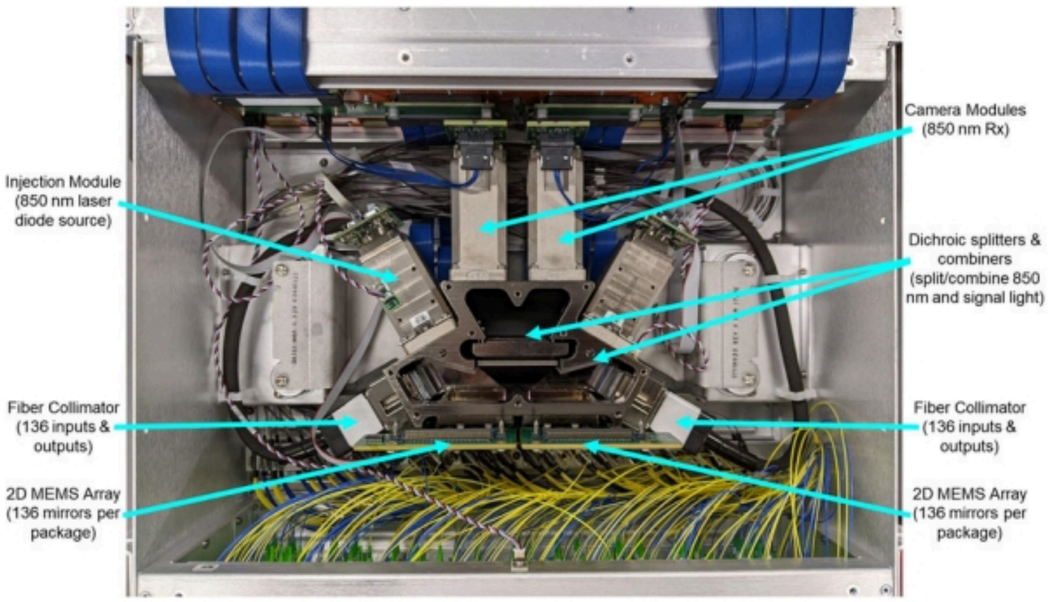

Source: Google

EXHIBIT 31: **Lumentum OCS R300**

OCS R300

Source: Lumentum

AI PCB PRIMER: TECHNOLOGY AND SUPPLY CHAIN

INDUSTRY OVERVIEW

The global PCB market size was $85Bn in 2025, and AI-driven demand is set to grow another double-digit percentage in 2026. Multi-layer PCBs (MLPCBs) are widely used in server CPU/GPU motherboards, power and thermal boards, and in switches. Layer counts are rising from a mainstream 16L+ layers to over 40L in Vera Rubin midplanes, and 70L+ in Rubin Ultra backplane. WUS (not covered) is the primary MLPCB supplier for NVIDIA's GB300 rack, followed by Unimicron and Victory Giant (VGT) (not covered). TTM, ISU, and Gold Circuit (all not covered) supply NVIDIA's High-Performance GPU eXtension (HGX) and ASIC servers and switch, alongside several emerging Tier 2 Taiwan vendors (Exhibit 1).

High-density interconnect (HDI) demand (~$16Bn) is also accelerating, and its share of the PCB market rose roughly 2ppt in 2025 to ~18%. HDI for AI servers (e.g., compute tray in the GB300 rack) requires more layers, more microvias, and higher thermal robustness than smartphone HDI. Compared to the Hopper server, the wider HDI adoption in Blackwell rack has roughly doubled the content value of HDI+PCB per GPU to $300, benefiting HDI suppliers VGT and Unimicron.

On the materials side, next-generation solutions for 2027 are now under evaluation. Copper-clad laminate (CCL) represents ~30% of PCB cost and is composed of copper foil, resin, and glass-fiber fabric. AI servers and high-speed switches increasingly use low- or ultra-low dielectric constant (Dk) glass fiber together with high-volume low-profile (HVLP) copper foil to enable cleaner high-speed signal transmission. Over 2022-25, the industry has migrated from Megtron 6 (M6) CCL to Megtron 8 (M8) to meet rising bandwidth, loss, and signal-integrity requirements. Looking ahead, leading platforms are expected to transition to M8.5 and M9 in 2026 as next-generation AI servers and network switches are deployed. These new-generation CCL materials should command significantly higher average selling prices (ASPs), reflecting more advanced resin systems, tighter process tolerances, and a slower yield ramp-up, which supports the pricing power of market leaders such as Doosan (not covered) and Elite Material Co. (EMC, not covered).

Rising package sizes and manufacturing complexity in AI accelerators are driving ABF substrate demand. ABF substrate cost depends on the package size, substrate layer count, and manufacturing techniques. Each AI accelerator generation increases both layer count and package area, as vendors integrate more logic dies and HBM stacks into a single package. Intel's Embedded Multi-die Interconnect Bridge with through-silicon vias (EMIB-T) should further enhance value capture for advanced substrate manufacturers by increasing interconnect density and packaging complexity. New materials and packaging approaches, including glass-core substrates and Chip-on-Wafer-on-Platform (CoWoP) architectures, could eventually enable larger high-performance computing (HPC) package sizes and/or improved thermal and power performance. However, given material readiness, yield challenges, and ecosystem maturity, we do not expect mass production before 2028-29.

ABF substrate ASPs are likely to rise in 2026 due to T-glass shortage. Low-CTE glass fiber (or T-glass) minimizes substrate expansion when temperature swings in AI chips. Market leader Nittobo (not covered) achieves a CTE of 2.8 x10^{-6}/°C, versus two or above for peers. We expect second sources such as Taiwan Glass and Taishan Fiberglass (both not covered) to have only limited supply for AI chips in 2026. **Ajinomoto is close to a monopolist supplier of the eponymous Build Up Film used in ABF substrate**, with mid 20's% top-line growth, driven by the combination of IC chip demand, substrate layer count and area expansion, and Ajinomoto's pricing power.

EXHIBIT 1: **Summary of PCB companies**

Ticker	Company	Mkt Cap (US$M)	Forward P/E	Revenue (US$M)		Gross Margin %		EPS
				2026E	25-28E CAGR	2026E	2027E	25-28E CAGR
Material Supplier								
600183 CH Equity	Shengyi Tech	27,703	30.2x	5,679	36%	28.9%	30.0%	49%
2802 JT Equity	Ajinomoto	31,784	33.0x	9,875	3%	36.7%	37.2%	32%
2383 TT Equity	EMC	53,862	43.5x	5,062	59%	32.0%	34.3%	82%
000150 KS Equity	Doosan	17,483	45.8x	14,364	9%	18.6%	19.5%	145%
5706 JP Equity	Mitsui Mining	15,529	31.2x	4,637	3%	26.3%	23.8%	9%
6274 TT Equity	TUC	10,133	43.3x	1,464	66%	26.3%	28.2%	88%
3110 JP Equity	Nittobo	6,876	47.8x	747	11%	38.9%	39.8%	27%
Key ABF Substrate Suppliers								
3037 TT Equity	Unimicron	45,374	48.3x	5,604	34%	22.6%	29.7%	119%
4062 JT Equity	Ibiden	24,253	59.4x	2,609	15%	31.8%	33.5%	33%
009150 KS Equity	SEMCO	42,114	47.3x	9,067	17%	23.1%	26.1%	48%
8046 TT Equity	Nan Ya PCB	20,141	67.8x	1,723	39%	18.1%	24.2%	138%
3189 TT Equity	Kinsus	8,845	44.7x	1,600	31%	26.2%	29.6%	109%
AUS GR Equity	AT&S	4,443	33.9x	2,103	23%	12.1%	20.3%	37%
Other PCB Suppliers								
300476 CH Equity	Victory Giant	47,110	26.8x	4,829	63%	40.7%	41.7%	74%
002463 CH Equity	WUS	28,694	31.5x	3,754	45%	37.5%	38.4%	54%
2368 TT Equity	Gold Circuit	23,240	36.3x	2,915	40%	36.3%	38.0%	66%
688183 CH Equity	Shengyi Electronics	13,886	33.6x	2,185	40%	33.6%	34.6%	54%
TTMI US Equity	TTM	16,353	39.3x	3,805	18%	22.1%	23%	32%

Notes: (1) Unimicron, Ibiden, and Ajinomoto are covered by Bernstein, the rest are not covered; (2) Ibiden, AT&S, Ajinomoto, Mitsui Mining, and Nittobo CAGRs based on FY26 ending March 2026; and (3) Nittobo's FY26 EPS boosted by a non-recurring real-estate asset sale; (4) All estimates from Bloomberg consensus; (5) Market data as of May 4, 2026.

Source: Bloomberg, Bernstein analysis

EXHIBIT 2: **Valuation of ABF and PCB companies**

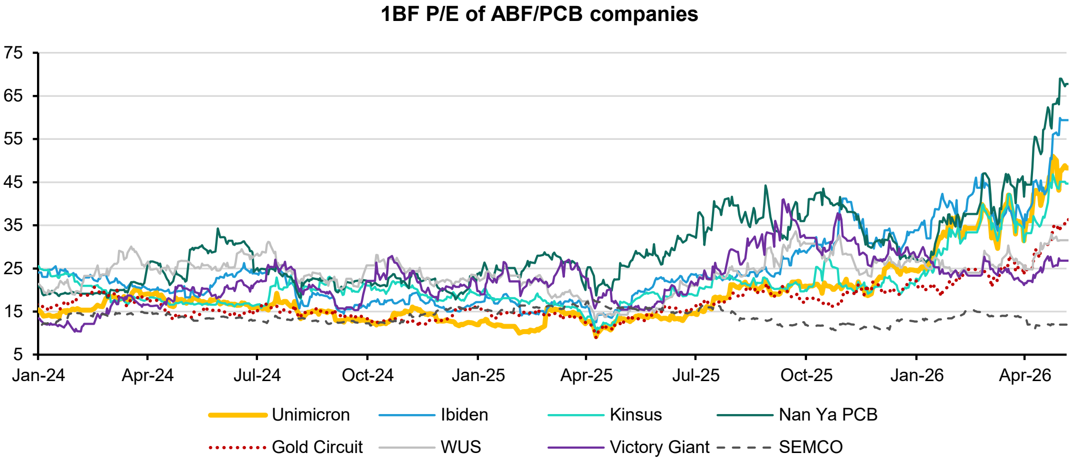

Note: Market data as of May 4, 2026; Unimicron and Ibiden are covered, the rest are not covered by Bernstein.
Source: Bloomberg, Bernstein analysis

EXHIBIT 3: **Stock price performance of ABF and PCB companies**

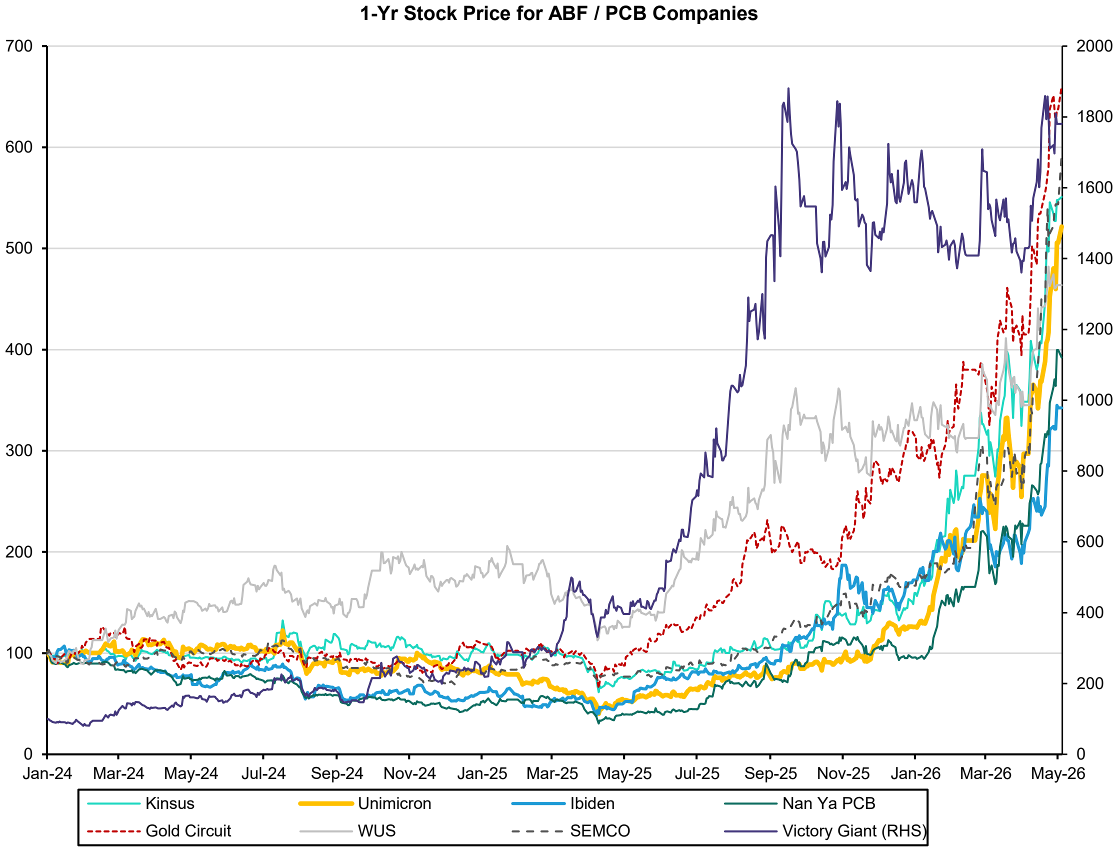

Note: (1) Stock price as of May 4, 2026; (2) Unimicron and Ibiden are covered, the rest are not covered by Bernstein; and (3) Victory Giant went up by 17x+ in past two years.
Source: Bloomberg, Bernstein analysis

OVERALL PCB MARKET

The global PCB market size was ~$85Bn in 2025, and AI-driven demand is set to grow another double-digit percentage in 2026. The PCB market, valued at $80-$90Bn, consists of a wide array of products. Roughly half of the market is taken by rigid PCB, and the remaining portion is split among flexible printed circuit (FPC), HDI, and IC substrate. Following significant corrections post-Covid-19, the PCB market saw steady growth since 2024. Following the 2025 strong rebound, Prismark expects the PCB market to maintain a mid-single-digit growth rate over 2026-27. The estimates have upside potential as the latest report from Taiwan Printed Circuit Association (TPCA) suggests the PCB market will grow in the mid-teens percentage in 2026 (Exhibit 4).

By application, the consumer sector (including smartphones and PCs) represents one-third of the PCB market, while the server and storage sector, 18% of the market in 2025, is expected to outpace the average with a 17% CAGR through 2030, reaching a 28% mix in 2030E, according to Prismark (Exhibit 5).

Exhibit 6 shows the use cases of different types of PCBs. Taking AI server as an example, the IC chip leverages ABF substrates in flip-chip BGA (FCBGA) packages. The substrate sits between IC chips (under an interposer in CoWoS packaging) and the PCB/HDI board to bridge the I/O mismatches between the two. For the main board, HDI (+5 and +6 structure) is heavily used for GPU and switch modules in both compute tray and switch tray in NVIDIA's GB300 server, while other boards such as the power distribution board (PDB) and controller boards use MLPCB. The broader adoption of HDI in the GB200/GB300 rack has driven the content value of HDI+PCB per GPU from $100-$150 to $300 (Exhibit 7 and Exhibit 8). Key suppliers in the AI PCB/ABF sector has experienced huge stock price increase (Exhibit 2 and Exhibit 3).

EXHIBIT 4: **Following the strong rebound in 2025 Prismark expects the PCB market to grow 9% in 2026**

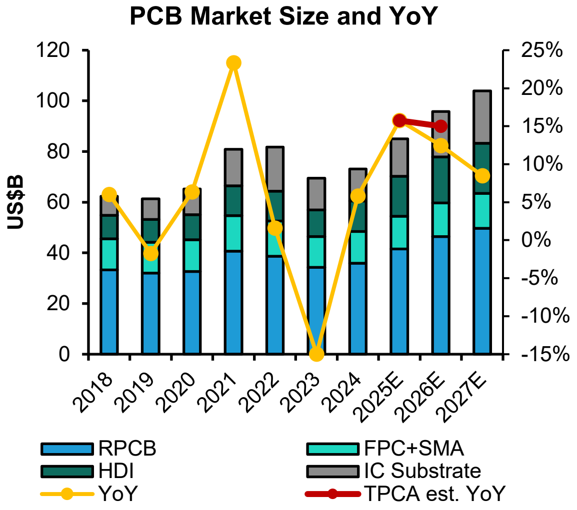

Source: Prismark data and estimates, Bernstein analysis

EXHIBIT 5: **PCB in server segment will likely grow faster than the overall market at a 17% CAGR in 2025-30E**

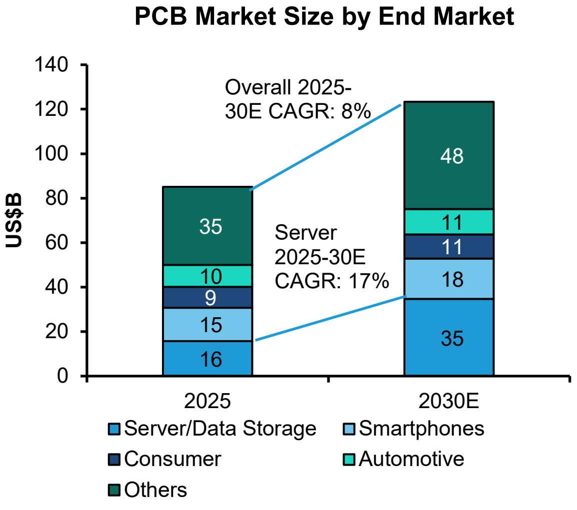

Source: Prismark data and estimates, Bernstein analysis

EXHIBIT 6: **IC substrates are mainly used in chips, while PCBs and HDIs are widely used in the main board**

Examples of PCB Type by Applications

Application	Product	PCB type	Characteristics
AI server	Chips	IC substrate (ABF)	multi-layer buildup; ideal for large-die, high-heat packages
	Board	MLPCB	2.5-10mm thickness; max aspect ratio 40:1
	Board	HDI +4 to +8	<4.5mm thickness
Networking	Chips	IC substrate (ABF+BT)	fine L/S, multi-layer buildup
	Board	MLPCB	Board thickness 2.0–4.0 mm; pattern accuracy of ±30 μm
	Board	HDI +1 to +3; 12-20L	
IoT	Chips	IC substrate (BT)	low-to-mid-pin-count BGAs packages
	Board	HDI (any-layer)	Min laser via diameter =75um; aspect ratio 0.8:1; 40μm L/S
	Board	FPC	
Auto electronics	Chips	IC substrate (BT + ABF)	
	Board	HDI +1 to +4; 6-20L	
	Board	FPC	4-6L rigid-flex
	Board	MLPCB	

Source: Victory Giant Tech, Yole, Bernstein analysis

EXHIBIT 7: **PCB market is very fragmented due to a wide range of product types and applications**

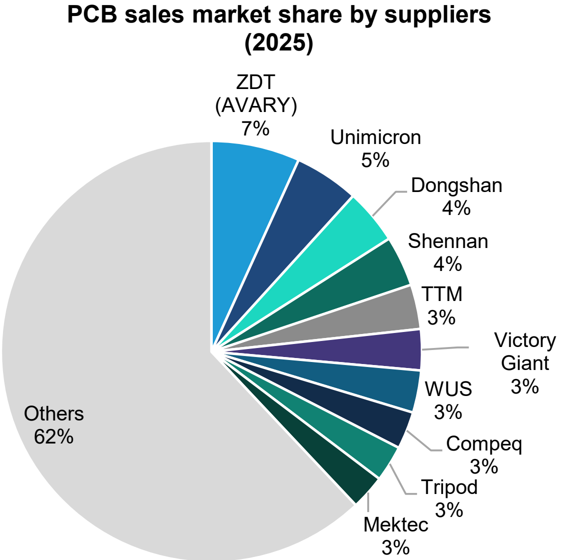

Note: Only Unimicron is covered.

Source: Prismark, Bernstein analysis

EXHIBIT 8: **The broader adoption of HDI in GB200 rack will likely drive the content value of HDI+PCB per GPU to $300**

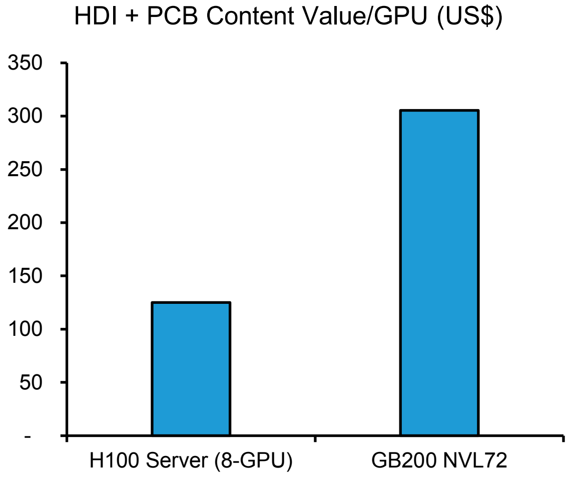

Source: Prismark, Bernstein analysis

MLPCB FOR AI SERVERS

Technology roadmap: To support more advanced, precise, and high-speed electronic functions, PCB design is trending toward higher layer counts and greater integration density. MLPCB can be grouped by layer count into mid- to low-layer boards (four to six layers) and high-layer boards (eight layers and above). MLPCB supports more complex circuits and higher-density component layouts, enabling greater functionality within a limited space, meeting the extreme density requirements of applications such as HPC and AI data centers. Many manufacturers now mass produce MLPCBs with 14+ layers, while ongoing R&D is pushing toward 30-layer and even 70-layer designs (e.g., 78L for Rubin Ultra's Backplane). Line width and spacing have also narrowed from the mainstream 100/100 μm to 75/75 μm and further to 50/50 μm or below, enabling denser routing. In 2025, the global high-layer-count MLPCB market was ~$13.7Bn, and is projected to grow at a ~6% CAGR into 2029, as per Frost & Sullivan (Exhibit 9 and Exhibit 10).

Key players: In terms of the competitive landscape, while the general MLPCB market is fragmented, the AI MLPCB market is relatively concentrated. In **NVIDIA's GB300 rack**, WUS (002463 CH, not covered) is the major supplier for MLPCB, followed by Victory Giant (VGT; 300476 CH, not covered) and Unimicron. In NVIDIA's HGX servers (with eight GPU) and **TPU** supply chain, TTM (TTMI US) and ISU (007660 KS, both not covered) are among the top suppliers. Tier 2 suppliers, including Lincstech (subsidiary of Global Brands (6191 TT, not covered)) and Dynamic Holding (3715 TT, not covered), have also made some breakthroughs in early 2026 in the TPU supply chain. In **the Trainium** supply chain, Gold Circuit (2368 TT, not covered) and Shengyi Electronics (688183 CH, not covered) are key suppliers.

