Morgan Stanley | RESEARCH

IDEA

October 19, 2026 01:16 AM GMT

Greater China Technology Hardware | Asia Pacific

2026 OCP Global Summit – MGX Upgrades Take Shape; Maia 300 on Display

Howard Kao
Equity Analyst
Howard.Kao@morganstanley.com +886 2 2730-2989

Irene Yen
Equity Analyst
Irene.Yen@morganstanley.com +886 2 2730-2869

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

GREATER CHINA TECHNOLOGY HARDWARE
Asia Pacific
Industry View **In-Line**

At the rack level, Maia 300 is the key incremental story, while Nvidia/AMD exhibits offered limited new system-level hardware. Nvidia's MGX component upgrades are the main supply-chain watch points, although PTFE and next-gen interconnect and cooling connector (ZQD) adoption remains uncertain.

Nvidia/AMD – focus on ramps rather than new platforms. Most ODM displays were similar to Computex Taipei (early June), so timing of ramp is key. For Nvidia, VR200 racks have started a small-volume ramp, with the Nvidia Groq 3 LPX rack following in late 2026/early 2027 and standalone Vera CPU racks in late 1Q27/early 2Q27, based on our checks. For AMD, Helios remained largely unchanged with recent checks, with racks ramping in 1Q27, with pilot-build observations pointing to Shennan Circuits for EAM PCBs and Fositek/Parker for quick disconnects.

Maia 300 - a larger in-house inference opportunity. Microsoft plans production deployments in 2027 as it scales Maia toward multi-gigawatt deployments. Each rack integrates 96 accelerators within 167 kW. Based on that, each gigawatt of deployment would imply 5 to 5.5K racks, and we believe Hon Hai and Wiwynn could be the ODM partners, while our pilot-build observations point to ISU Petasys for the accelerator PCB, AVC for fans, and Fositek/Staubli for quick disconnects.

PTFE - technical upside, but no confirmed timing of adoption. Nvidia is evaluating PTFE for Rubin Ultra on Oberon following the Kyber delay and showcased its NVLink PTFE PCB with seven suppliers (though we think there are suppliers missing from this list). Better signal performance and reduced reliance on Q glass are potential benefits, but processing complexity, higher costs and yield/