EXHIBIT 9: **MLPCB manufacturing process**

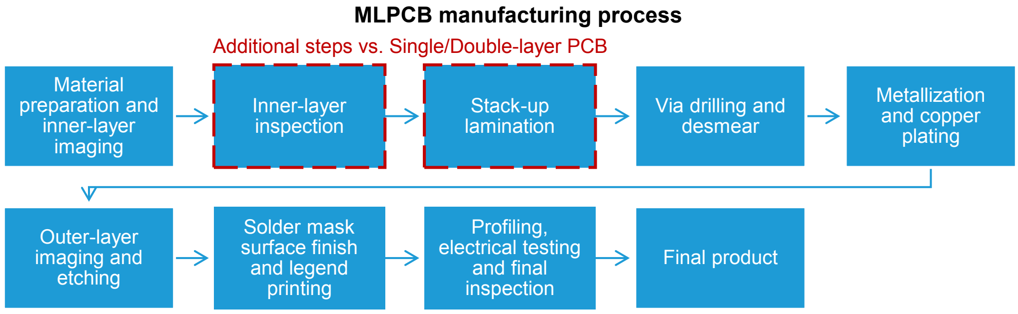

Source: VGT, Bernstein analysis

EXHIBIT 10: **With the development of AI servers, the number of layers in MLPCB has been increasing**

Application	Typical # layers	Notes for # layer count
AI server	>=16	- High-density GPU interconnect requirements (e.g. PCIe 6.0 / 112 Gbps PAM4) - Multiple power-plane distribution (AI chips consume hundreds of watts) - Extremely high signal-integrity and EMI-shielding requirements
General server	8~16	- Lower signal rates vs. AI server - Lower power-density - Relatively lower interconnect complexity
Consumer electronics	4~12	- Strict space constraints (pursuit of thin and light designs) - Medium signal speeds (≤ 32 Gbps) - Low power (single chip < 10 W)
Industrial control	4~10	- Simple signals (mostly low-speed control signals) - Moderate anti-interference requirements - Cost-sensitive

Source: iPCB, Prismark, Bernstein analysis

HDI

Technology roadmap: The HDI market size was about $16Bn in 2025, representing 18% of the total PCB market. Compared to conventional PCB, HDI's higher circuit density allows for superior data rates and compact design, making it widely used in smartphones and AI servers. HDI boards are characterized by the use of microvias/buried blind vias to connect copper layers inside (unlike conventional through-hole multilayer boards, where copper layer interconnection is achieved via through-holes). Microvias and buried blind vias are typically implemented through a build-up process (adding layers incrementally). In HDI manufacturing, the spec of HDI (e.g., 3+/4+/5+) indicates the number of build-up layers, or the number of laser drilling and lamination processes it goes through. A 4+ HDI requires four laser drilling and lamination procedures (e.g., 4+N+4 structure).

Technical requirements for HDI in the AI server (e.g., Open Accelerator Module (OAM) in HGX server structure and compute tray and switch tray in GB200/GB300 rack structure) differ significantly from those used in smartphones, as server HDI is thicker (over 2-3mm) and must withstand higher temperatures (Exhibit 11 and Exhibit 12).

Key players: Among the top five players, Compeq (2311 TT) and AT&S (AUS GR, both not covered) focus on consumer electronics, and are expanding their business in aerospace and networking. Regarding AI HDI, while Unimicron was the dominant supplier for NVIDIA's Hopper servers in 2024, Victory Giant has become the largest HDI supplier for GB200/GB300, thanks to its investment in new facilities and equipment and strong management execution. VGT has locked in many laser drilling machines and microvia filling/copper plating equipment from Mitsubishi (6503 JP, not covered) and other Japanese and German suppliers, ensuring precision in critical HDI steps that speed up yield ramp-up at new sites. In addition to having more sophisticated tools, VGT's capacity is more concentrated, with the current 20k sqm 6+ HDI located at its headquarters in Huizhou, with plans to expand into Vietnam in 2027. In contrast, Unimicron's production lines are spread across different locations in mainland China and Taiwan. Other than the NVIDIA supply chain, VGT also aims to capture the growing HDI/PCB demand from ASIC projects (e.g., Google TPU) starting from 2026. As Unimicron gradually ramps up its Taiwan HDI lines in 1Q26, we expect its market share to increase from 32% in 2025 to 35% in 2026 (Exhibit 13).)

EXHIBIT 11: **Laser drilling and copper plating are vital steps in HDI manufacturing process**

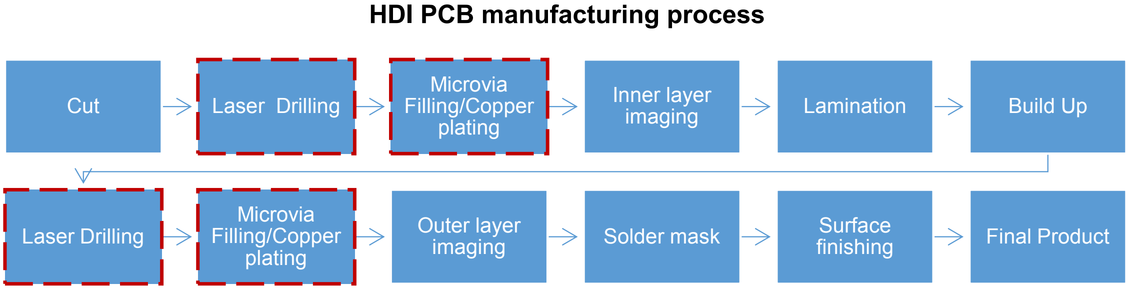

Source: VGT, Bernstein analysis

EXHIBIT 12: **Server/data center's mix in HDI increased to 14% in 2025 from ~5% historically**

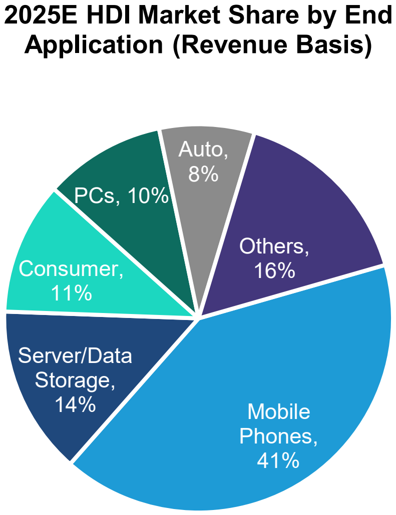

Source: Prismark data and estimates, Bernstein analysis

EXHIBIT 13: **Victory Giant's share in HDI increased meaningfully in 2025**

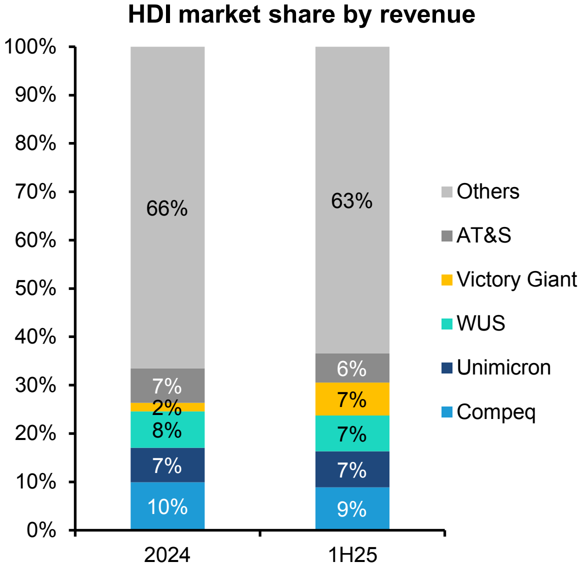

Note: Only Unimicron is covered.

Source: Prismark, Bernstein analysis

SUPPLY CHAIN OF HDI AND PCB RAW MATERIAL

The PCB/HDI supply chain starts with three key upstream materials that make up ~90% of CCL cost: Copper foil (~40% of CCL) provides the conductive layer, resin (~25%) binds everything together, and glass fiber fabric (20%) adds structural strength. These materials flow into CCL manufacturers that layer and laminate them to create CCL that makes ~30% of PCB. PCB/HDI manufacturers then use multiple CCL layers and different processes to build the final circuit boards (Exhibit 14 and Exhibit 15).

CCL is made by impregnating a reinforcing material with resin, covering it with copper foil on one or both sides, and then hot-pressing it. Glass fiber cloth serves as the reinforcing structure, and its dielectric properties along with that of the resin, together determine the overall dielectric constant (Dk) and dielectric loss (measured by Dissipation Factor (Df)) of the CCL. **A lower Dk** value allows signals to transmit faster within the material, reducing delay and making it more suitable for high-speed signal transmission. **A lower Df** value corresponds to reduced signal transmission loss. The industry currently mostly uses Panasonic's (6752 JP, not covered) resin grading system to classify CCL types (Megtron 2 or M2 to Megtron 9 or M9); the higher the grade, the lower the signal transmission loss (Exhibit 16 and Exhibit 17).

AI server PCB requires extremely clean signal transmission, necessitating the use of special low Dk/ultra-low dielectric glass fiber and low surface roughness copper foil (high-frequency, high-speed ultra-low-profile (hyper very low profile; HVLP)). As AI server demand surges, **the supply of high-end glass fibers (e.g., low-Dk2) and high-end copper foil (HVLP3 and above) are tight** as only a handful of manufacturers globally possess the capability to mass-produce those materials. This trend benefits companies that are capable of producing high-end materials. Glass fiber providers such as Nittobo (3110 JP), Asahi Kasei (3407 JP), Taiwan Glass (1802 TT), Fulltech Fiber Glass (1815 TT), and Sinoma (002080 CH) (all not covered) and HVLP producers Mitsui Mining (5706 JP), Co-Tech (8358 TT), Furukawa (5801 JP), and Fukuda (1303 JP) (all not covered), and Circuit Foil (private) are likely beneficiaries (Exhibit 18).

For the upcoming Rubin series, while M8.5 and M9 will likely be potential options, the CCL choice has not yet been determined, and different providers may choose different solutions (including a mix of M9 and M8). To our understanding, the PCB backplane, if adopted in Rubin Ultra, is likely to use M9, given its higher complexity.

Driven by the AI wave, leading CCL suppliers, including EMC (2383 TT), TUC (6274 TT), Shengyi Tech (600183 CH), Doosan (000150 KS), and ITEQ (6213 TT) (all not covered), have seen sharp share price gains since 2023, along with multiple expansion. Most suppliers have also delivered meaningful growth in both sales and gross profit (Exhibit 19 to Exhibit 22).

EXHIBIT 14: **PCB/HDI supply chain**

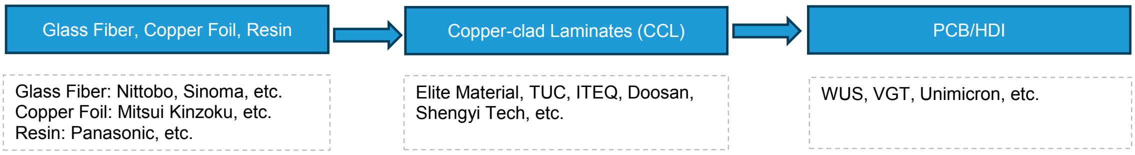

Note: Only Unimicron is covered.

Source: Company reports, Bernstein analysis

EXHIBIT 15: **CCL is the major part of PCB costs, which is mostly made of copper foil, resin, and glass fiber cloth**

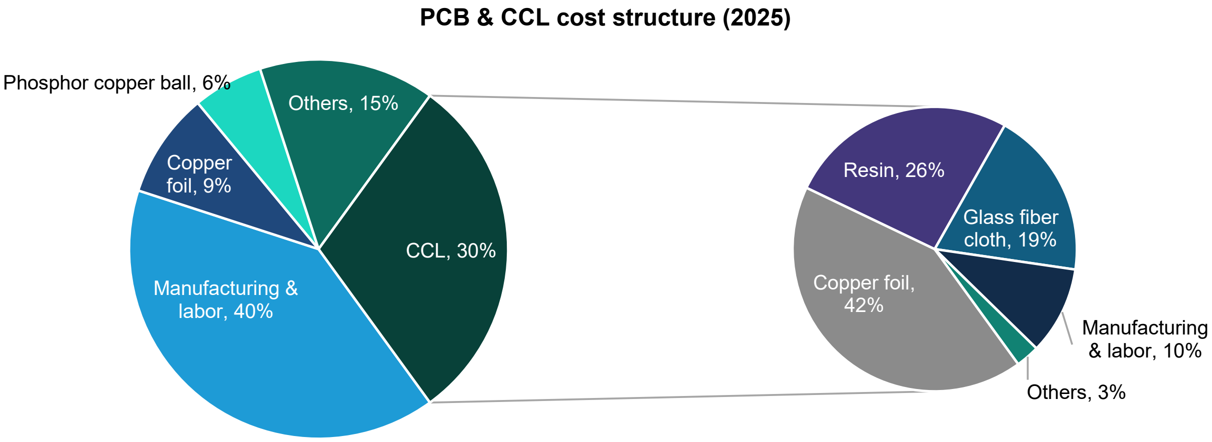

Source: CHNCI, Bernstein analysis

EXHIBIT 16: **Technology roadmap of CCL**

Generation	Material			Tech spec	Signal speed	Cost	Typical application
	Resin	Glass fabric	Copper foil				
Megtron 2	Modified EP	E-glass	ST	Dk≈4.1, Df≈0.010@1GHz, Tg≈170℃	10Gbps	1x	4G base stations, datacenter industrial control equipment
Megtron 4	Modified PPE	Low Dk	RT	Dk≈3.8, Df≈0.005@1GHz, Tg≈200℃	25-50Gbps		Traditional data center servers, switch
Megtron 6	PPE (PPO)	Low Dk	HVLP3	Dk≈3.4-3.7, Df≈0.002@1GHz, Tg≈185℃	112Gbps	1.5-2x	112G PAM4 AI Server, 5G Base Station
Megtron 8	PPE (PPO) + PCH	Low Dk	HVLP3/HVLP4?	Dk≈3.1, Df≈0.0012@14GHz, Tg≈220℃	200Gbps	3-5x	High-end AI server & switch
Megtron 9	PPE (PPO) + PCH (higher mix)	Low Dk? Quartz cloth?	HVLP4/HVLP5?	Dk≈2.8-3.1, Df<=0.002@28GHz, Tg>=250℃	224Gbps+	5-10X?	224Gbps AI server (Rubin) high-layer or high-density PCB

Note: EP= Epoxy, PPE = Polyphenylene Ether, PPO = Polyphenylene Oxide, PCH = Polycyclic Conjugated Hydrocarbon, Tg = glass transition temperature, Dk = dielectric constant, Df = dissipation factor, ST = Standard, RT = Reverse Treated, HVLP = Hyper Very Low Profile.

Source: Panasonic, Bernstein analysis

EXHIBIT 17: **Tech roadmap for CCL**

Company	PCB type	Series	Layer count	CCL
Nvidia	HDI	Hopper	18-24L	M6, M7
		Blackwell	18-22L	M8
		Rubin	24L+	M8, M9?
AWS	HDI	Trainium 2	24-26L	M6
		Trainium2.5/3	36-44L	M8
Google	MLPCB	TPU v6e/p	22-24L	M7
		TPU v7e/p	36-44L	M8

Source: TrendForce, Bernstein analysis

EXHIBIT 18: **CCL market is relatively concentrated, with top 10 suppliers capturing ~77% of market**

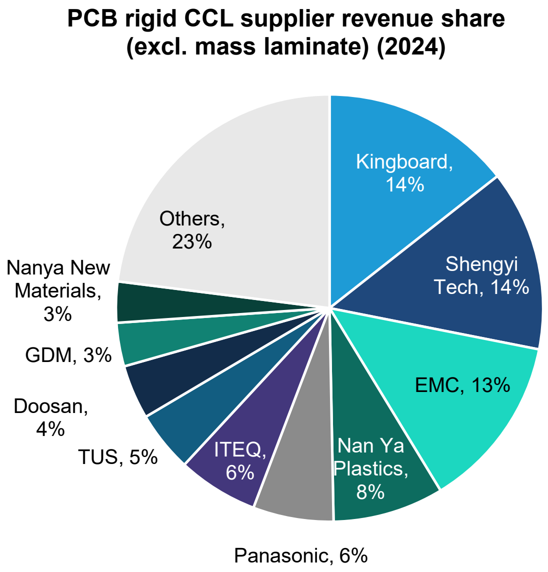

Note: None of the companies are covered.

Source: Prismark, Bernstein analysis

EXHIBIT 19: **CCL companies' stock prices have experienced significant increases in past three years...**

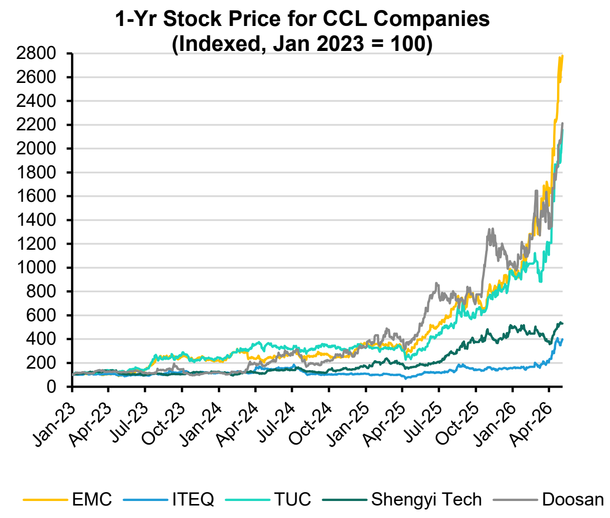

Note: None of the companies are covered. Market data as of May 4, 2026.

Source: Bloomberg, Bernstein analysis

EXHIBIT 20: **... so have their valuations**

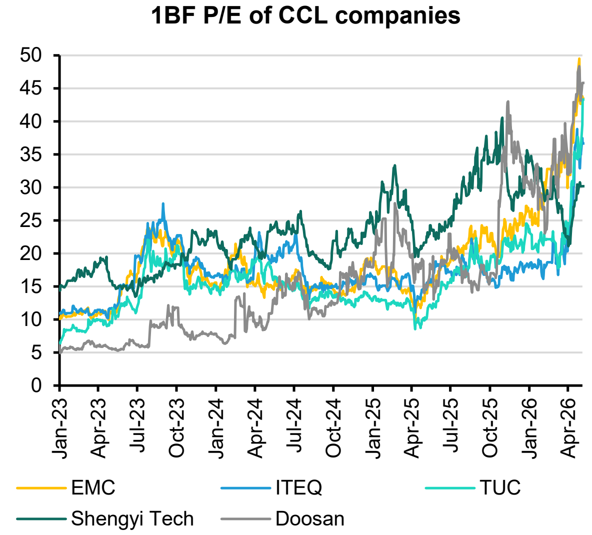

Note: None of the companies are covered. Market data as of May 4, 2026.

Source: Bloomberg, Bernstein analysis

EXHIBIT 21: **Most AI CCL suppliers' sales have experienced meaningful expansion in past three years...**

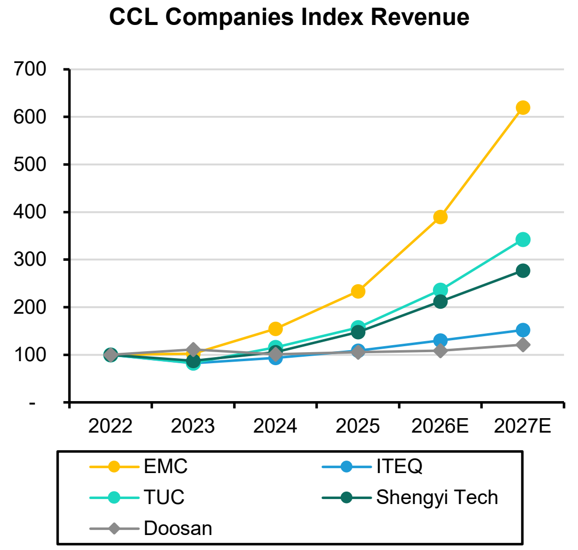

Source: Bloomberg consensus estimates, Bernstein analysis

EXHIBIT 22: **... so have their gross profits**

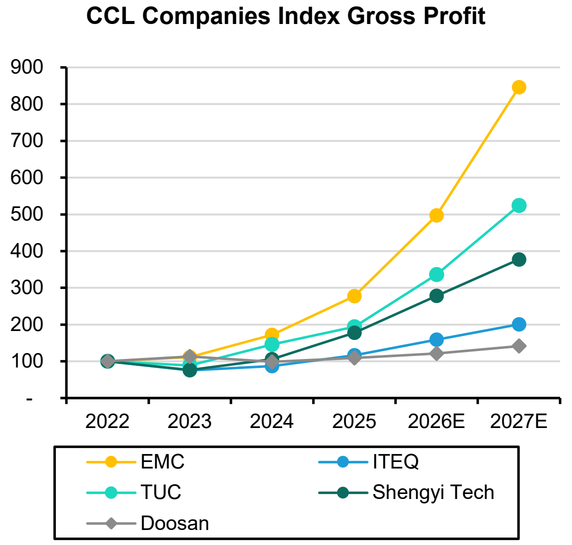

Source: Bloomberg consensus estimates, Bernstein analysis

IC SUBSTRATE MARKET

IC substrate represents 15-20% of the overall PCB market. The substrate is a material that sits between the IC chip and the board to bridge the I/O mismatches between the two. This market is much more consolidated compared to the overall PCB market, with the top 3 players — Unimicron, SEMCO (009150 KS, not covered), and Ibiden — holding ~35% share (Exhibit 23).

Different materials are used in the substrate depending on the packaging requirements for different products. For example, FCBGAs are used in the main chips of HPCs such as CPUs, GPUs, and ASICs. The average package sizes of mainstream products are around 50mm x 50mm, while some high-end chips are approaching 100mm x 100mm. The I/O counts in FCBGA can easily pass 2,000. **ABF substrate** is prominently used in this type of packaging as it allows for a high number of pins and higher transmission. Compared to that, Flip Chip Scale Package (FCCSP) is typically seen in smaller-sized chips such as smartphone CPU and DRAM. Bismaleimide triazine (BT) material is often used in FCCSP packaging (Exhibit 24).

The semiconductor industry has been exploring ways other than scaling to extend Moore's Law, such as advanced packaging. Thanks to growing server demand and the increasingly complex design in computing chips, the demand for ABF substrate surged rapidly since 2018 and, as a result, substrate's contribution to the overall PCB market has become increasingly substantial. Supply shortage emerged when Covid-19 began and persisted until the end of 2022. Following a two-year price correction, demand for high-end ABF substrate (for AI accelerators) gained momentum in 2025 and the tight supply in upstream materials is likely driving a new round of price hikes starting early 2026. We, therefore, expect overall substrate mix in the PCB market to grow to ~20% by 2027 (Exhibit 25).

EXHIBIT 23: **Substrate market is more consolidated compared to overall PCB market, with top 3 players holding ~35% of share**

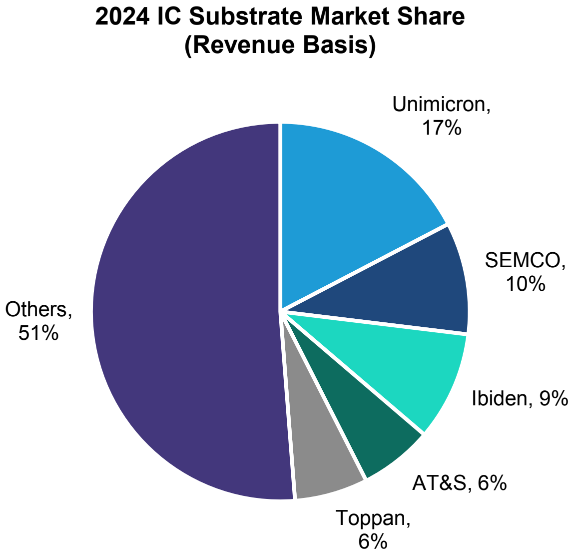

Note: Only Unimicron and Ibiden are covered.

Source: Yole, Bernstein analysis

EXHIBIT 24: **IC substrate by packaging technology**

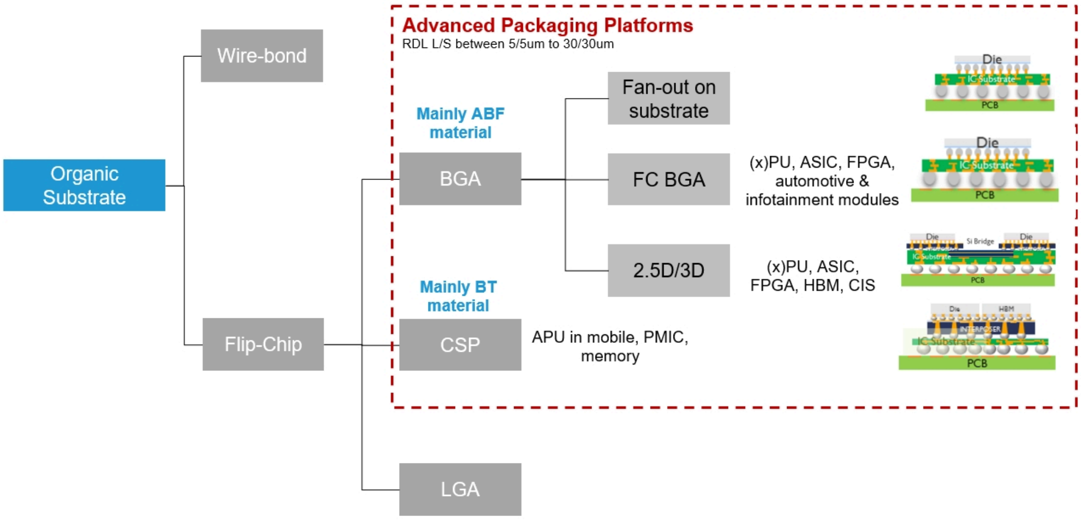

Note: BGA = Ball Grid Array, CSP = Chip Scale Package, LGA = Land Grid Array, APU = Application Processor Unit.
Source: Yole, Bernstein analysis

EXHIBIT 25: **We expect overall substrate mix in PCB market to grow to ~20% by 2027**

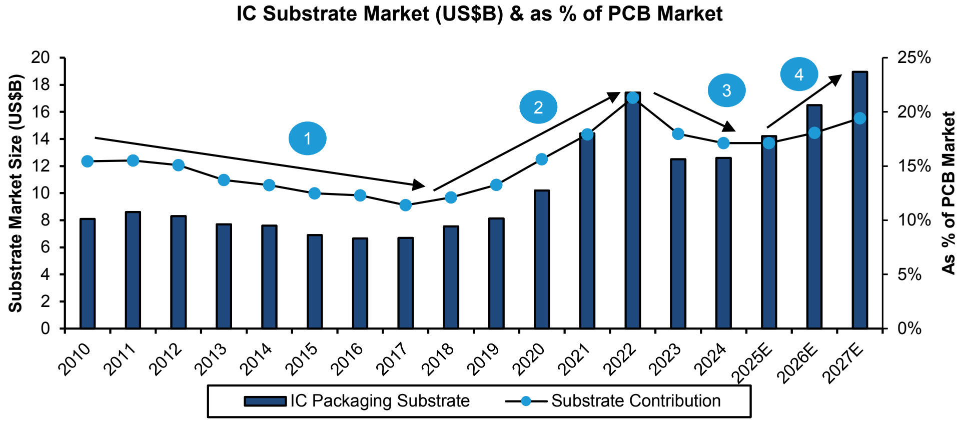

1. 2010-17:
- Increasing BT substrate demand from smartphone can't fully offset the decreasing ABF substrate demand on decaling PC market

2. 2018-22:
- Growing server market and 5G base station
- Strong PC recovery amid COVID
- Emergence of advanced packaging

3. 2023-24:
- Correction of PC & traditional server market post COVID
- Capacity overbuilt starting from 2022

4. 2025 onwards
- ABF substrate demand and ASP recover gradually driven by Gen-AI

Source: Prismark, Bernstein analysis and estimates

ABF SUBSTRATE

Technology roadmap: The cost of ABF substrate per chip depends on the processing time and yield, which are determined by the package size and substrate layer count, as well as manufacturing techniques. Compared to server CPUs from AMD and Intel, while AI accelerators feature fewer layers in substrate, they still carry decent ASP and gross margin, as it requires sophisticated manufacturing techniques due to the need for stability under high temperatures. The complexity of AI accelerators is also increasing, as evidenced by different generations of NVIDIA chips in Exhibit 26. This combined with the increasing volume of chips, is fueling the demand for high-end ABF substrate.

As of 2025, the majority of AI accelerators are leveraging TSMC's CoWoS packaging, which adds silicon interposer in the front-end. In contrast, Intel's Embedded Multi-die Interconnect Bridge (EMIB) adds silicon bridge in the substrate, which increases the substrate manufacturing complexity. The Ponte Vecchio GPU and Sapphire Rapids CPU have embedded 11 and 14 EMIB bridges, respectively, to connect different dies. The next variant of EMIB is called **EMIB-T**, which adds vertical TSVs within the bridge die and leads to better power delivery and signal routing in the z-direction (Exhibit 27).

Update for glass-core substrate: Today, the majority of the substrates are organic and leverage PCB-like cores with woven-glass laminates. The semiconductor industry is exploring new materials in response to the increasing size of the advanced packaging. Glass core substrates (GCS) feature a layer of glass as the core material, which is sandwiched between build-up dielectric structures (e.g., ABF) on opposing sides. Intel introduced GCS in 2023 and expects the new technology will offer 50% less pattern distortion and potentially 10x increase in interconnect density compared to organic core substrate (Exhibit 28 to Exhibit 30).

Today, the industry faces many challenges toward commercialization, including glass handling, lack of standardization of glass panel size, and glass singultation resulting from the brittleness of glass. Substrate manufacturers need to use laser ablation (CO_2 or UV) or ultrasonic drilling to build TGV (through glass via), and face challenges to metalize TGV due to poor adhesion of metal to glass. A specialized tool for bar glass inspection and TGV vias metrology is also needed as existing tools struggle to perform with transparent glass.

In January 2026, Wccftech reported that **Intel's** pilot line for GCS in Arizona will embed glass core in its EMIB technology.[1] The substrate will support roughly twice the silicon size (2x 858mm^2) to hold logic and memory and has 10 layers (10-2-10 structure). **Absolics**, a subsidiary of SK group (034730 KS, not covered), has invested $600Mn in Georgia for a manufacturing site dedicated to GCS, with the first phase completed in 2025 and ready for product sampling in 2026. Both SEMCO and LG Innotek (011070 KS, not covered) have built pilot production lines for GCS.

Despite rising interest, GCS adoption is limited by high production costs and an immature supply chain. We think product samples may be available in 2026, but mass production is unlikely before 2028. Yole forecasts GCS revenue of $160Mn in 2027 and $400Mn by 2030.

Is CoWoP a threat to substrate demand? In mid-2025, the market discussed a new packaging design termed Chip on Wafer on Platform PCB (CoWoP), which aims to eliminate substrates between chips and PCBs to enhance signal transmission speed and thermal stability. We believe there are significant challenges ahead toward commercialization, including managing CTE (Coefficient of Thermal Expansion) mismatch and mechanical stress between PCB and chips, improving PCB densities close to ABF substrate, and yield and reliability concerns. Our channel

---

[1] https://wccftech.com/intel-showcases-glass-core-substrates-with-emib-advanced-packaging/

check suggests this new technology will stay in R&D and prototype phase in 2026-27.

Key players in AI ABF substrate market: Ibiden and Unimicron are the core ABF substrate suppliers for NVIDIA's AI GPU, switch chip, and CPU, with Ibiden as the lead supplier, especially for the new Rubin platform. Unimicron is also the major supplier across all different ASIC platforms (Google TPU, AWS Trainium, and Meta MTIA). Other suppliers for AI ASIC chips include Nan Ya PCB (8046 TT), SEMCO (009150 KS), and Toppan (7911 JP, not covered) (Exhibit 31 and Exhibit 32).

EXHIBIT 26: **GPUs and ASICs generally feature less layer counts compared to server CPUs; there's a trend of increasing substrate layer and size for AI chips**

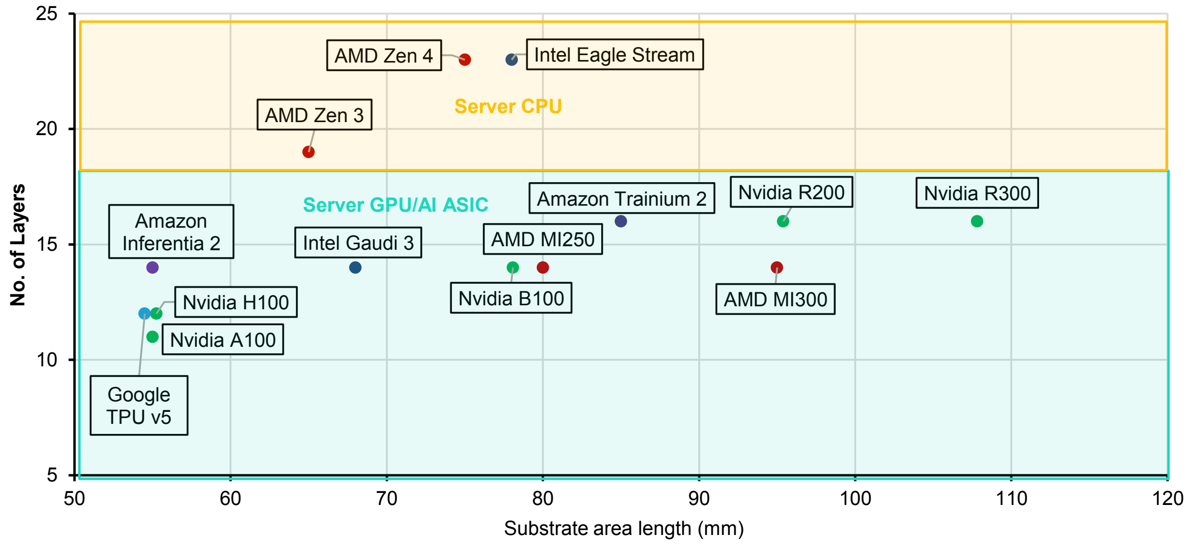

Source: Yole, Prismark, Bernstein analysis

EXHIBIT 27: **Compared to TSMC's CoWoS that adds silicon interposer in the front-end, Intel's EMIB silicon bridge is embedded in the substrate, which increases substrate manufacturing complexity**

CoWoS (Chip on Wafer on Substrate)

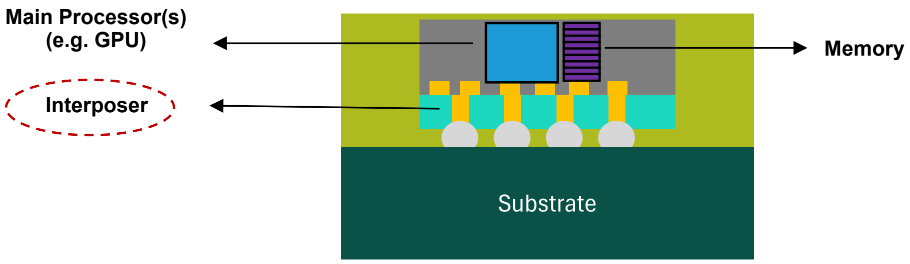

EMIB (Embedded Multi-die Interconnect Bridge)

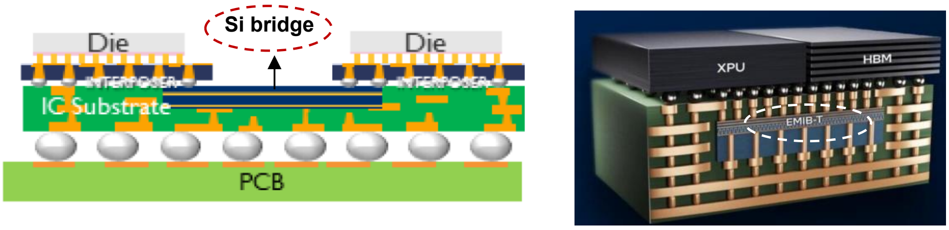

Source: Intel, TSMC, Bernstein analysis

EXHIBIT 28: **Organic core substrate versus glass core substrate**

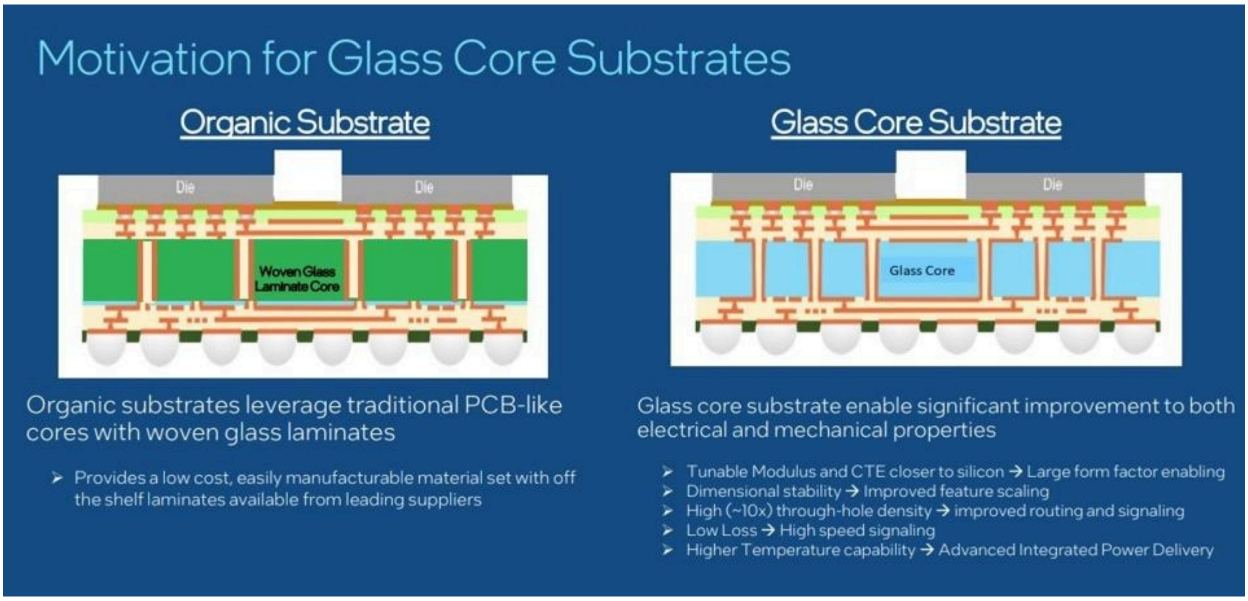

Motivation for Glass Core Substrates

Organic Substrate

Organic substrates leverage traditional PCB-like cores with woven glass laminates

- Provides a low cost, easily manufacturable material set with off the shelf laminates available from leading suppliers

Glass Core Substrate

Glass core substrate enable significant improvement to both electrical and mechanical properties

- Tunable Modulus and CTE closer to silicon → Large form factor enabling
- Dimensional stability → Improved feature scaling
- High (~10x) through-hole density → improved routing and signaling
- Low Loss → High speed signaling
- Higher Temperature capability → Advanced Integrated Power Delivery

Source: Intel

EXHIBIT 29: **Advantages of glass core substrates and the challenges of mass adoption**

Advantages (vs. Organic Materials)	Challenges for Adoption
Excellent Thermal Properties	**Brittleness**
Lower coefficient of thermal expansion (CTE) and better dimensional stability under temperature variation. Its CTE can be fine-tuned to closely match the CTE of chips	More likely to break compared to other materials. It requires careful handling during manufacturing and assembly
Electrical Capability	**High Manufacturing Cost**
Lower dielectric constant, which reduces signal loss and interference and make it suitable for high-speed data transmission applications (e.g. Gen-AI.).	1) Higher material cost vs. organic core materials 2) Require specialized equipment and know-how 3) Through-glass-via (TGV) manufacturing is challenging
Mechanical Stability	**Limited Application Range**
Higher stiffness and is thus more stable, which enables ultra-large form-factor packages with high assembly yields	The relatively higher cost and brittleness makes them less suitable for applications incl. smartphone, auto, etc.

Source: Intel, Bernstein analysis

EXHIBIT 30: **Glass core substrates are better in performance but less cost-effective than organic core IC substrates**

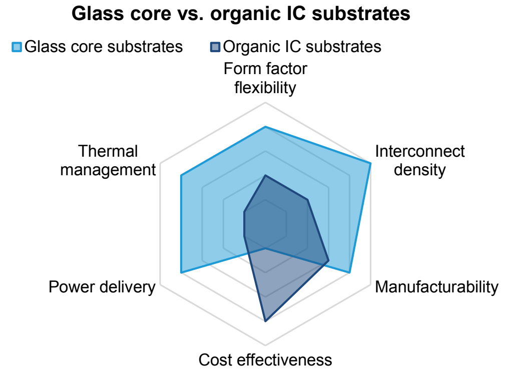

Source: Yole, Bernstein analysis

EXHIBIT 31: **High-end ABF substrate represents over one-third of market demand**

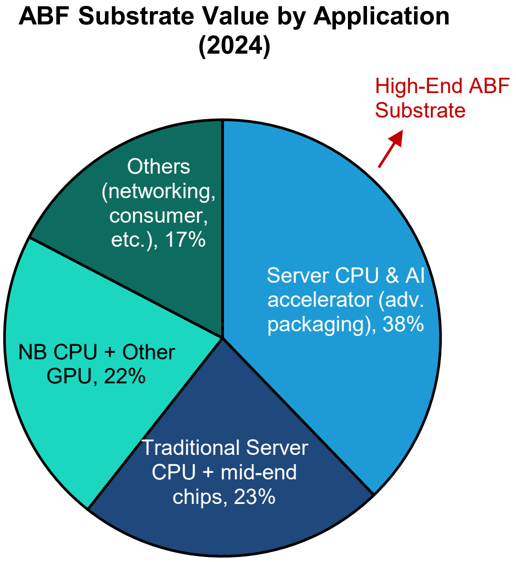

Source: Prismark, Bernstein analysis

EXHIBIT 32: **Top 3 represent over 50% of ABF substrate market**

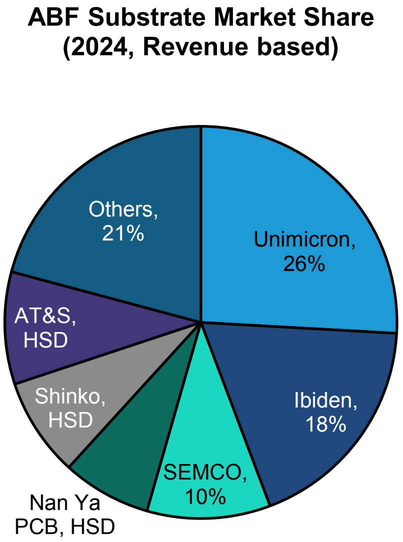

Note: HSD = high-single digit; only Unimicron and Ibiden are covered.
Source: Company reports, Bernstein analysis and estimates (all data)

RAW MATERIAL FOR ABF SUBSTRATE

Update on T-glass shortage situation: One of the most important substrate materials is low-CTE (Coefficient of Thermal Expansion) glass fiber (or T-glass), which minimizes substrate expansion when temperature swings in AI chips. The material is essential to prevent ball grid array (BGA) solder joint failure and package warpage in AI accelerators. **The CTE of industry leader Nittobo reaches 2.8 x10^-6/°C (every 1°C increase in temperature, the material's length increases by 2.8×10^{-6} of its original length), while its peers products are at 3.2 x10-6/°C or above.**

The T-glass shortage is likely to persist until 4Q26, as new capacity from the incumbent leader Nittobo will not be available until end-2026 to 2027. Nittobo's new factory in Taiwan, focusing on glass fiber yarn production, is set to begin mass production in late 2026 at the earliest, but it will still take time to convert to fiberglass fabric and obtain customer certification. Its new factory in Japan is guided to start production in 2027 and begin shipment in 2028. Second sources such as Taiwan Glass (1802 TT, not covered) and Taishan Fiberglass (a subsidiary of Sinoma (002080 CH, not covered)) have limited supply for AI chips in 1H26. Sinoma plans to quadruple its monthly Low-CTE glass (3.2 x10-6/°C) output from 50k meter per month in end-2025 to 200-300k meter per month by the end of 2Q26, though it requires time to ramp up the yield.

As suppliers optimize their product portfolio by allocating T-glass to high-end products (e.g., AI accelerators), ABF substrate suppliers negotiated new prices across customers in late-2025, and we think the new pricing should reflect in 1H26. The T-glass shortage has created a ripple effect on BT substrate, as ~20% of T-glass is used on BT substrate for high-end smartphones and networking ICs. Kinsus (3189 TT, not covered) noted that BT prices have seen a sequential increase every quarter since 3Q25, and that this is likely to continue in 1H26. Within the PCB business segment, Kinsus, Nanya PCB, LG Innotek, and SEMCO have relatively large sales exposure to BT substrate (Exhibit 33 and Exhibit 34).

ABF film market: ABF is an amino acid-derived film that is used as a thermal barrier within IC chip packaging substrates and has been adopted by the substrate industry to describe the overall category.

Ajinomoto is close to a monopolist supplier of the eponymous Build Up Film used in ABF substrate, with mid-20% top-line growth, driven by the combination of IC chip demand, substrate layer count and area expansion, and Ajinomoto's pricing power. The company generates a 61% margin from ABF, with consistent 1.5-1.7x operational leverage (Exhibit 36), and its ABF division accounts for ~25% of earnings.

In mid 2025, the company roughly doubled the scale of its Gunma plant (one of two sites in Japan) and management still has the ability to add a second production shift at both sites. We see very limited near-term risk of an ABF film supply crunch. Ajinomoto is working to expand its ABF TAM into mobile devices (which tend to use BT substrate) via its ABF-RCC product (Exhibit 35).

EXHIBIT 33: **Substrate supply chain**

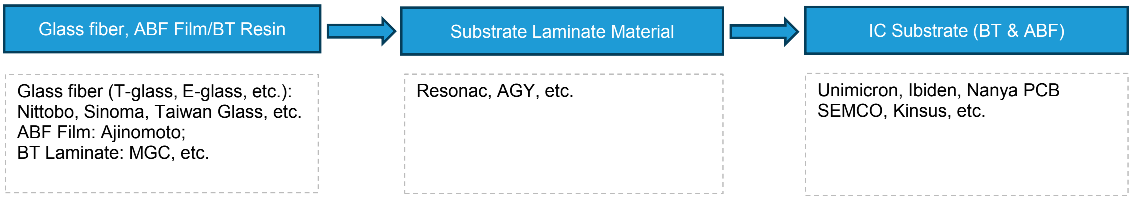

Note: Only Ajinomoto, Unimicron, and Ibiden are covered.

Source: Company reports, Bernstein analysis

EXHIBIT 34: **Nittobo's new T-glass capacity will not be available until late 2026-27**

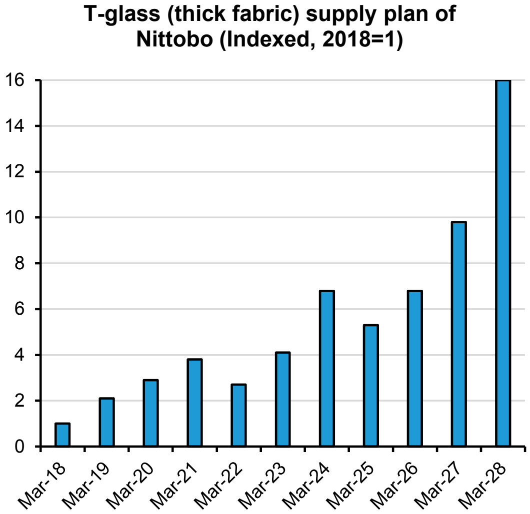

Source: Prismark data and estimates (2026-28), Nittobo, Bernstein analysis

EXHIBIT 35: **ABF is produced at two plants in Japan and the company is materially expanding capacity**

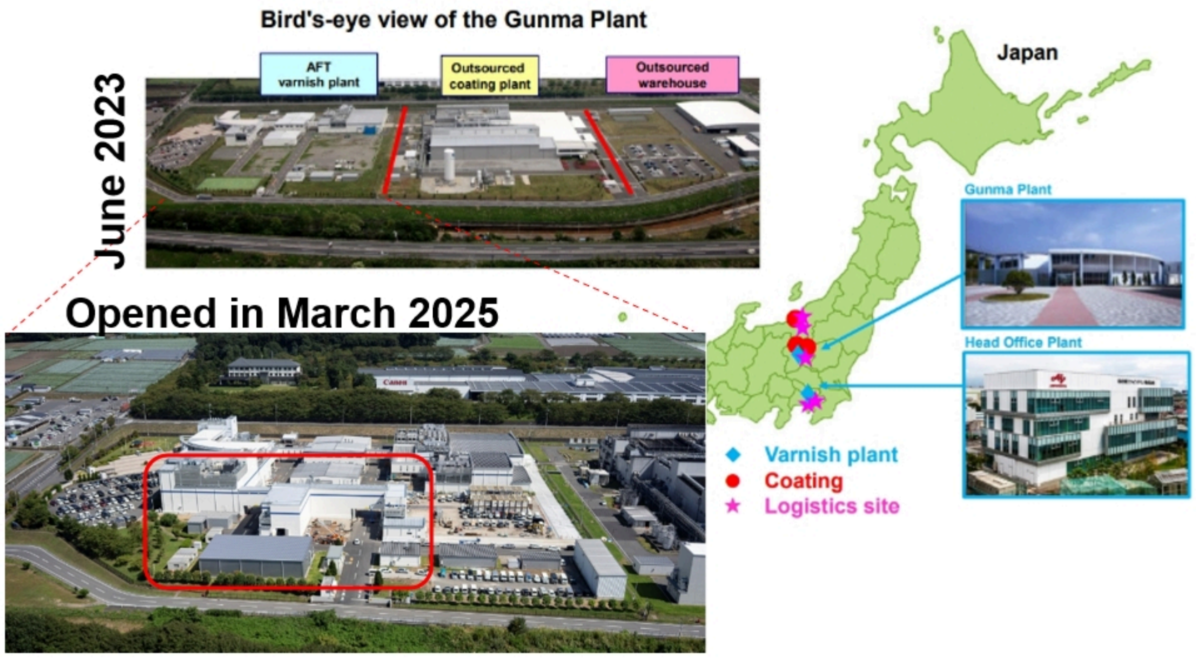

Source: Company presentation

EXHIBIT 36: **We expect ABF operational leverage to remain in the typical 1.5-1.7x range, driving margin expansion for the division**

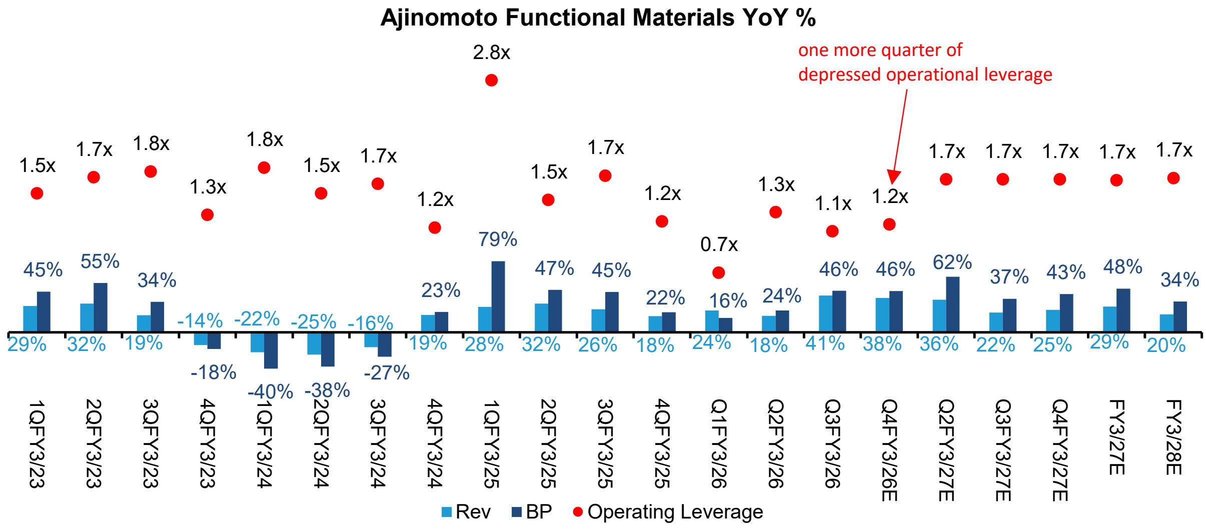

Source: Company reports, Bernstein analysis and estimates

AI PCB PRIMER: PLAYERS, POSITIONING, AND PROFITABILITY

INDUSTRY OVERVIEW

ABF substrate supply remains concentrated in Taiwan, South Korea, and Japan, while China plays a critical role in MLPCB, HDI, and materials thanks to scale, capital investment, and environmental compliance. Unlike IC substrates tied closely to chip manufacturing, PCBs are sub-components shipped to ODMs for assembly, making it less geopolitically sensitive. AI PCBs already account for 20-60% of revenue for companies such as WUS, Victory Giant, and Shengyi Electronics.

AI is fueling strong top-line growth across the PCB supply chain. NVIDIA-driven HDI demand and Amazon's PCB orders supported 50-100% revenue growth in 2025 for VGT, Shengyi Electronics, and Gold Circuit. Substrate makers such as Unimicron and Ibiden saw slower growth due to non-AI weakness, but should accelerate in 2026. Rising PCB complexity and higher AI accelerator shipments support a double-digit percentage market growth outlook. CCL suppliers such as Shengyi Tech and EMC will likely remain major AI beneficiaries.

AI-related products currently enjoy elevated margins, though upside may narrow. Most AI PCB and CCL companies saw meaningful gross margin (GM) expansion, ranging from ~1ppt (TTM) to 14ppts (VGT). EMC and WUS delivered ~20% ROIC, the highest in the supply chain. In our 2026 outlook, we are cautious on further GM expansion into 2026-27 as competition intensifies, overseas capacity ramps slowly, and depreciation rises. New product cycles — midplane MLPCB (late 2026) and backplane MLPCB (under testing) — and expanding ASIC demand should partially offset pressure. ABF substrate makers remain better positioned in the near term due to tight supply and larger body sizes for Vera Rubin chips.

Unimicron (Outperform): We expect AI to account for roughly 50% of revenue in 2026. Unimicron is likely to capture ~35% ABF market share in NVIDIA high-end GPUs in 2026, along with a 50%+ share for ASIC chips including Google TPU and AWS Trainium. Yield improvements at Taiwan HDI sites in 2Q26 should lift AI HDI/PCB revenue by ~50% in 2026 and mid-teens in 2027. The key debate remains whether Unimicron's GM can return above 30% as in the 2022 cycle. We think this is unlikely as the current ASP hikes are raw-material-driven rather than substrate-driven.

Ibiden (Outperform): We believe the outlook for Ibiden remains strong, given NVIDIA Rubin GPU ramp-up and Intel server CPU and AMD MI400 GPU. The 3Q FY26/3 miss was due to lackluster Intel chip and higher depreciation costs, but with Rubin ramping up from 4Q FY26/3, profit should start growing again.

Ajinomoto (Market-Perform): ABF film accounts for close to one-third of Ajinomoto's business profit (business profit is an accounting term in Japan) where the company has a ~95% market share, and its tight R&D relationships with key chip designers represents a solid moat. ABF generates a 54% business profit margin (BPM) for Ajinomoto, and we expect the FY3/27E 40%

+ growth rate to continue well into FY3/28E on the back of AI demand, primarily.

KEY PCB, HDI, SUBSTRATE, AND MATERIAL SUPPLIERS

Stocks of companies operating in PCB and substrate markets have been on fire, driven by rising content in AI servers and expectations of ASP hikes in 2026. Exhibit 1 summarizes the major businesses of key PCB, HDI, substrate, and material suppliers with AI exposure, and Exhibit 2 quantifies their revenue exposure. In particular:

- In IC substrate (17% of PCB market), **Ibiden** and **Unimicron** are the primary ABF substrate suppliers for NVIDIA's AI GPUs, switch chips, and CPUs. Ibiden leads in substrates for NVIDIA, particularly for the new Rubin platform; Unimicron is a key supplier for major ASIC platforms including Google TPU, AWS Trainium, and Meta MTIA. Ibiden's PCB business is largely substrate, while Unimicron has a more balanced ABF/HDI exposure. Nanya PCB (not covered) and SEMCO also have exposures to ASIC ABF substrates.

- In HDI (18% of PCB market), Unimicron dominated supply for NVIDIA's Hopper servers in 2024, but **Victory Giant** (VGT, 300476 CH, not covered) has become the largest HDI supplier for GB200 and GB300, driven by new facility investments and strong execution. VGT has secured extensive laser drilling and microvia filling/copper plating equipment from Mitsubishi (6503 JP, not covered) and other Japanese/German vendors, enabling high-precision HDI processes and faster yield ramp-up at new sites. Its HDI capacity is also more concentrated, with a 20k-sqm 6+ HDI line in Huizhou and expansion planned in Vietnam for 2027. In contrast, Unimicron's production is dispersed across multiple sites in China and Taiwan.

- For MLPCB (a part of the PCB market), suppliers include WUS, ISU, TTM, Shengyi Electronics, and Gold Circuit (all not covered). **WUS** is the primary supplier for NVIDIA's GB300 rack, followed by Victory Giant and Unimicron. For NVIDIA's HGX (8-GPU) servers and the TPU supply chain, **TTM** and **ISU** are leading suppliers. VGT also captures MLPCB demand from Google TPU, and volume shipment should start in mid-2026. For AWS Trainium, key MLPCB suppliers include **Gold Circuit** and **Shengyi Electronics**.

- In high-end material used in PCB (Exhibit 3), Doosan, EMC, and Shengyi Tech (all not covered) are major suppliers for AI CCL. Doosan is a conglomerate with diversified businesses. EMC is a pure-play CCL materials supplier. Shengyi Tech is also mainly a CCL producer, with PCB revenue contributed by its subsidiary Shengyi Electronics. Because single-material markets are smaller than downstream PCB, leading high-end materials suppliers — **Ajinomoto** (ABF film), **Mitsui Mining** (HVLP4 copper foil; 5706 JT, not covered), and **Nittobo** (T-glass; 3110 JP, not covered) — maintain significant revenue from other business segments while serving as sole suppliers or technology leaders in their respective areas. **Ajinomoto** has a ~95% market share in the supply of ABF film to the IC substrate industry.

EXHIBIT 1: **Summary of major PCB, HDI, substrate, and materials suppliers with AI exposure**

Sector		PCB				Materials			
Sub-sector		RPCB	HDI	IC Substrate	FPC+SMA	CCL	ABF Film	Copper Foil	Glass fiber
PCB companies	As % of total PCB market (2025E)	48%	19%	17%	16%				
	As % of Total PCB BOM					30%	8%	22%	6%
	Victory Giant	○	★						
	Unimicron	○	★	★	○				
	WUS Group	★							
	TTM	★							
	Gold Circuit	★							
	Ibiden	○		★	○				
	SEMCO			★					
	Shengyi Electronics	★							
	Kinsus			★					
	Nan Ya PCB	○		★					
	ISU	★							
Material company	Doosan					★			
	Shengyi Tech	○				★			
	EMC					★			
	Ajinomoto						★		
	Mitsui Mining							★	
	Nittobo								★

○ Industry exposure; ★ Major business

Note: Unimicron, Ibiden, Aginomoto are covered; all other companies are not covered.

Source: Prismark data and estimates (2025E PCB revenue), company reports, Bernstein analysis

Investment

EXHIBIT 2: **Among AI PCB suppliers, Unimicron and Ibiden are major substrate players, VGT is a major HDI supplier; among material suppliers, EMC and Shengyi Tech are pure-play PCB material suppliers while others have diversified businesses**

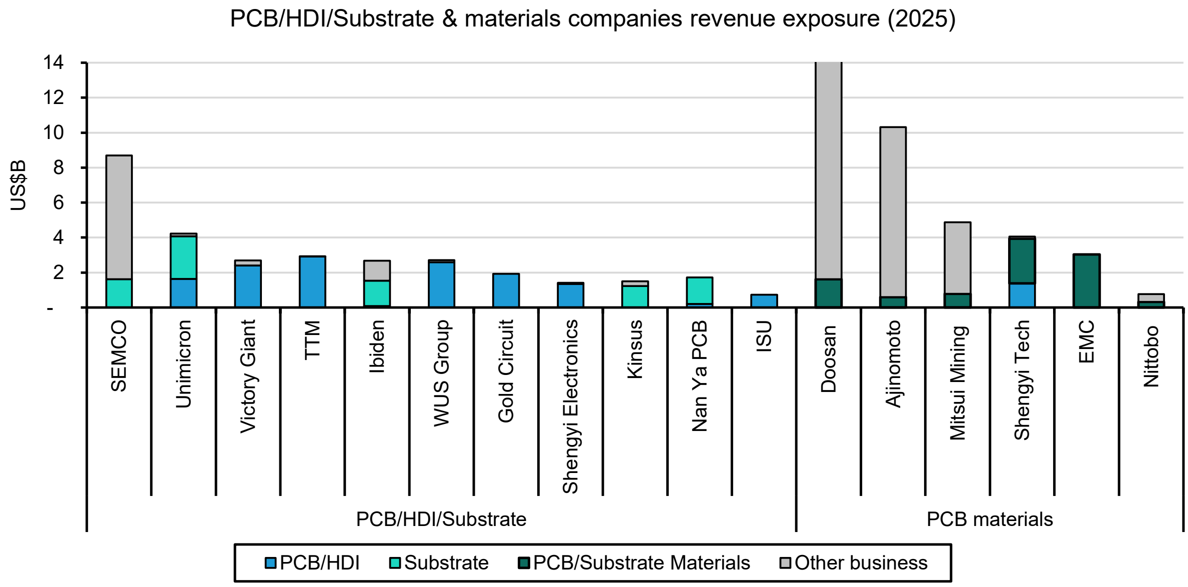

Note: The 2025 mix is estimated based on historical numbers; Victory Giant, WUS Group, Shengyi Electronics, and Shengyi Tech 2025 numbers are based on Bloomberg consensus, others are actual. Unimicron, Ibiden, Ajinomoto are covered; all other companies are not covered.

Source: Bloomberg, Bernstein analysis and estimates

EXHIBIT 3: **The major part of PCB costs is CCL, which is mostly made of copper foil, resin, and glass fiber cloth**

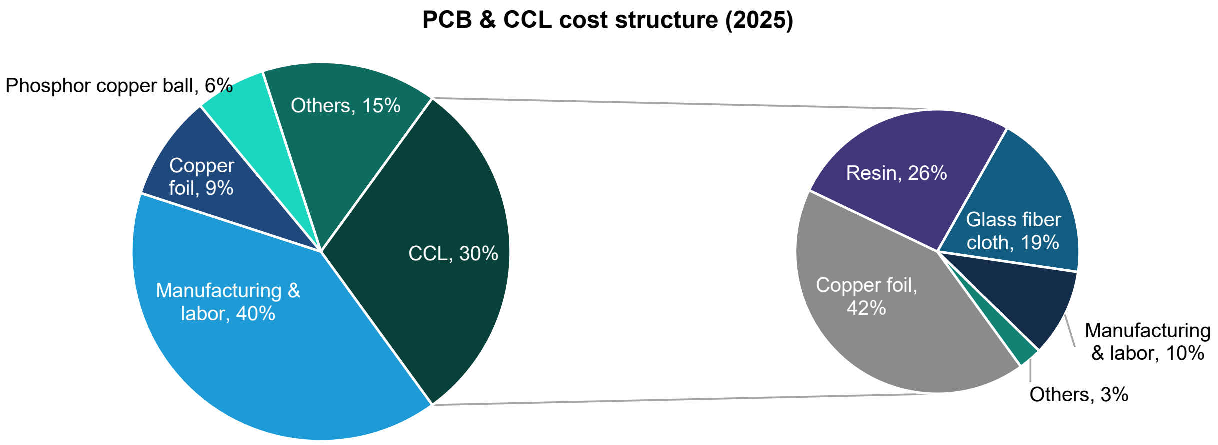

Source: CHNCI, Bernstein analysis

SERVER PRODUCTS TEND TO GENERATE STRONGER ROIC

Exhibit 4 shows major AI PCB suppliers' AI mix in their total PCB revenue; Exhibit 5 and Exhibit 6 present their ROIC versus invested capital. The PCB industry requires heavy capital investment as manufacturers must continually expand capacity to keep pace with cyclical demand. ROIC across the supply chain generally ranges from the low single digits to the low 20s.

Companies with higher exposure to the server market, such as EMC, Gold Circuit, and WUS, tend to generate stronger ROIC. VGT, with significant AI mix increase in 2025, is also expected to see meaningful ROIC improvement in 2026 based on consensus forecasts. Conversely, ABF substrate suppliers such as **Unimicron** and **Ibiden**, despite their sizable AI-server exposure, faced margin pressure due to oversupply in low- to mid-end ABF, Intel-related burdens at the high end, and headwinds in other business lines in the past. Other substrate companies — SEMCO, Nan Ya PCB, and Kinsus — had limited AI exposure (low-single-digit) as of 2025, and their ROICs are at the lower end as well. **Ajinomoto** dominates the supply of ABF film, and the material accounts for ~30% of the company's operating profit. **Nittobo** dominates T-glass supply, but the material only accounts for ~15% of company revenues.

With the ongoing shortage of high-end materials (especially T-glass), we expect Unimicron and Ibiden will be able to increase prices and margins. We model Unimicron's financials to bottom out starting from 1Q26 with ROIC back to double-digits in 2026.

EXHIBIT 4: **Major AI supply chain PCB suppliers all have over 20% AI mix**

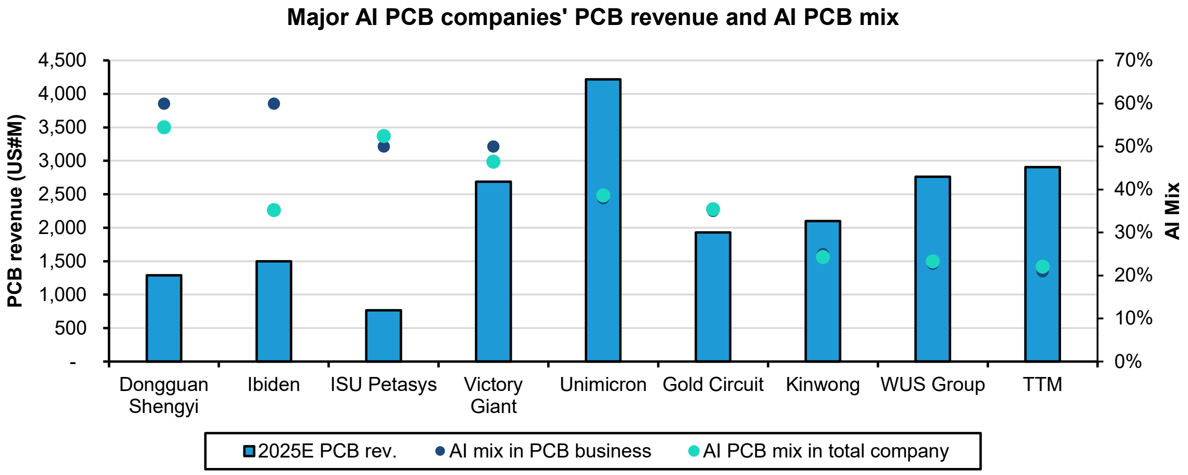

Note: 2025E PCB revenues are Prismark estimates; AI mix and Unimicron revenue estimated by Bernstein. Unimicron and Ibiden are covered; all other companies are not covered.

Source: Prismark, Bernstein analysis and estimates

EXHIBIT 5: **PCB supply chain companies with higher exposure to server market and AI would likely present better ROIC, while financials of ABF substrate suppliers are dragged by oversupply of low-to mid ABF, Intel burden for high-end ABF, and other business**

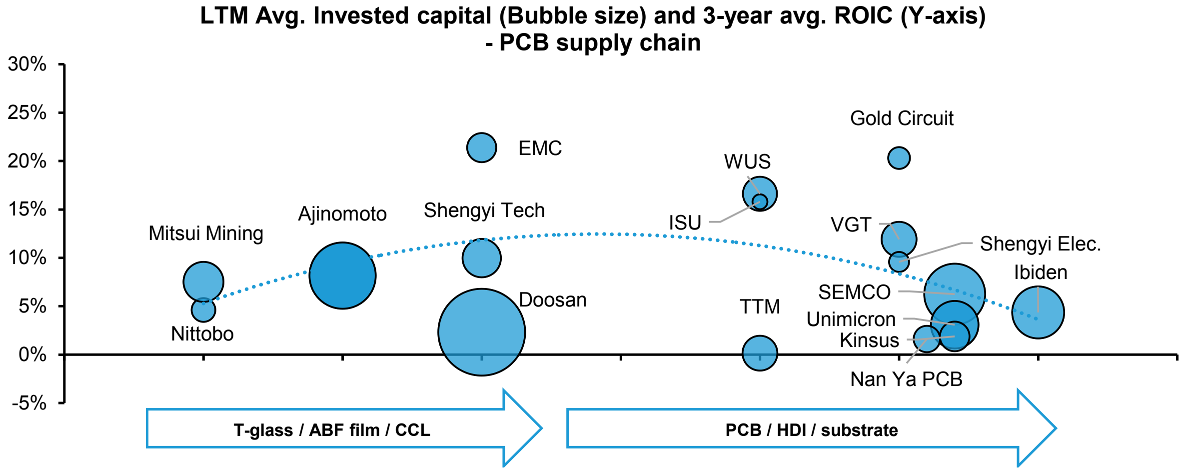

Note: Three-year ROIC refers to 2023-3Q25 average ROIC. Unimicron, Ibiden, and Ajinomoto are covered by Bernstein; all other companies are not covered.

Source: Bloomberg, Bernstein analysis

EXHIBIT 6: **Companies with notable AI mix increase such as VGT, EMC, Shengyi Electronics, and TTM also observed an increase in ROIC**

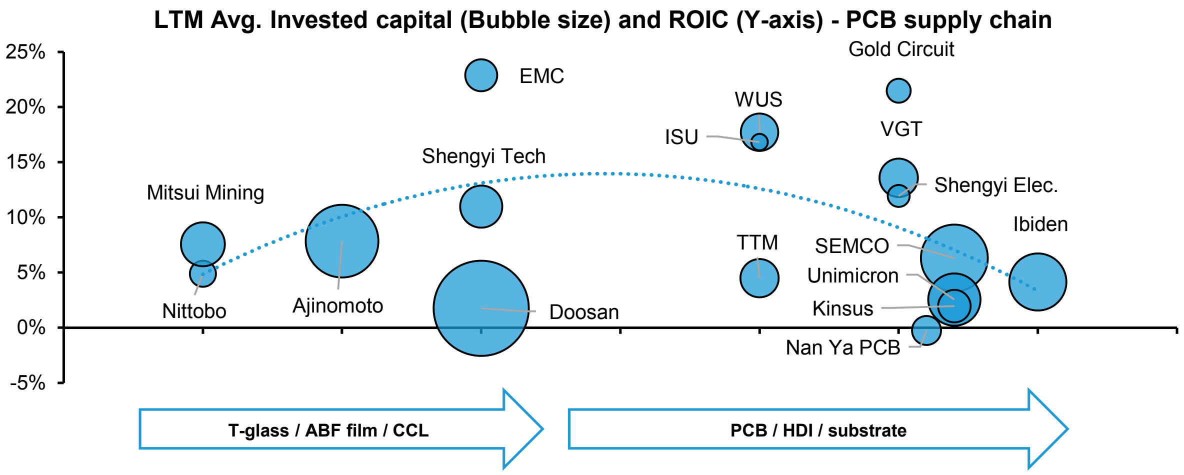

Note: LTM ROIC refers to CY4Q24-3Q25 average ROIC. Unimicron, Ibiden, and Ajinomoto are covered by Bernstein; all other companies are not covered.
Source: Bloomberg, Bernstein analysis

AI IS FUELING STRONG TOP-LINE GROWTH ACROSS PCB SUPPLY CHAIN

WUS, the key MLPCB supplier for NVIDIA, posted 40%+ revenue growth in 2025. NVIDIA-driven HDI demand and Amazon's PCB orders support 50-100% revenue growth for VGT, Shengyi Electronics, and Gold Circuit, according to consensus estimates. Substrate makers such as Unimicron and Ibiden delivered slower growth in 2025 due to weakness in non-AI segments, but both are likely to see accelerated growth in 2026. Looking ahead, rising PCB complexity and higher AI accelerator volumes should support double-digit revenue growth across the PCB market. In materials, CCL suppliers such as Shengyi Tech and EMC will likely remain major beneficiaries of the AI cycle.

NVIDIA's Rubin platform, expected to ramp in 2H26, is shifting to a cableless interconnect design, benefiting suppliers of high-end MLPCBs, HDI, and advanced materials. The Rubin **midplane** is central to this architecture, replacing copper cables for high-speed links between the switch tray and compute tray, and is likely to use a A 44-layer plated through-hole (PTH) MLPCB with M8-class CCL to reduce loss and latency. For Rubin Ultra in the Kyber architecture, a PCB **backplane** is also under consideration, potentially using 78-layer or higher MLPCBs with M9 CCL, which would further raise PCB value per server. We understand WUS, VGT, and Unimicron are potential suppliers for the midplane and backplane. Rubin Ultra's final design is not yet fixed, and backplane adoption remains under discussion at this stage (Exhibit 9 and Exhibit 10).

Ajinomoto has been capitalizing on this demand, growing its Functional Materials revenues and business profit (BP) by 26% and 46% in FY3/25, respectively, and on track to deliver 31% and 34% growth, respectively, in FY3/26 (operating leverage dampened temporarily due to major capacity expansion), and by 29% and 49%, respectively, in FY3/27, by our estimates.

The growth of upstream materials suppliers, such as Mitsui's copper foil and Nittobo's T-glass/low-Dk glass are capped by their capacity expansion pace; however, their overall revenue growth is also dragged by moderate growth from other businesses (Exhibit 7 and Exhibit 8).

EXHIBIT 7: **Consensus expects companies in PCB supply chain to see revenue increase in 2025-27E...**

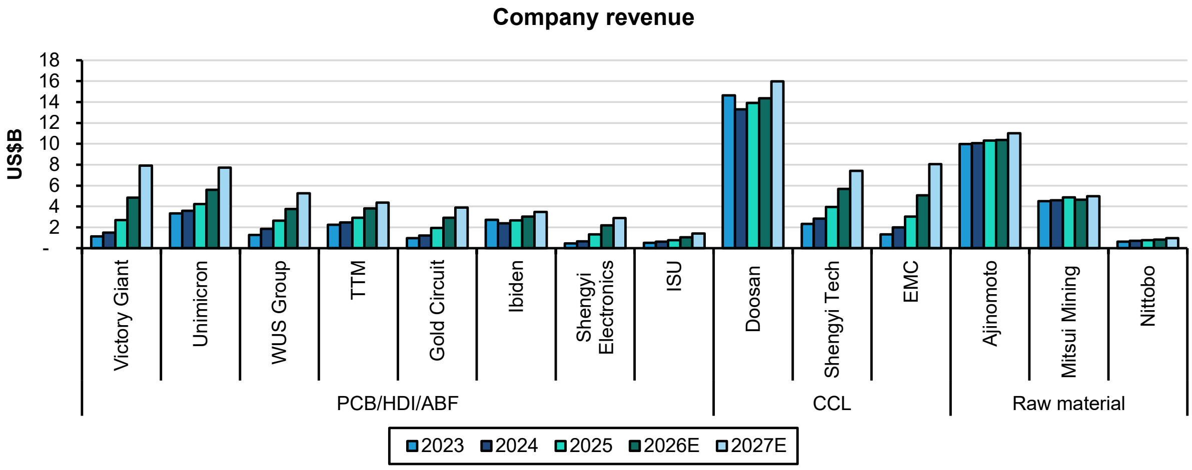

Note: All forecasts are Bloomberg consensus, adjusted to calendar year. Only Unimicron, Ibiden, and Ajinomoto are covered.

Source: Bloomberg, Bernstein analysis

EXHIBIT 8: **... and forecast companies with higher AI and server exposure to experience more significient revenue growth in 2025-27E**

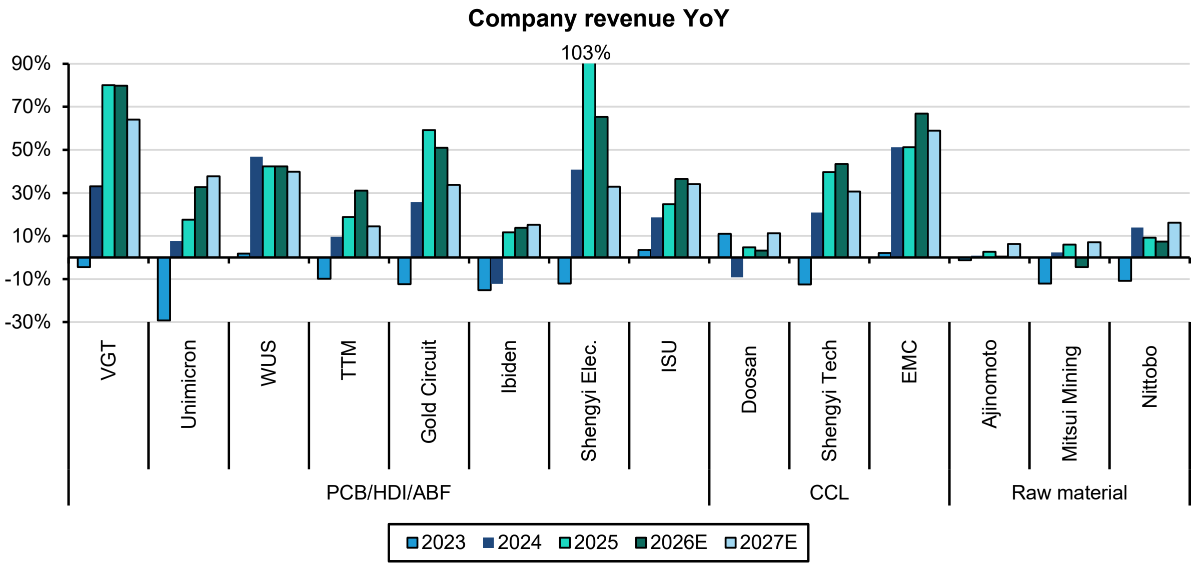

Note: All forecasts are Bloomberg consensus, adjusted to calendar year. Only Unimicron, Ibiden, and Ajinomoto are covered.

Source: Bloomberg, Bernstein analysis

EXHIBIT 9: **Midplane is an MLPCB that replaces cables to connect components in the compute tray**

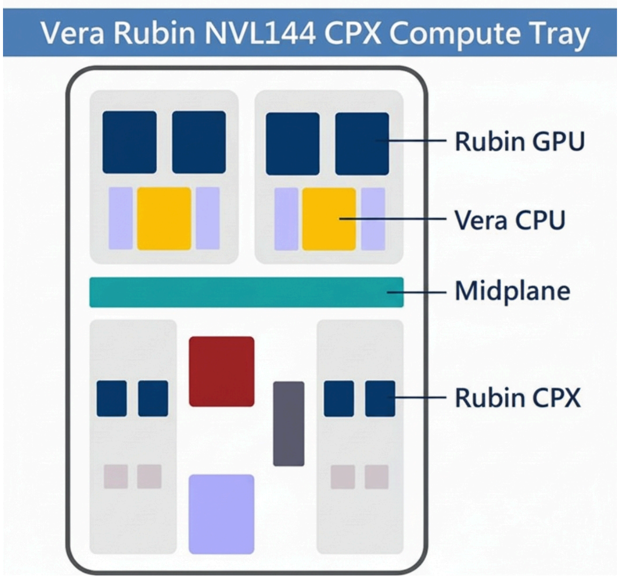

Source: SemiVision, Bernstein analysis

EXHIBIT 10: **Potential design for Kyber: the backplane is an MLPCB that connects multiple compute blades and switch trays**

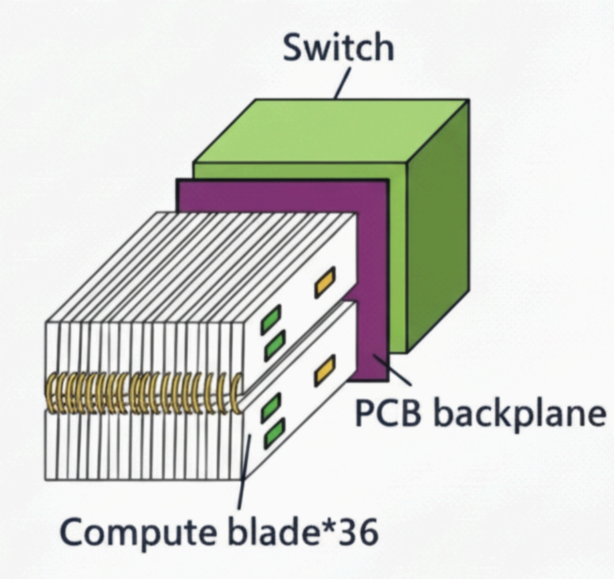

Source: SemiVision, Bernstein analysis

PCB MARKET IS CAPEX-INTENSIVE

The overall PCB market is capex-intensive as companies need to expand capacities to meet changing demand, often cyclical in nature. To capture the booming AI PCB demand, many companies have announced new rounds of fund-raising for capacity expansion plans in 2026-27E. **VGT** stands out, with consensus forecasting a 7.8x capex increase in 2026E versus 2024, indicating major AI-driven PCB capacity build-outs both in China (Huizhou HQ) and overseas (Thailand and Vietnam). VGT is currently ramping up its Thailand site; its Vietnam site just started construction with the first HDI phase targeting to operate in 2H26. **WUS** and **Shengyi Electronics** are also expanding capacity in China and Southeast Asia to support growth and reduce geopolitical supply chain risk. WUS announced plans in January 2026 to invest ~$300Mn to develop mSAP and CoWoP-related technology platforms, with the potential to increase investment to $500Mn, depending on the first phase's progress. Shengyi Electronics announced in November 2025 plans to build new MLPCB/HDI capacity, commencing in 2026-27. Note that these new investments may not be fully reflected in consensus numbers yet. Outside Asia, the main new capacity comes from **TTM's** Syracuse Phase 1 plant in New York, which is expected to start volume production in 2H26 and will primarily serve the aerospace & defense market, leaving most AI PCB capacity concentrated in Asia.

On the ABF substrate side, **Unimicron** has maintained $700Mn-$1Bn capex per year since 2021 (~NT$26Bn on average), focused on HDI capacity conversions and ABF substrate capacity. Although the company guides roughly flat capex for 2026 (about NT$20Bn) versus 2025, there is potential upside, given tight supply at its Yangmei (Intel + ASICs) and Guangfu sites (NVIDIA). The first phase of its Guangfu site is fully ramped, with the second phase scheduled to be fully ramped by end-2026; in November 2025, Unimicron announced a NT$145Mn (~$5Mn) plant lease in Taiwan to prepare for further expansion.

Ibiden, in February 2026, announced plans to raise capex in 2026-28 to ¥500Bn, up from consensus of ¥376Bn. The capex is to be spent on two major projects: (1) ¥280Bn on Ono fab, which we believe is for a major GPU customer; and (2) ¥220Bn on Gama fab, which we believe is dedicated to EMIB-T. Based on our understanding, both fabs will have the most advanced technology and high ASP and, hence, a high return on investment. With these investments, we estimate the total capacity of Ibiden to almost double. Based on the rule of thumb of $1 revenue for $1 capex, and assuming 20-25% OPM, the investment may bring ¥100-¥125Bn OP annually after full ramp up, almost tripling the current OP.

Materials suppliers are also scaling up, though shortages in key upstream inputs such as T-glass are likely to persist. Shengyi Tech plans to invest about RMB1.4Bn (~$196Mn) in a new CCL plant in Thailand. EMC also bought new land in January 2026 to expand, targeting to increase overall capacity by 30% after completion. Nittobo is adding ¥15Bn of T-glass capacity in Fukushima, but this will not come online until early 2027, implying tight T-glass supply through 2026. Mitsui Kinzoku (not covered) is increasing investment in its engineered materials segment (including copper foil, engineered powders, and rare materials) and in new business creation, particularly solid electrolytes for all-solid-state batteries and High-Reliability Density Product (HRDP) (Exhibit 11 to Exhibit 14).

In mid 2026, **Ajinomoto** will likely roughly double the scale of its Gunma ABF plant in Japan, adding around 50% to its overall capacity after a total ¥95Bn capex. The company is currently only running one production shift out of a possible two, and management estimates that it has sufficient capacity to take it beyond 2030.

EXHIBIT 11: **The overall PCB market is capex-intensive as companies need to expand capacities to meet changing demand, often cyclical in nature**

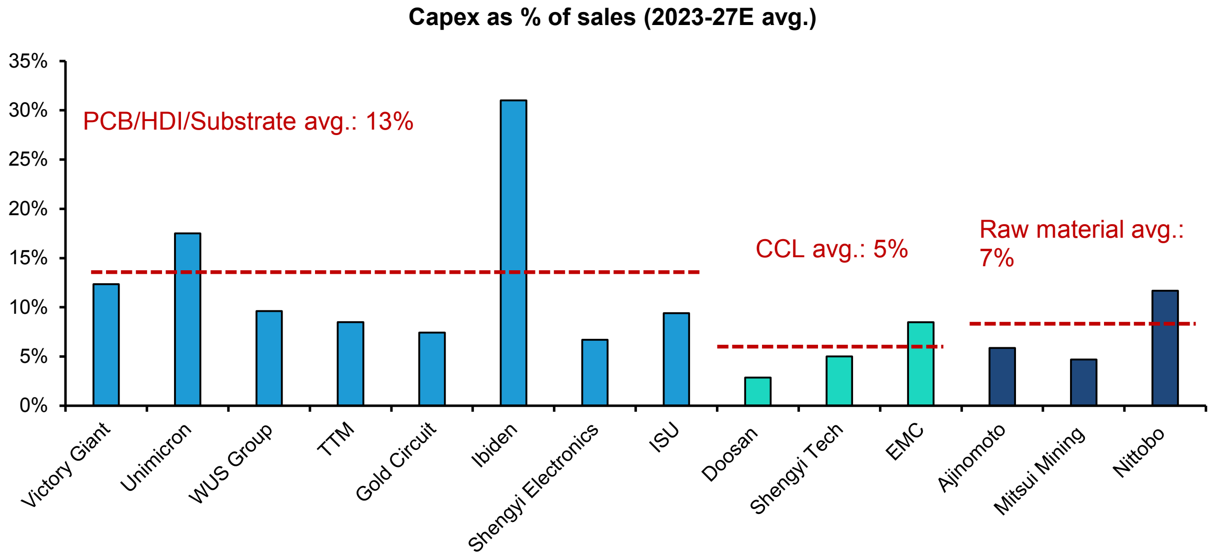

Note: All forecasts are from Bloomberg consensus, adjusted to calendar year. Only Unimicron, Ibiden, and Ajinomoto are covered.
Source: Bloomberg, Bernstein analysis

EXHIBIT 12: **Chinese players are aggressive in investing new capacity to capture emerging AI opportunities**

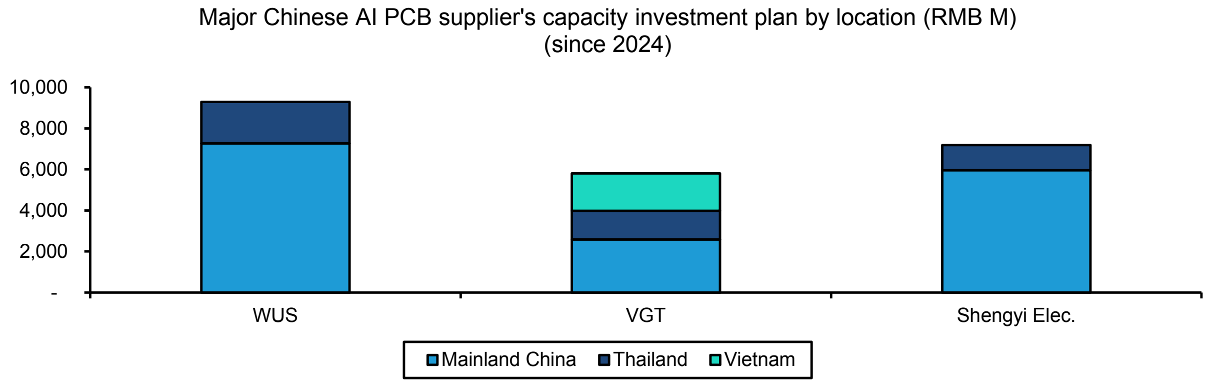

Source: Company reports, Bernstein analysis

EXHIBIT 13: **ABF substrate companies' expansion plan**

Company	Site	Product	Capacity	Production Start (Calendar Year)
Unimicron	Yangmei site, TW	ABF (Server)	Yangmei and Guangku site together is c.35% of total Unimicron ABF	2021 (fully ramp in 2024)
	Guangfu site, TW	ABF (Server)		End-2024 (1st phase fully ramp in 2025, 2nd phase starts 2026)
Ibiden	Ogaki	ABF	Will almost double ibiden's capacity with these investments	Expansion in 2023
	Ono Plant, JP	ABF (Server)		2H25
	Gama Plant, JP	ABF (Server)		2027
Nanya	Shulin, TW (1st Phase)	ABF	Total ABF production capacity: 50M Units/M by End-2023	3Q22
	Shulin, TW (2nd Phase)	ABF (HPC)		3Q23
	Kunshan, CN (2nd Phase)	ABF		3Q22
	Jinxing, TW	ABF		Upgrade in 2Q23-1Q24
Shinko	Chikuma Plant (New)	Flip Chip	Increase capacity by c.50% by 1H25, compared to 2H21	4Q24-1Q25
	Kohoku Plant	Flip Chip		By 1Q24
	Wakaho Plant	Flip Chip		By 1Q24
SEMCO	Vietnam Plant	Flip Chip		Expansion in 2022
	TBD	Flip Chip		Under consideration

Source: Company reports, Bernstein analysis

EXHIBIT 14: **Multiple PCB companies have adjusted up their capex in 2026-27E to meet booming AI demand**

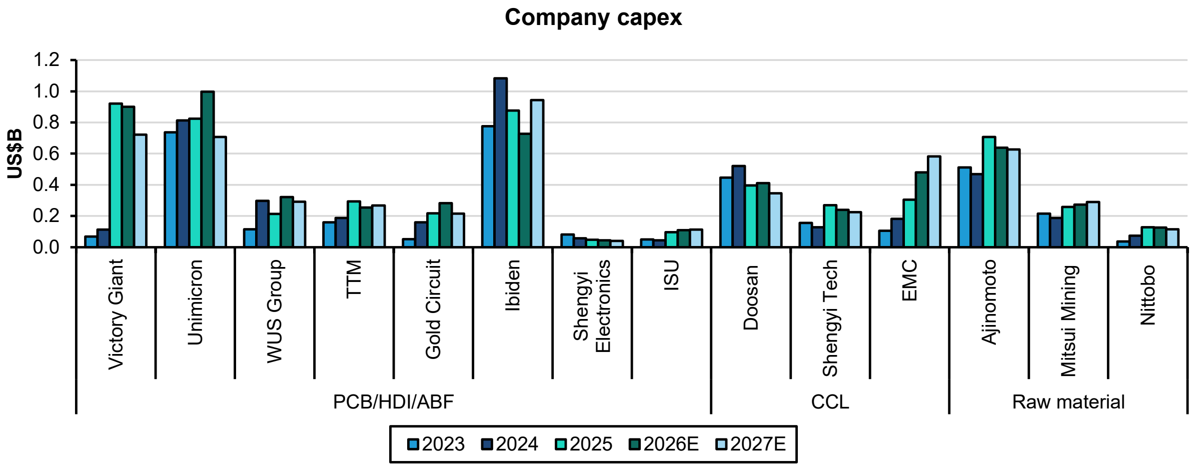

Note: All forecasts are Bloomberg consensus, adjusted to calendar year; Victory Giant, WUS Group, Shengyi Electronics, and Shengyi Tech's 2025 numbers are based on Bloomberg consensus, others are actual.
Source: Bloomberg consensus estimates, Bernstein analysis

AI-RELATED PRODUCTS TEND TO CARRY HIGH MARGINS

AI-related products currently enjoy elevated margins, though upside may be limited from here. Most AI PCB and CCL companies saw meaningful GM expansion in 2025— ranging from ~1ppt (TTM) to 14ppts (VGT). In our 2026 outlook, we are cautious on further GM gains for leading AI PCB players as competition intensifies from late-2026 into 2027. Companies are adding overseas capacity, and slower yield ramp-ups plus rising depreciation will likely weigh on margins. New product cycles — midplane MLPCB (late 2026) and backplane MLPCB (late 2027) — could partially offset competitive pressure. ABF substrate makers (Unimicron and Ibiden) are better positioned in the near term, thanks to tight supply and larger body sizes used in Vera Rubin chips (Exhibit 15 and Exhibit 16).

Ajinomoto typically generates a 1.5-1.7x operational leverage ratio in its Functional Materials division with strong pricing realized on the back of specification upgrades tied to new generation IC chip requirements. In FY3/26, the operational leverage was temporarily depressed due to increased depreciation and fixed cost relating to its Gunma plant expansion, but we expect operational leverage return to the historical range in FY3/27 on the back of Rubin-related price increases, primarily.

EXHIBIT 15: **Most AI PCB and CCL companies saw meaningful expansion in 2025, and consensus expects most of them to see sequential growth in 2026-27E...**

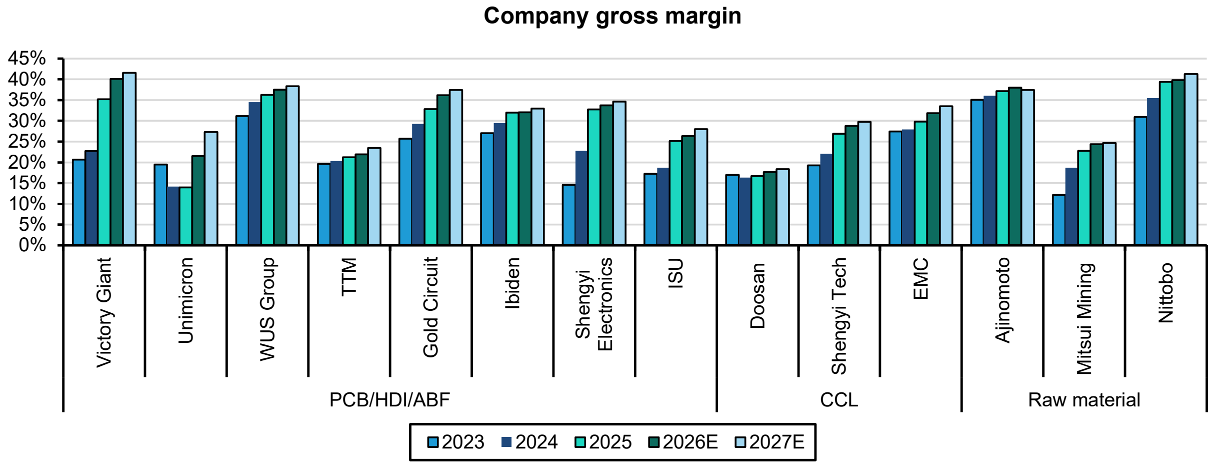

Note: All forecasts are Bloomberg consensus, adjusted to calendar year; Victory Giant, WUS Group, Shengyi Electronics, and Shengyi Tech 2025 numbers are based on Bloomberg consensus, others are actual.
Source: Bloomberg consensus estimates, Bernstein analysis

EXHIBIT 16: **... similarly for operating margin**

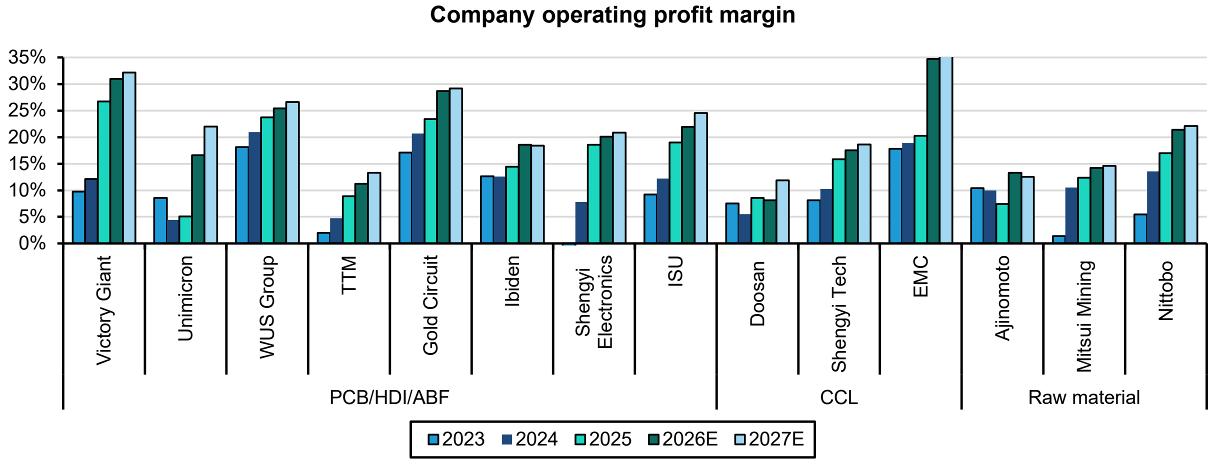

Note: All forecasts are from Bloomberg consensus, adjusted to calendar year; Victory Giant, WUS Group, Shengyi Electronics, and Shengyi Tech 2025 numbers are based on Bloomberg consensus, others are actural.
Source: Bloomberg consensus estimates, Bernstein analysis

UNIMICRON INVESTMENT HIGHLIGHTS

We expect Unimicron can grow 2027 revenue and EPS significantly on a favorable product mix and the slow ramp-up of Tier 2 T-glass suppliers. We project 2027 revenue growth of 23% YoY, driven by robust demand for AI-related ABF substrates and HDI. We estimate ABF substrate ASP to rise 5-7% every quarter in 2026, supported by T-glass constraints through 3Q26 and the 4Q26 ramp of Vera Rubin chips (about two quarters behind Ibiden). Second-source suppliers should have limited supply for AI chips in 2026.

We model AI to represent half of Unimicron's revenue in 2026. ABF capacity at Yangmei and Guangfu Phase 1 should reach full utilization in 1H26, with KF Phase 2 ramping in 2027 and Yangmei Phase 2 in 2028. We expect AI to account for roughly 50% of revenue in 2026. Unimicron is likely to capture ~5% ABF market share in NVIDIA high-end GPUs in 2026, along with a 50%+ share for ASIC chips, including Google TPU and AWS Trainium. Yield improvements at Taiwan HDI sites in 2Q26 should lift AI HDI/PCB revenue by ~50% in 2026 and mid-teens percentage in 2027.

We forecast a 25% revenue CAGR for 2025-27E. EPS is expected to nearly triple to NT $11.8 in 2026 from a low base, and rise another 60%+ in 2027 on continued margin recovery. We model1Q26 gross margin to improve to18.4%, and to reach ~25% by 4Q26, driven by HDI yield improvement and favorable substrates pricing/mix. AI ABF GM will likely improve to 35% and HDI GM will likely recover to low-20s% by the end of 2026.

The key debate remains whether Unimicron's GM can return above 30% as in the 2022 cycle. We think this is unlikely as the current ASP hikes are raw-material driven rather than substrate-driven. Given a better visibility for Unimicron's margin recovery in the next 12 months, we apply a target P/E of 32x P/E to 2027 EPS of NT$19.1, setting the target price at NT$610.00; Outperform (Exhibit 17 toExhibit 20).

EXHIBIT 17: **We expect Unimicron's substrate ASP to rise in 2026-27, driven by tight supply and favorable product mix**

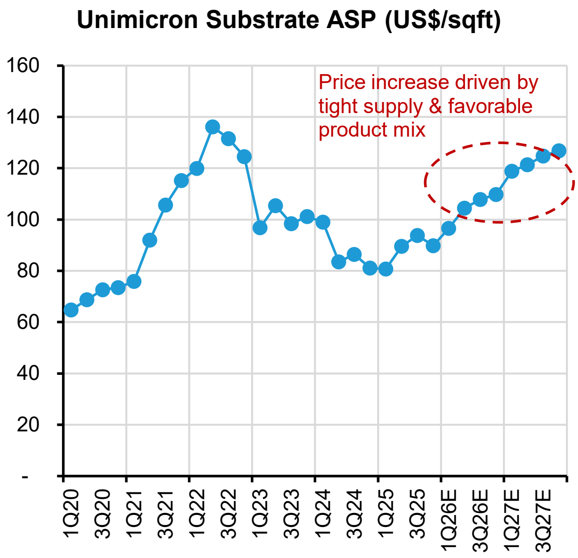

Source: Company disclosures, Bernstein analysis and estimates

EXHIBIT 18: **We expect 63% of Unimircon's 2026E revenue to come from IC substate (ABF+BT)**

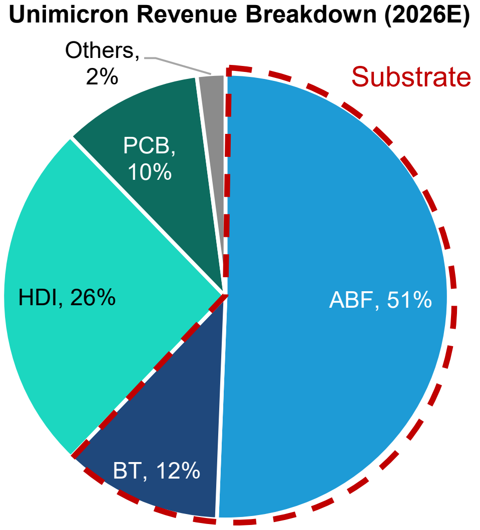

Source: Company reports, Bernstein analysis and estimates

EXHIBIT 19: **We expect substrate ASP to recover amid T-glass shortage and ramp-up of Vera Rubin chips**

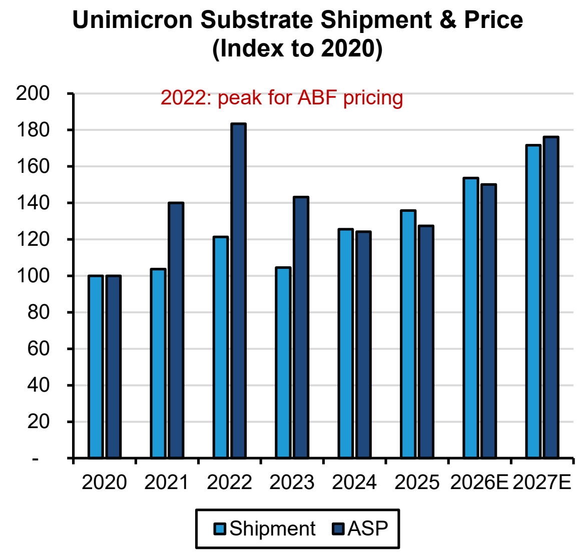

Source: Company disclosures, Bernstein analysis and estimates

EXHIBIT 20: **We model EPS to have a strong rebound in the next two years**

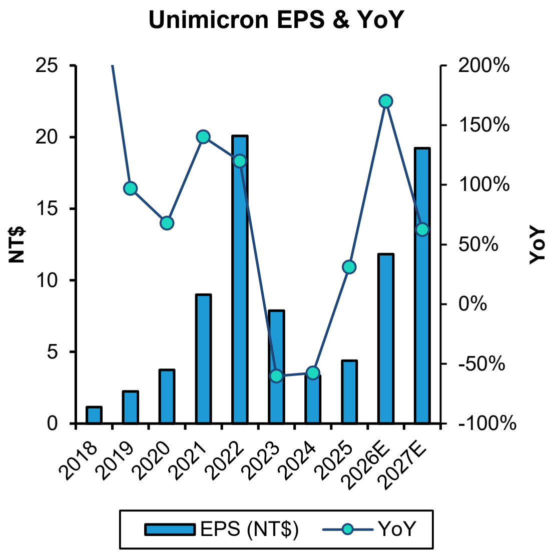

Source: Company disclosures, Bernstein analysis and estimates

IBIDEN INVESTMENT HIGHLIGHTS

Focused on IC packages, Ibiden is the biggest player in high-end ABF substrates. Its customers include Intel, AMD, and NVIDIA (all three covered). Ibiden dominates the market with an 80%+ market share with key customer NVIDIA. In all, 50% of Ibiden's revenue and 56% of its operating profits in FY25/3 came from flip chip (FC) packages (mainly ABF substrates) (Exhibit 21 and Exhibit 22).

NVIDIA's Rubin substrate migration is approaching, with ABF content dollar value doubling

NVIDIA's main AI GPU is migrating from Blackwell to Rubin, which doubles substrate content. The substrate area is projected to grow from 4,785mm^2 to 6,699mm^2, increasing by 40%. The number of layers is also projected to increase from 14 to 18. Multiplying them yields a total area increase of 1.8x. With the increased complexity and, subsequently, lower yield, we believe the ASP will also increase, resulting in total dollar value for Rubin substrate at ~2x that of Blackwell (Exhibit 23).

The transition was on track for the end of 2025. Despite some industry chatter on a potential delay of Rubin, we saw substrate manufacturing starting by the December 2025 quarter, with significant volume expected in the March 2026 quarter.

The next generation, Rubin Ultra, will likely have a total substrate area (chip area x number of layers) 2x as big as Rubin, and the value will likely be more than 2x (Exhibit 24). This would allow substrate suppliers to still benefit from the substantial value growth, even if Rubin Ultra's volume is expected to be lower than Rubin's.

We believe the total market for NVIDIA's GPU substrate will grow by 127% in FY26/3 to $898Mn and will continue to grow at a 44% CAGR in the next two years to reach $1.8Bn, driven by the migration (Exhibit 25).

EXHIBIT 21: **Ibiden FY25/3 revenue breakdown**

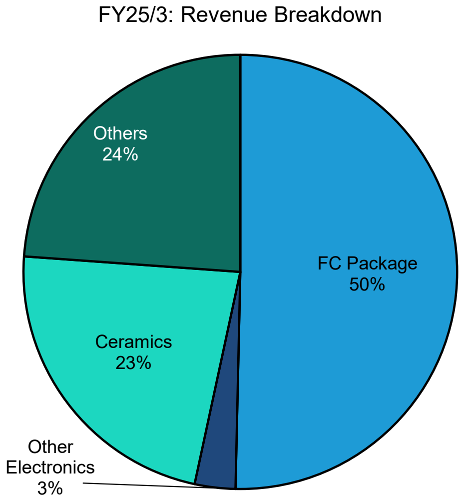

Source: Company reports, Bernstein analysis

EXHIBIT 22: **Ibiden FY25/3 operating profit breakdown**

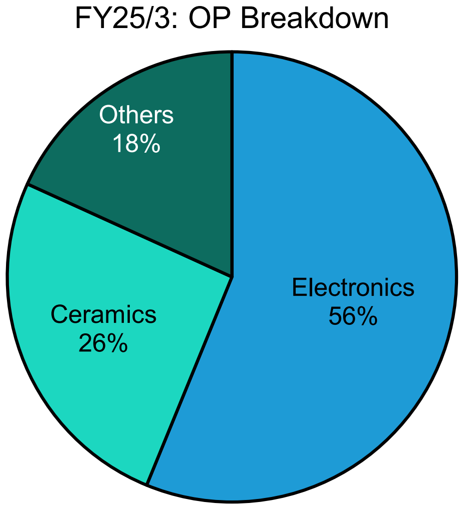

Source: Company reports, Bernstein analysis

EXHIBIT 23: **We estimate total substrate area to grow by 80-100% in each generation for NVIDIA AI GPUs**

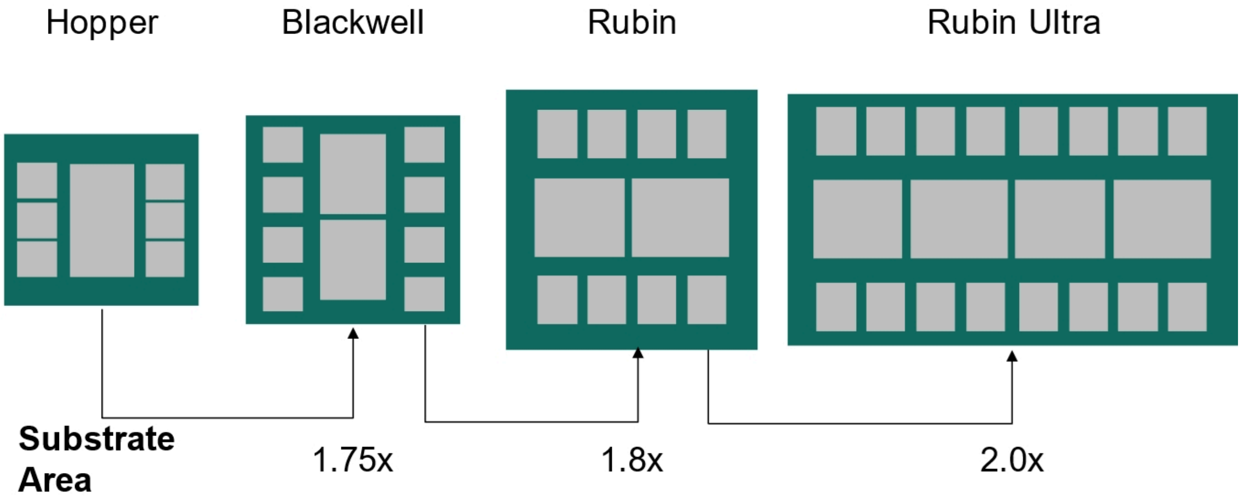

Source: NVIDIA, Bernstein analysis and estimates (substrate area growth rate)

EXHIBIT 24: **We expect significant ASP increase for substrates as generation migrates**

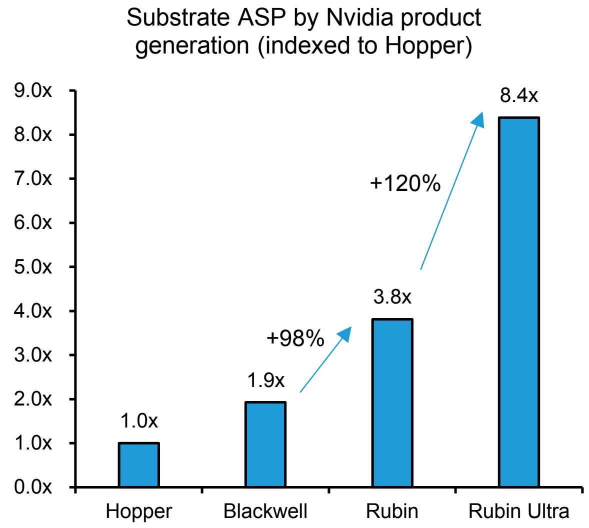

Source: Bernstein analysis and estimates

EXHIBIT 25: **We expect ABF substrate TAM for NVIDIA AI GPUs to continue strong growth of ~40% over the next two years**

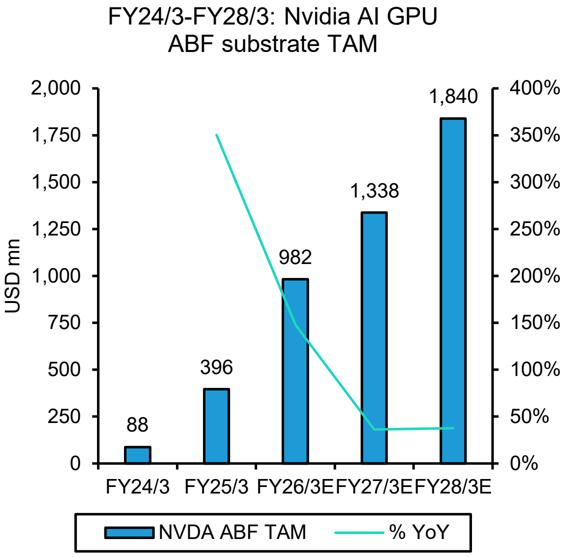

Source: TrendForce, Bernstein analysis and estimates

EMIB-T adoption as additional upside to Ibiden

Intel has offered EMIB-T technology as an alternative to TSMC's CoWoS, which is currently the de facto packaging method for AI chips. Reportedly[1], Google-MediaTek is considering it for its 2027 TPU v9 (some also call it v8E), and Meta is considering it for its MTIA accelerators (link). If this happens, we could see AI chips made using TSMC front-end and Intel EMIB-T back-end packaging. On the recent Intel earnings call in early 2026, management said advanced packaging revenue opportunities could be "hundreds of millions of dollars" to even "billions."

Intel's EMIB-T technology is based on EMIB packaging, which has been used by Intel internally for years, but modified to have TSV in the substrate and silicon bridges inserted in the substrate for AI GPU/ASIC packaging. EMIB-T is better positioned to support a larger reticle size: while CoWoS-S can support ~3.3× and CoWoS-L can extend to 5.5x and 9.5x later; Intel claimed EMIB already supported 6× in 2024 and aims to scale that to 8-12× by 2026-27. By eliminating the unused area of a round wafer as a production carrier, EMIB promises a more cost-effective solution for AI customers requiring very large packages. Another benefit is that EMIB-T packaging can be done by Intel in the US, pairing up with TSMC's front-end fabs in the US to keep the entire production in the US. However, the main weakness, in our view, is the lack of a proven track record, and possibly lower production yield due to the difficulty of embedding silicon bridges in the substrate, as two different materials are difficult of integrate. We believe MediaTek, as well as other customers such as Broadcom, is assessing Intel's EMIB in parallel with TSMC's CoWoS for possible production in late 2027 and more in 2028 (Exhibit 26).

---

[1] https://www.trendforce.com/news/2025/11/24/news-intels-emib-reportedly-gains-traction-with-ai-asic-smartphone-clients-could-package-tsmc-dies/

As indicated in Global Semis: Can Intel challenge TSMC with EMIB-T? And who benefits in the supply chain?, we believe the best way to invest in the EMIB-T opportunity might be through Ibiden. With EMIB-T, the complexity of packaging shifts from interposer to substrate and, hence, adds value to Ibiden in higher ASP (more than 50+% higher) and margins. We expect EMIB-T substrate value to rise to ~$300 or more, significantly higher versus Blackwell substrate ($80-$100) or Rubin ($180-$200) (Exhibit 27). On our estimates, every 1 million chips shifted from CoWoS (TSMC) to EMIB-T (Intel) means ~8% additional revenue to Ibiden in FY28/3E and more than 10% OP, making it a better candidate to invest in for the shift, in our view (Exhibit 28). Ibiden in 1Q26 raised capex over the next three years to ¥500Bn, of which ¥220Bn will be on Gama fab, which we believe is dedicated to EMIB-T.

EXHIBIT 26: **Process comparison: CoWoS versus EMIB**

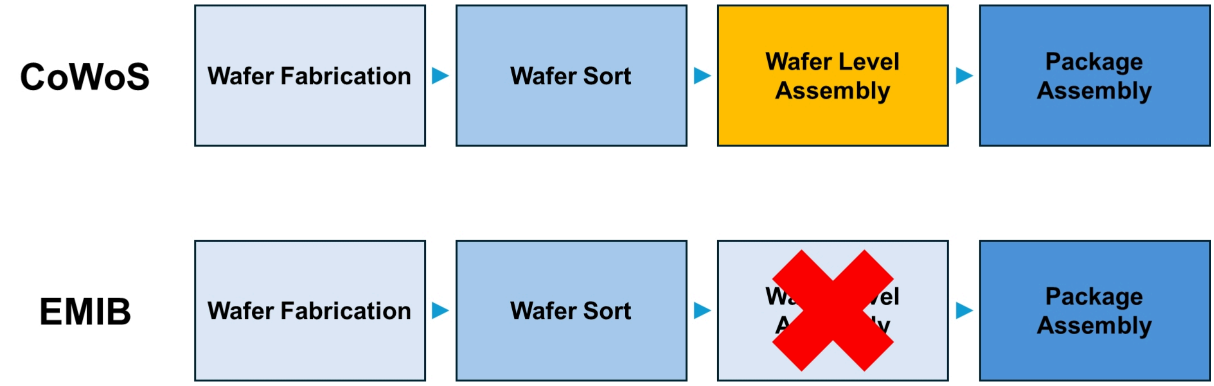

Source: Intel, Bernstein analysis

EXHIBIT 27: **AI GPUs' BOM costs on substrates and advanced packaging process comparison**

BOM Costs	Blackwell	Rubin	EMIB-T (Rubin equivalent)
Substrate (revenue to substrate makers)	~$90-100	~$180-200	~$300?
CoWoS/EMIB Process (revenue to TSMC/Intel)	~$530	~$920	Lower than CoWoS
Total Costs	~$620-630	~$1,100	Lower

Source: Company reports, Bernstein analysis and estimates (all data)

EXHIBIT 28: **Financial impact: Potential impact of every one million chip shifted from CoWoS (TSMC) to EMIB-T (Intel)**

Financial Impacts	TSMC	Intel	Ibiden
Revenue impact	-1mn x $920	A few hundred $mns	+1mn x $300
As % of 2027 revenue	-0.5%	~1-2%	7.7%

Note: FY28/3E revenue for Ibiden.
Source: Company reports, Bernstein analysis and estimates

Ibiden's recently raised capex indicates customer commitment

Ibiden raised capex for 2026-28 to ¥500Bn, up from consensus of ¥376Bn (see Ibiden 3Q FY26/3 a miss as expected, yet growth story remains intact). The capex is to be spent on two major projects: (1) ¥280Bn on Ono fab, which we believe is for a major GPU customer; and (2) ¥220Bn on Gama fab, which we believe is dedicated to EMIB-T. Based on our understanding, both fabs will have the most advanced technology and high ASP and, hence, high return on investment. With these investments, we estimate the total capacity of Ibiden to almost double. Based on the rule of thumb of $1 revenue for $1 capex, and assuming 20-25% OPM, the investment may bring ¥100-¥125Bn OP annually after full ramp-up, almost tripling the current OP (Exhibit 29).

EXHIBIT 29: **Ibiden's capex plan versus Bernstein and consensus estimates**

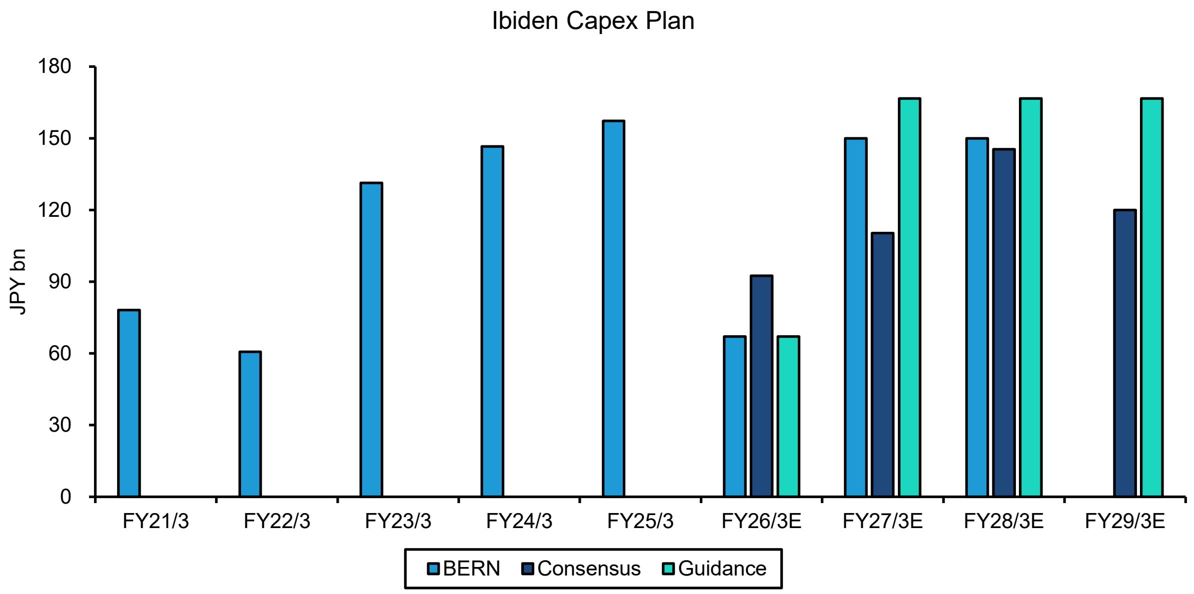

Note: Ibiden announced JPY500Bn capex over FY27/3E-FY29/3E, assuming spent evenly. Bernstein FY21/3-FY25/3 represent historical data.
Source: Bloomberg, company disclosures, Bernstein analysis and estimates

We believe the outlook for Ibiden remains strong, given ramp-up of NVIDIA Rubin GPU, Intel server CPU and AMD MI400 GPU. The 3Q FY26/3 miss was due to lackluster Intel chip shipment and higher depreciation costs, but with Rubin ramping up from 4Q FY26/3, profit should start growing again (Exhibit 30).

We envision margin expansion as larger revenue scale can dilute the negative impact from high depreciation, along with better product mix (more NVIDIA and more Intel servers versus PCs) (Exhibit 31 and Exhibit 32).

We now see a 48% EPS growth CAGR over FY26/3-FY28/3, making the current valuation of 30x Q5-Q8 EPS undemanding (implies 0.6 PEG).

Based on a 30x target P/E on Q5-8 EPS, our target price is set at ¥9,200.00. We rate Ibiden Outperform.

EXHIBIT 30: **FY21/3-FY28/3E: Ibiden revenue**

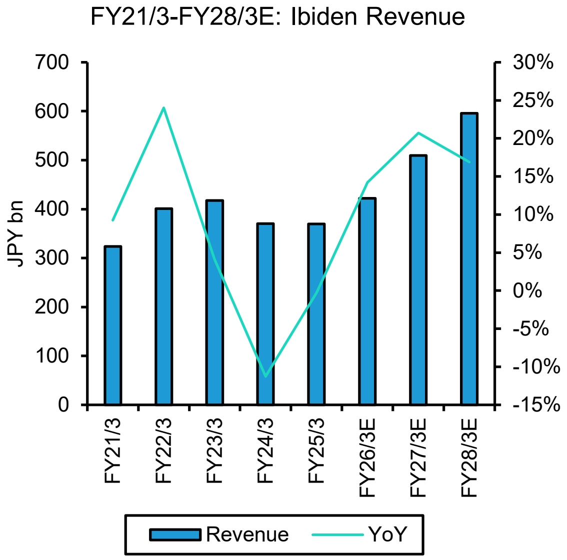

Source: Company disclosures, Bernstein analysis and estimates

EXHIBIT 31: **FY21/3-FY28/3E: Ibiden gross profit**

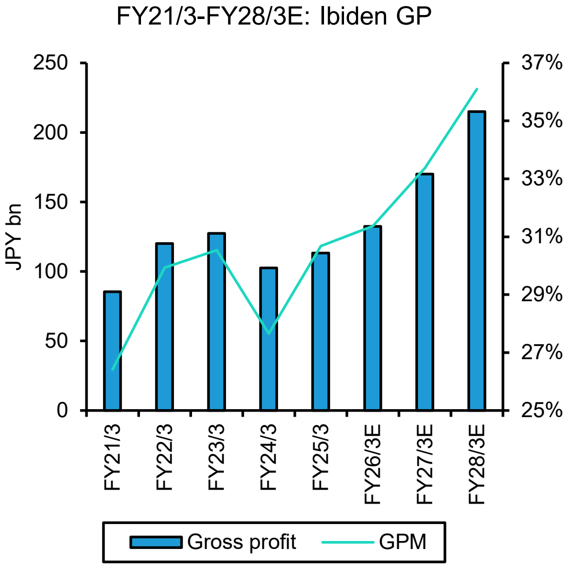

Source: Company disclosures, Bernstein analysis and estimates

EXHIBIT 32: **FY21/3-FY28/3E: Ibiden gross profit**

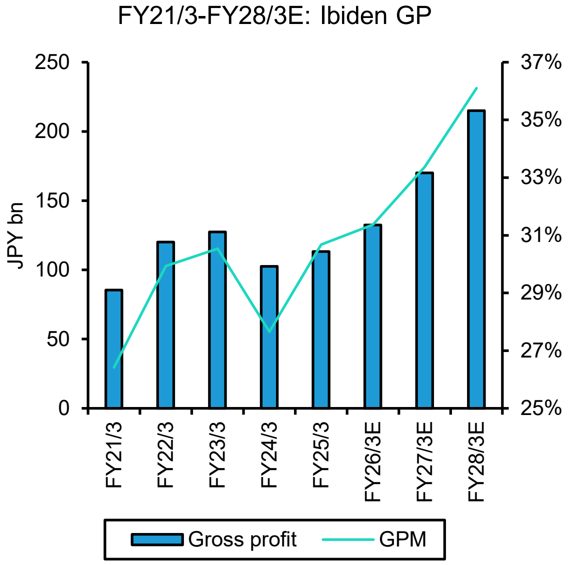

Source: Company disclosures, Bernstein analysis and estimates

AJINOMOTO INVESTMENT HIGHLIGHTS

Ajinomoto is a seasoning and food company that developed the thermal barrier ABF product in the late 1990s as part of its MSG R&D process. Its Functional Materials division (largely all ABF) is its second-largest profit center now, accounting for just under one-third of business profit. The division generates by far the highest business profit margin (BPM) in the group (54%), and is set to drive ~60% of the company's business profit growth through FY3/28 (Exhibit 33 and Exhibit 34).

Ajinomoto generates ~47% of its business profit from Asia Seasonings & Foods, where the company earns a low-20s% BPM, and where we see a high-single-digit revenue growth outlook driven by middle class consumers upgrading to increasingly convenient seasonings products in order to cut food preparation time. While margins are lower in its Japan Seasonings & Foods operations (10%), management is decisively restructuring its portfolio to focus on higher-margin, higher-growth segments and SKUs. Management is also taking full advantage of Japan's new inflationary environment to raise prices. We expect to see ~130bps of margin expansion in the Japan Seasonings & Foods business over 2026-27, on top of mid-single-digit revenue growth.

Ajinomoto management is among the most capable in Japan, with a strong focus on ROIC enhancement and a track record of selling assets and brands, and restructuring business units to enhance the margin and returns profile.

EXHIBIT 33: **Functional Materials or ABF accounts for close to one-third of Ajinomoto business profit, generating a 54% business profit margin**

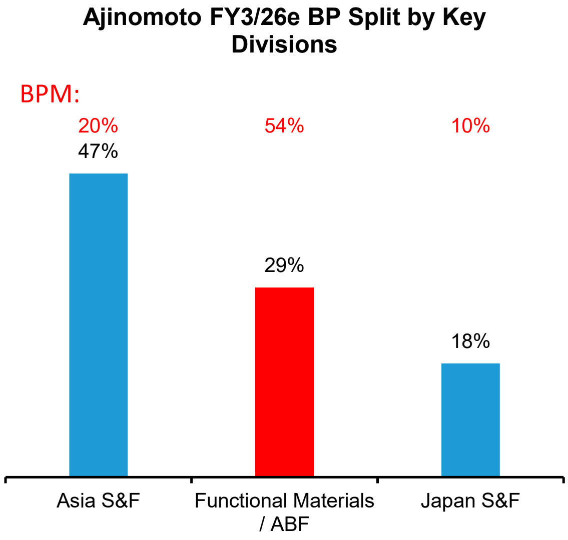

Source: Company reports, Bernstein analysis and estimates

EXHIBIT 34: **We expect ABF to drive ~60% of the company's absolute business profit growth through the next three years**

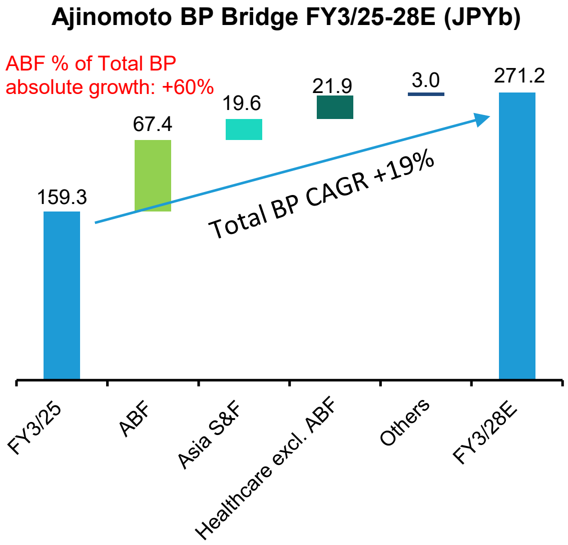

Source: Company reports, Bernstein analysis and estimates

Ajinomoto Functional Materials is an upstream supplier to substrate manufacturers such as Unimicron and Ibiden, so the impact of AI-related substrate demand hits Ajinomoto's P&L first. We also expect Ajinomoto's Functional Materials revenues to grow faster than the substrate manufacturers' as a result of its monopoly position and ongoing specification upgrades, which allow Ajinomoto to take price, as well as the increasing ABF film area tied to IC chip specification upgrades.

In FY3/24, we estimate that AI end-uses accounted for ~11% of ABF volumes, and we expect this to reach ~23% in FY3/27, driven primarily by >100% growth in FY3/26 (Exhibit 35).

We expect Ajinomoto Functional Materials revenues to grow by 38% in 4Q FY3/26 (net of a 3% forex headwind) and 26% in FY3/27, with divisional operating profit growing 33% and 44% in FY3/26 and FY3/27, respectively, driven primarily by demand for AI CPU + GPU substrates (Exhibit 36) and complemented by ~15% annual price increases on cutting-edge ABF specifications.

EXHIBIT 35: **We expect AI end-uses to increase from ~11% to ~23% of ABF sales volumes, and a higher proportion of revenues**

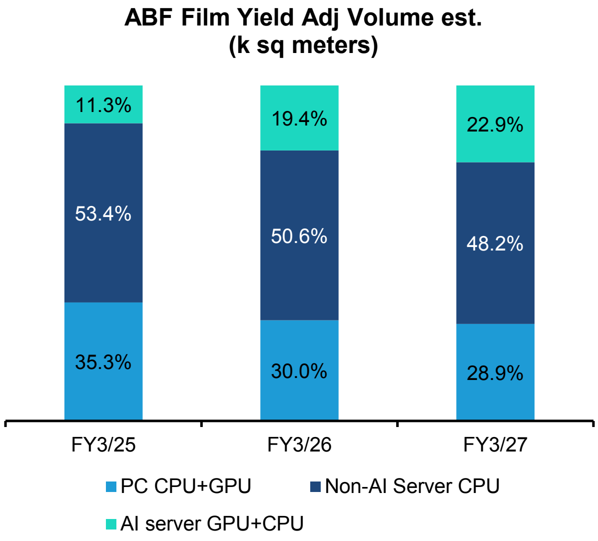

Source: Bernstein analysis and estimates (all)

EXHIBIT 36: **We expect AI Server CPU+GPU demand-related growth to drive ABF revenues**

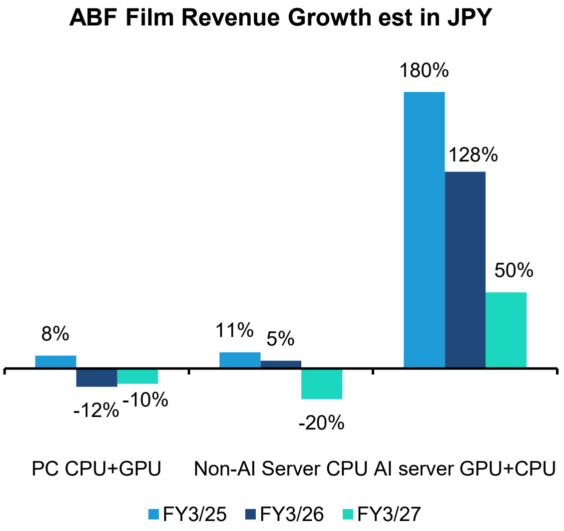

Source: Bernstein analysis and estimates (all)

We expect the 32% Functional Materials business profit growth in FY3/26 to be driven primarily by the impact of volume and pricing growth, coupled with increasing operational leverage (Exhibit 37).

While ABF operational leverage has been slightly depressed in FY3/26 (1.1x average for the first nine months, Exhibit 38) due to additional depreciation from the Gunma capacity expansion, we expect the leverage to tend back toward the historical 1.7x average in FY3/27.

EXHIBIT 37: **We expect ABF business profit growth to be driven by a combination of revenue growth and operational leverage in FY3/26, partially offset by depreciation from the new Gunma varnish production line**

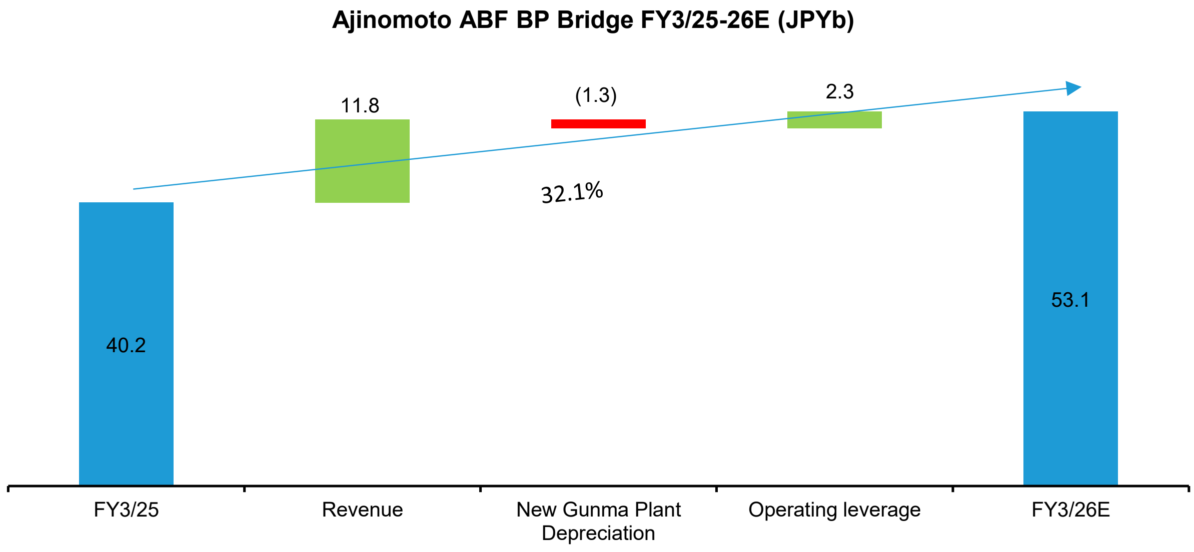

Source: Company reports, Bernstein analysis and estimates

EXHIBIT 38: **We expect Ajinomoto ABF business profit to grow 51% YoY, with a significant margin expansion in FY3/27**

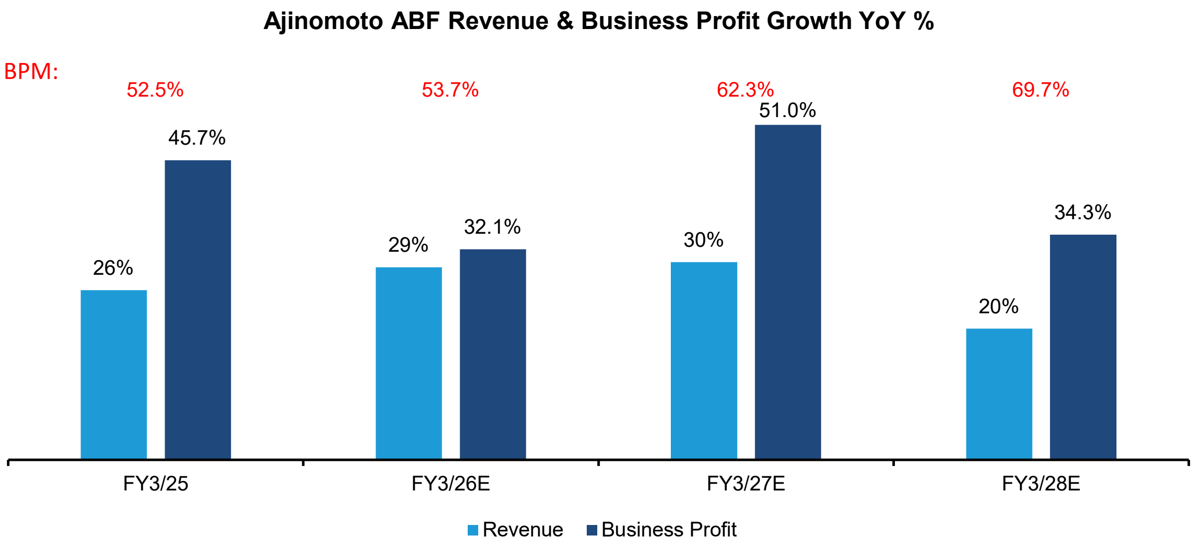

Source: Company reports, Bernstein analysis and estimates

INVESTMENT OUTLOOK

INVESTMENT IMPLICATIONS

We rate Chroma Outperform, target price NT$1,660.00.

We rate Largan Market-Perform, target price NT$2,600.00.

We rate Luxshare Outperform, target price CNY86.00.

We rate Unimicron Outperform, target price NT$610.00.

We rate NVIDIA Outperform, target price $300.00.

We rate Broadcom Outperform, target price $525.00.

We rate TSMC Outperform, target price $351.00; NT2,200.

We rate Ajinomoto Market-Perform, target price ¥5,100.00.

We rate Ibiden Outperform, target price ¥9,200.00.

VALUATION METHODOLOGY

See Disclosure Appendix of this *Whitebook* for details.

RISKS

See Disclosure Appendix of this *Whitebook* for details.

BERNSTEIN TICKER TABLE

			8 May 2026		TTM	Reported EPS				Reported P/E (x)		
			Closing	Price	Rel.							
Ticker	Rating	Cur	Price	Target	Perf.	Cur	2025A	2026E	2027E	2025A	2026E	2027E
2360.TT (Chroma ATE)	O	TWD	2,230.00	1,660.00	559.6%	TWD	27.50	33.22	43.66	81.1	67.1	51.1
3008.TT (Largan)	M	TWD	2,570.00	2,600.00	(27.9)%	TWD	192.41	154.39	178.50	13.4	16.6	14.4
002475.CH (Luxshare)	O	CNY	71.29	86.00	69.5%	CNY	2.26	2.66	3.45	31.5	26.8	20.7
3037.TT (Unimicron)	O	TWD	818.00	610.00	742.0%	TWD	4.36	11.78	19.13	187.6	69.4	42.8
NVDA (NVIDIA)	O	USD	211.50	300.00	49.9%	USD	4.77	8.88	11.94	44.3	23.8	17.7
AVGO (Broadcom)	O	USD	412.56	525.00	68.3%	USD	6.82	11.38	17.73	60.5	36.3	23.3
TSM (TSMC)	O	USD	414.15	351.00	106.1%	USD	10.37	14.52	17.28	12.7	9.6	7.6
2330.TT (TSMC)	O	TWD	2,290.00	2,200.00	101.2%	TWD	66.25	92.78	110.40	11.0	8.3	6.6
2802.JP (Ajinomoto)	M	JPY	5,030.00	5,100.00	23.6%	JPY	270,077	323,727	366,031	19.6	16.4	14.5
4062.JP (Ibiden)	O	JPY	15,365	9,200.00	602.7%	JPY	113.10	151.51	238.90	135.9	101.4	64.3
ASIAX			1,975.48									
SPX			7,337.11									
JPL			2,493.75									

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

NVDA estimate is Adjusted EPS EBITDA Margin (%); AVGO estimate is Adjusted EPS; 2802.JP estimate is EBITDA (M); NVDA, AVGO valuation is Adjusted P/E (x); TSM, 2330.TT valuation is P/B (x); 2802.JP valuation is EV/EBITDA (x); 3008.TT, 4062.JP base year is 2024; NVDA, 2802.JP base year is 2026; Source: Bloomberg, Bernstein estimates and analysis.

Note: US stocks data as of close May 7, 2026.

I. REQUIRED DISCLOSURES

References to "Bernstein" or the “Firm” in these disclosures relate to the following entities: Bernstein Institutional Services LLC (April 1, 2024 onwards), Sanford C. Bernstein & Co., LLC (pre April 1, 2024), Bernstein Autonomous LLP, BSG France S.A. (April 1, 2024 onwards), Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司, Sanford C. Bernstein (Canada) Limited, Sanford C. Bernstein (India) Private Limited (SEBI registration no. INH000006378), Sanford C. Bernstein (Singapore) Private Limited, Sanford C. Bernstein Japan KK（サンフォード・C・バーンスタイン株式会社）and analysts employed by Société Générale Africa Technologies & Services to produce Bernstein research under a Global Services Agreement in place between Bernstein and Société Générale.

Bernstein is part of a joint venture between Société Générale (SG) and AllianceBernstein, L.P. (AB). Unless specifically noted otherwise, for purposes of these disclosures, references to Bernstein’s “affiliates” relate to both SG and AB and their respective affiliates.

VALUATION METHODOLOGY

This research publication covers six or more companies. For valuation methodology and other company disclosures:
Please visit: https://bernstein-autonomous.bluematrix.com/sellside/Disclosures.action.
Or, you can also write to the Director of Compliance, Bernstein Institutional Services LLC, 245 Park Avenue, New York, NY 10167.

RISKS

This research publication covers six or more companies. For risks and other company disclosures:
Please visit: https://bernstein-autonomous.bluematrix.com/sellside/Disclosures.action.
Or, you can also write to the Director of Compliance, Bernstein Institutional Services LLC, 245 Park Avenue, New York, NY 10167.

RATINGS DEFINITIONS, BENCHMARKS AND DISTRIBUTION

EQUITY RATINGS DEFINITIONS

Bernstein brand

The Bernstein brand rates stocks based on forecasts of relative performance for the next 12 months versus the S&P 500 for stocks listed on the U.S. and Canadian exchanges, versus the Bloomberg Europe Developed Markets Large and Mid Cap Price Return Index EUR (EDME) for stocks listed on the European exchanges and emerging markets exchanges outside of the Asia Pacific region, versus the Bloomberg Japan Large and Mid Cap Price Return Index USD (JPL) for stocks listed on the Japanese exchanges, and versus the Bloomberg Asia ex-Japan Large and Mid Cap Price Return Index (ASIAX) for stocks listed on the Asian (ex-Japan) exchanges -unless otherwise specified.

The Bernstein brand has three categories of ratings:

- Outperform: Stock will outpace the market index by more than 15 pp
- Market-Perform: Stock will perform in line with the market index to within +/-15 pp
- Underperform: Stock will trail the performance of the market index by more than 15 pp

Coverage Suspended: Coverage of a company under the Bernstein research brand has been suspended. Ratings and price targets are suspended temporarily, are no longer current, and should therefore not be relied upon.

Not Rated: A rating assigned when the stock cannot be accurately valued, or the performance of the company accurately predicted, at the present time. The covering analyst may continue to publish research reports on the company to update investors on events and developments.

Not Covered (NC) denotes companies that are not under coverage.

Bernstein brand stock ratings are based on a 12-month time horizon.

Autonomous brand – common stocks

The Autonomous brand rates common stocks as indicated below. As our benchmarks we use the Bloomberg Europe 500 Banks And Financial Services Index (BEBANKS) and Bloomberg Europe Dev Mkt Financials Large and Mid Cap Price Ret Index EUR (EDMFI) index for developed European banks and Payments, the Bloomberg Europe 500 Insurance Index (BEINSUR) for European insurers, the S&P 500 and S&P Financials for US banks and Payments coverage, S5LIFE for US Insurance, the S&P Insurance Select Industry (SPSIINS) for US Non-Life Insurers coverage, and the Bloomberg Emerging Markets Financials Large, Mid and Small Cap Price Return Index (EMLSF) for emerging market banks and insurers and Payments. Ratings are stated relative to the sector (not the market).

The Autonomous brand has three categories of common stock ratings:

- Outperform (OP): Stock will outpace the relevant index by more than 10 pp
- Neutral (N): Stock will perform in line with the market index to within +/-10 pp
- Underperform (UP): Stock will trail the performance of the relevant index by more than 10 pp

Coverage Suspended: Coverage of a company under the Autonomous research brand has been suspended. Ratings and price targets are suspended temporarily, are no longer current, and should therefore not be relied upon.

Not Rated: A rating assigned when the stock cannot be accurately valued, or the performance of the company accurately predicted, at the present time. The covering analyst may continue to publish research reports on the company to update investors on events and developments.

Those denoted as ‘Feature’ (e.g., Feature Outperform FOP, Feature Under Outperform FUP) are our core ideas.

Not Covered (NC) denotes companies that are not under coverage.

Autonomous brand common stock ratings are based on a 12-month time horizon.

Autonomous brand – preferred stocks

The Autonomous brand has three categories of preferred stock ratings:

- Outperform (OP): The total return of the preferred instrument is expected to outperform preferred securities of other issuers operating in similar sectors or rating categories over the next six months.
- Neutral (N): The total return of the preferred instrument is expected to perform in line with preferred securities of other issuers operating in similar sectors or rating categories over the next six months.
- Underperform (UP): The total return of the preferred instrument is expected to underperform preferred securities of other issuers operating in similar sectors or rating categories over the next six months.

Autonomous preferred stock ratings are based on a 6-month time horizon.

AUTONOMOUS CREDIT RESEARCH

Where this report contains investment recommendations for credit instruments, as defined in article 3(1)(35) of the Market Abuse Regulation, the information below is presented to comply with its disclosure requirements.

The report may also include reference(s) to published opinions by other Autonomous or Bernstein analysts covering the equity securities of the issuer(s) referenced herein. Please note an investment recommendation for credit instruments published by the author(s) of this report may differ from the published view of the analyst covering equity securities for the issuer(s) contained in this report and vice versa.

CREDIT RATINGS DEFINITIONS

The Autonomous brand has three categories of credit ratings:

- Credit Outperform (C-OP): The total return of the Reference Credit Instrument is expected to outperform the credit spread of bonds of other issuers operating in similar sectors or rating categories over the next six months.

- Credit Neutral (C-N): The total return of the Reference Credit Instrument is expected to perform in line with the credit spread of bonds of other issuers operating in similar sectors or rating categories over the next six months.

- Credit Underperform (C-UP): The total return of the Reference Credit Instrument is expected to underperform the credit spread of bonds of other issuers operating in similar sectors or rating categories over the next six months.

Autonomous credit ratings are based on a 6-month time horizon.

A list of all investment recommendations produced by the author(s) of this report alongside credit ratings history are available upon request.

It is at the sole discretion of the Firm as to when to initiate, update and cease research coverage. The Firm has established, maintains and relies on information barriers to control the flow of information contained in one or more areas (i.e. the private side) within the Firm, and into other areas, units, groups or affiliates (i.e. public side) of the Firm

DISTRIBUTION OF EQUITY RATINGS/INVESTMENT BANKING SERVICES

Equity Rating	Market Abuse Regulation (MAR) and FINRA Rating Category	Global Rating Distribution	Investment Banking Relationships*
Outperform	BUY	51.1%	16.5%
Market-Perform (Bernstein Brand) Neutral (Autonomous Brand)	HOLD	36.3%	17.8%
Underperform	SELL	12.6%	14.9%

* These figures represent the percentage of companies within each equity rating category for which affiliates of Bernstein have provided investment banking services within the previous 12 months.
As of March 31, 2026. All figures are updated quarterly.

Prior to April 1, 2024, Sanford C. Bernstein & Co., LLC. issued the ratings and price target information in the graph(s) below for the following companies: Broadcom Inc and NVIDIA Corp.

PRICE CHARTS/ RATINGS AND PRICE TARGET HISTORY

This research publication covers six or more companies. For price chart and other company disclosures, please visit https://bernstein-autonomous.bluematrix.com/sellside/Disclosures.action or you can write to the Director of Compliance, Bernstein Institutional Services LLC, 245 Park Avenue, New York, NY 10167.

CONFLICTS OF INTEREST

All statements in this report attributable to Gartner represent Bernstein's interpretation of data, research opinion or viewpoints published as part of a syndicated subscription service by Gartner, Inc., and have not been reviewed by Gartner. Each Gartner publication speaks as of its original publication date (and not as of the date of this report). The opinions expressed in Gartner publications are not representations of fact, and are subject to change without notice.

Stacy A. Rasgon maintains long positions in various crypto currencies.
Bernstein Autonomous LLP or BSG France SA, beneficially, has either a net long or short position of 0.5% or more of the total issued share capital of a class of common equity securities of the following MiFID eligible securities: Broadcom Inc.

SG and/or its affiliates beneficially own 1% or more of a class of common equity securities of the following company: Ibiden Co Ltd.

Bernstein has received compensation for non-investment banking securities-related products or services in the previous twelve months from the following clients: NVIDIA Corp.

Bernstein and/or affiliates have received compensation for non-investment banking securities-related products or services in the previous twelve months from the following clients: Ajinomoto Co Inc.

Certain affiliates of Bernstein act as market maker or liquidity provider in the equities securities of: Broadcom Inc and NVIDIA Corp.

Certain affiliates of Bernstein act as market maker or liquidity provider in the debt securities of: Broadcom Inc.

OTHER MATTERS

The legal entity(ies) employing the analyst(s) listed in this report, and their location, can be determined by the country code of their phone number, as follows:

+1 Bernstein Institutional Services LLC; New York, New York, USA

+44 Bernstein Autonomous LLP; London UK

+212 Société Générale Africa Technologies & Services; Casablanca, Morocco

+33 BSG France S.A.; Paris, France

+34 BSG France S.A.; Madrid, Spain

+41 Bernstein Autonomous LLP; Geneva, Switzerland

+49 BSG France S.A.; Frankfurt, Germany

+91 Sanford C. Bernstein (India) Private Limited; Mumbai, India

+852 Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司; Hong Kong, China

+65 Sanford C. Bernstein (Singapore) Private Limited; Singapore

+81 Sanford C. Bernstein Japan KK; Tokyo, Japan

Where this report has been prepared by research analyst(s) employed by a non-US affiliate, such analyst(s), is/are (unless otherwise expressly noted below) not registered as associated persons of Bernstein Institutional Services LLC or any other SEC-registered broker-dealer and are not licensed or qualified as research analysts with FINRA. Accordingly, such analyst(s) may not be subject to FINRA's restrictions regarding (among other things) communications by research analysts with a subject company, interactions between research analysts and investment banking personnel, participation by research analysts in solicitation and marketing activities relating to investment banking transactions, public appearances by research analysts, and trading securities held by a research analyst account.

Where this report has been prepared by research analyst(s) employed by Société Générale Africa Technologies & Services (part of the Société Générale group of companies), it has been prepared on behalf of a Bernstein company under a Global Services Agreement in place between Bernstein and Société Générale.

CERTIFICATION

Each research analyst listed in this report, who is primarily responsible for the preparation of the content of this report, certifies that all of the views expressed in this publication accurately reflect that analyst's personal views about any and all of the subject securities or issuers and that no part of that analyst's compensation was, is, or will be, directly or indirectly, related to the specific recommendations or views in this publication.

II. ADDITIONAL GLOBAL CONFLICT DISCLOSURES

It is at the sole discretion of the Firm as to when to initiate, update and cease research coverage. The Firm has established, maintains and relies on information barriers to control the flow of information contained in one or more areas (i.e., the private side) within the Firm, and into other areas, units, groups or affiliates (i.e., public side) of the Firm.

III. OTHER IMPORTANT INFORMATION AND DISCLOSURES

Separate branding is maintained for “Bernstein” and “Autonomous” research products.

- Bernstein produces a number of different types of research products including, among others, fundamental analysis and quantitative analysis under both the “Autonomous” and “Bernstein” brands. Recommendations contained within one type of research product may differ from recommendations contained within other types of research products, whether as a result of differing time horizons, methodologies or otherwise. Furthermore, views or recommendations within a research product issued under one brand may differ from views or recommendations under the same type of research product issued under the other brand. The Research Ratings System for the two brands and other information related to those Rating Systems are included in the previous section.

- Autonomous operates as a separate business unit within the following entities: Bernstein Institutional Services LLC, Bernstein Autonomous LLP, Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司 and Sanford C. Bernstein (India) Private Limited. For information relating to “Autonomous” branded products (including certain Sales materials) please visit: www.autonomous.com. For information relating to Bernstein branded products please visit: www.bernsteinresearch.com.

Analysts are compensated based on aggregate contributions to the research franchise as measured by account penetration, productivity and proactivity of investment ideas. No analysts are compensated based on performance in, or contributions to, generating investment banking revenues.

This report has been produced by an independent analyst as defined in Article 3 (1)(34)(i) of EU 596/2014 Market Abuse Regulation (“MAR”) and the same article of MAR as it forms part of United Kingdom domestic law by virtue of the European Union (Withdrawal) Act 2018.

To our readers in the United States: Bernstein Institutional Services LLC, a broker-dealer registered with the U.S. Securities and Exchange Commission (“SEC”) and a member of the U.S. Financial Industry Regulatory Authority, Inc. (“FINRA”) is distributing this publication in the United States and accepts responsibility for its contents. Where this material contains an analysis of debt product(s), such material is intended only for institutional investors and is not subject to the US independence and disclosure standards applicable to debt research prepared for retail investors.

Bernstein Institutional Services LLC may act as principal for its own account or as agent for another person (including an affiliate) in sales or purchases of any security which is a subject of this report. This report does not purport to meet the objectives or needs of any specific individuals, entities or accounts.

To our readers in Canada: If this publication pertains to a Canadian domiciled company, it is being distributed in Canada by Sanford C. Bernstein (Canada) Limited, which is licensed and regulated by the Canadian Investment Regulatory Organization. If the publication pertains to a non-Canadian domiciled company, it is being distributed by Bernstein Institutional Services LLC, which is licensed and regulated by both the SEC and FINRA, into Canada under the International Dealers Exemption.

This document may not be passed onto any person in Canada unless that person qualifies as "permitted client" as defined in Section 1.1 of NI 31-103.

To our readers in Brazil: This report has been prepared by Bernstein Institutional Services LLC, and Banco BTG Pactual S.A. ("BTG") is responsible for the distribution of this report in Brazil.

To readers in the United Kingdom: This publication has been issued or approved for issue in the United Kingdom by Bernstein Autonomous LLP, authorised and regulated by the Financial Conduct Authority and located at 60 London Wall, London EC2M 5SH, +44 (0)20-7170-5000. Registered in England & Wales No OC343985.

This document is for distribution only to persons who (i) have professional experience in matters relating to investments falling within Article 19(5) of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005 (as amended, the “Financial Promotion Order”), (ii) are persons falling within Article 49(2)(a) to (d) (“high net worth companies, unincorporated associations, etc.”) of the Financial Promotion Order, (iii) are outside the United Kingdom, or (iv) are persons to whom an invitation or inducement to engage in investment activity (within the meaning of section 21 of the FSMA) in connection with the issue or sale of any securities may otherwise lawfully be communicated or caused to be communicated (all such persons together being referred to as “relevant persons”). This document is directed only at relevant persons and must not be acted on or relied on by persons who are not relevant persons. Any investment or investment activity to which this document relates is available only to relevant persons and will be engaged in only with relevant persons.

To our readers in the member states of the EEA: This publication is being distributed by BSG France SA, which is authorised and regulated by the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and Autorité des Marchés Financiers (AMF).

To our readers in Hong Kong: This publication is being distributed in Hong Kong by Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司, which is licensed and regulated by the Hong Kong Securities and Futures Commission (Central Entity No. AXC846) to carry out Type 4 (Advising on Securities) regulated activities and subject to the licensing conditions mentioned in the SFC Public Register (https://www.sfc.hk/publicregWeb/corp/AXC846/details)). This publication is solely for professional investors, as defined in the Securities and Futures Ordinance (Cap. 571). The purpose of this report is solely to provide an analysis of the issuers referred to in this report and is not intended for any purpose contrary to the laws of Hong Kong.

To our readers in Singapore: This publication is being distributed in Singapore by Sanford C. Bernstein (Singapore) Private Limited, only to accredited investors or institutional investors, as defined in the Securities and Futures Act 2001 of Singapore ("SFA"). Recipients in Singapore should contact Sanford C. Bernstein (Singapore) Private Limited in respect of matters arising from, or in connection with, this publication. Sanford C. Bernstein (Singapore) Private Limited is regulated by the Monetary Authority of Singapore and licensed under the SFA as a capital markets services licence holder for dealing in capital markets products that are securities and collective investment schemes and an exempt financial adviser for advising on, issuing and promulgating analyses and reports on securities. Sanford C. Bernstein (Singapore) Private Limited is registered in Singapore with Company Registration No. 20213710W and located at One Raffles Quay, #27-11 South Tower, Singapore 048583, +65-6230-4612.

To our readers in the People's Republic of China: The securities referred to in this document are not being offered or sold and may not be offered or sold, directly or indirectly, in the People's Republic of China (for such purposes, not including the Hong Kong and Macau Special Administrative Regions or Taiwan, the "PRC") in contravention of any applicable laws of the PRC.

This document does not constitute an offer to sell or the solicitation of an offer to buy any securities in the PRC to any person to whom it is unlawful to make the offer or solicitation in the PRC.

We do not represent that this document may be lawfully distributed, or that any securities may be lawfully offered, in compliance with any applicable registration or other requirements in the PRC, or pursuant to an exemption available thereunder, or assume any responsibility for facilitating any such distribution or offering. In particular, no action has been taken by us which would permit a public offering of any securities or distribution of this document in the PRC. Accordingly, the securities are not being offered or sold within the PRC by means of this document or any other document. Neither this document nor any advertisement or other offering material may be distributed or published in the PRC, except under circumstances that will result in compliance with any applicable laws and regulations.

To our readers in Japan: This publication is being distributed in Japan by Sanford C. Bernstein Japan KK (サンフォード・C・バーンスタイン株式会社), which is registered in Japan as a Financial Instruments Business Operator with the Kanto Local Finance Bureau (registration number: The Director-General of Kanto Local Finance Bureau (FIBO) No.3387) and regulated by the Financial Services Agency. It is also a member of Japan Investment Advisers Association. This publication is solely for qualified institutional investors in Japan only, as defined in Article 2, paragraph (3), items (i) of the Financial Instruments and Exchange Act.

For the institutional client readers in Japan who have been granted access to the Bernstein website by Daiwa Securities Group Inc. ("Daiwa"), your access to this document should not be construed as meaning that Bernstein is providing you with investment advice for any purposes. Whilst Bernstein has prepared this document, your relationship is, and will remain with, Daiwa, and Bernstein has neither any contractual relationship with you nor any obligations towards you.

To our readers in Australia: Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司 is responsible for distributing research in Australia. It is regulated by the Securities and Exchange Commission under U.S. laws, by the Financial Conduct Authority under U.K. laws, which differs from Australian laws. Sanford C. Bernstein (Hong Kong) Limited 盛博香港有限公司 is exempt from the requirement to hold an Australian financial services license under the Corporations Act 2001 in respect of the provision of the following financial services to wholesale clients:

- providing financial product advice;
- dealing in a financial product;
- making a market for a financial product; and
- providing a custodial or depository service.

To our readers in India: This publication is being distributed in India by Sanford C. Bernstein (India) Private Limited (SCB India) which is licensed and regulated by Securities and Exchange Board of India ("SEBI") as a research analyst entity under the SEBI (Research Analyst) Regulations, 2014, having registration no. INH000006378 and as a stock broker having registration no. INZ000213537. SCB India is currently engaged in the business of providing research and stock broking services. Please refer to www.bernsteinresearch.in for more information.

- SCB India is a Private limited company incorporated under the Companies Act, 2013, on April 12, 2017 bearing corporate identification number U65999MH2017FTC293762, and registered office at Level 3A, 4th Floor, First International Financial Centre, Plot Nos C-54 and C-55, G Block, Near CBI Office, Bandra Kurla Complex, Bandra (East), Mumbai 400098, Maharashtra, India (Phone No: +91-22-68421401).
- For details of Associates (i.e., affiliates/group companies) of SCB India, kindly email MUM-BERNSTEIN-InCompliance@bernsteinsg.com.
- SCB India does not have any disciplinary history as on the date of this report.
- Except as noted above, SCB India and/or its Associates (i.e., affiliates/group companies), the Research Analysts authoring this report, and their relatives
  - do not have any financial interest in the subject company
  - do not have actual/beneficial ownership of one percent or more in securities of the subject company;
  - is not engaged in any investment banking activities for Indian companies, as such;
  - have not managed or co-managed a public offering in the past twelve months for any Indian companies;
  - have not received any compensation for investment banking services or merchant banking services from the subject company in the past 12 months;
  - have not received compensation for brokerage services from the subject company in the past twelve months;
  - have not received any compensation or other benefits from the subject company or third party related to the specific recommendations or views in this report; and
  - do not currently, but may in the future, act as a market maker in the financial instruments of the companies covered in the report.
  - do not have any conflict of interest in the subject company as of the date of this report.
- Except as noted above, the subject company has not been a client of SCB India during twelve months preceding the date of distribution of this research report. Neither SCB India nor its Associates (i.e., affiliates/group companies) have received compensation for products or services other than investment banking, merchant banking or brokerage services from the subject company in the past twelve months.
- The principal research analyst(s) who prepared this report, members of the analysts' team, and members of their households are not an officer, director, employee or advisory board member of the companies covered in the report.
- Our Compliance officer / Grievance officer is Ms. Rupal Talati, who can be reached at +91-22-68421451, or MUM-BERNSTEIN-InCompliance@bernsteinsg.com / Scbin-investorgrievance@bernsteinsg.com
- The Research investor charter and Terms & Conditions of SCB India are available on its website and may be accessed at Sanford C. Bernstein (India) Private Limited (https://bernsteinresearch.in/) for your reference.

- Disclaimer: Registration granted by SEBI, and certification from NISM, is in no way a guarantee of performance of the intermediary or provide any assurance of returns to investors. Investments in securities market are subject to market risks. Read all the related documents carefully before investing.

To our readers in Switzerland: This document is provided in Switzerland by or through Bernstein Autonomous LLP, and is provided only to qualified investors as defined in article 10 of the Swiss Collective Investment Scheme Act (“CISA”) and related provisions of the Collective Investment Scheme Ordinance and in strict compliance with applicable Swiss law and regulations. The products mentioned in this document may not be suitable for all types of investors. This document is based on the Directives on the Independence of Financial Research issued by the Swiss Bankers Association (SBA) in January 2008.

To our readers in the Middle East: Bernstein Autonomous LLP, DIFC branch has its principal office at Gate Village 06, DIFC, Dubai, UAE. Bernstein Autonomous LLP, DIFC branch is regulated by the Dubai Financial Services Authority (DFSA) with the registration number CL10040 and is provisioned for Arranging Deals in Investments and Advising on Financial Products. All communications and services are directed at Professional Clients and Market Counterparties only (as defined in the DFSA rulebook). Persons other than Professional Clients and Market Counterparties, such as Retail Clients, are not the intended recipients of our communications or services.

LEGAL

All research publications are disseminated to our clients through posting on the firm's password protected websites, bernsteinresearch.com and autonomous.com. Certain, but not all, research publications are also made available to clients through third-party vendors or redistributed to clients through alternate electronic means as a convenience.

This publication has been published and distributed in accordance with the Firm's policy for management of conflicts of interest in investment research, a copy of which is available from Bernstein Institutional Services LLC, Director of Compliance, 245 Park Avenue, New York, NY 10167. Additional disclosures and information regarding Bernstein's business are available on our website www.bernsteinresearch.com.

The securities described herein may not be eligible for sale in all jurisdictions or to certain categories of investors where that permission profile is not consistent with the licenses held by the entities noted herein. This document is for distribution only as may be permitted by law. This publication is not directed to, or intended for distribution to or use by, any person or entity who is a citizen or resident of, or located in any locality, state, country or other jurisdiction where such distribution, publication, availability or use would be contrary to law or regulation or which would subject any of the entities referenced herein or any of their subsidiaries or affiliates to any registration or licensing requirement within such jurisdiction. This publication is based upon public sources we believe to be reliable, but no representation is made by us that the publication is accurate or complete. We do not undertake to advise you of any change in the reported information or in the opinions herein. This publication was prepared and issued by entity referred to herein for distribution to eligible counterparties or professional clients. This publication is not an offer to buy or sell any security, and it does not constitute investment, legal or tax advice. The investments referred to herein may not be suitable for you. Investors must make their own investment decisions in consultation with their professional advisors in light of their specific circumstances. The value of investments may fluctuate, and investments that are denominated in foreign currencies may fluctuate in value as a result of exposure to exchange rate movements. Information about past performance of an investment is not necessarily a guide to, indicator of, or assurance of, future performance.

This report is directed to and intended only for our clients who are “eligible counterparties”, “professional clients”, “institutional investors” and/or “professional investors” as defined by the aforementioned regulators, and must not be redistributed to retail clients as defined by the aforementioned regulators. Retail clients who receive this report should note that the services of the entities noted herein are not available to them and should not rely on the material herein to make an investment decision. The result of such act will not hold the entities noted herein liable for any loss thus incurred as the entities noted herein are not registered/ authorised/ licensed to deal with retail clients and will not enter into any contractual agreement/arrangement with retail clients. This report is provided subject to the terms and conditions of any agreement that the clients may have entered into with the entities noted herein. All research reports are disseminated on a simultaneous basis to eligible clients through electronic publication to our client portal.

The information in this report was prepared by Bernstein solely for the internal business use of our clients. Clients may store, display, analyze, reformat and print the information in this report for this limited use only. Clients may not copy, alter, create derivative works, resell, reverse engineer, commercially exploit, share or distribute any part of the information contained herein for any purpose without Bernstein’s express written consent. These restrictions include extracting data or using the content to develop indices or other products. Further, you may not use this report, or any portion of this report, to train or finetune any third-party machine learning or artificial intelligence system, or as a prompt or input into any such system. You also may not, without Bernstein’s express written consent, do any of the foregoing in connection with your own internal machine learning or artificial intelligence system.

Bernstein may use artificial intelligence tools in the preparation of its materials. Any such materials are reviewed by Bernstein's research analysts prior to publication.

This report has been prepared for information purposes only and is based on current public information that we consider reliable, but the entities noted herein do not warrant or represent (express or implied) as to the sources of information or data contained herein are accurate, complete, not misleading or as to its fitness for the purpose intended even though the entities noted herein rely on reputable or trustworthy data providers, it should not be relied upon as such. Opinions expressed are the author(s)' current opinions as of the date appearing on the material only and we do not undertake to advise you of any change in the reported information or in the opinions herein.

This publication was prepared and issued by the entity referred to herein for distribution to eligible counterparties or professional clients. The information in this report is intended for general circulation and does not constitute an offer to buy or sell any security, investment, legal or tax advice nor a personal recommendation, as defined by any of the aforementioned regulators. It does not take into account the particular investment objectives, financial situations, or needs of individual investors. The report has not been reviewed by any of the aforementioned regulators and does not represent any official recommendation from the aforementioned regulators. The investments referred to herein may not be suitable for you. Investors must make their own investment decisions in consultation with advice sought from a financial adviser regarding the suitability of the investment product, taking into account the specific investment objectives, financial situation or particular needs of any recipient of the recommendation, before the recipient makes a commitment to purchase the investment product.

The analysis contained herein is based on numerous assumptions. Different assumptions could result in materially different results. The information in this report does not constitute, or form part of, any offer to sell or issue, or any offer to purchase or subscribe for shares, or to induce engage in any other investment activity. The value of any securities or financial instruments mentioned in this report may fluctuate subject to market conditions. Information about past performance of an investment is not necessarily a guide to, indicative of, or assurance of future performance. Estimates of future performance mentioned by the research analyst in this report are based on assumptions that may not be realized due to unforeseen factors like market volatility/fluctuation. In relation to securities or financial instruments denominated in a foreign currency other than the clients' home currency, movements in exchange rates will have an effect on the value, either favorable or unfavorable. Before acting on any recommendations in this report, recipients should consider the appropriateness of investing in the subject securities or financial instruments mentioned in this report and, if necessary, seek for independent professional advice.

The securities described herein may not be eligible for sale in all jurisdictions or to certain categories of investors where that permission profile is not consistent with the licenses held by the entities noted herein. This document is for distribution only as may be permitted by law. It is not directed to, or intended for distribution to or use by, any person or entity who is a citizen or resident of or located in any locality, state, country or other jurisdiction where such distribution, publication, availability or use would be contrary to law or regulation or would subject the entities noted herein to any regulation or licensing requirement within such jurisdiction.

Source: Bloomberg Index Services Limited. BLOOMBERG® is a trademark and service mark of Bloomberg Finance L.P. and its affiliates (collectively "Bloomberg"). Bloomberg or Bloomberg's licensors own all proprietary rights in the Bloomberg Indices. Neither Bloomberg nor Bloomberg's licensors approves or endorses this material, or guarantees the accuracy or completeness of any information herein, or makes any warranty, express or implied, as to the results to be obtained therefrom and, to the maximum extent allowed by law, neither shall have any liability or responsibility for injury or damages arising in connection therewith.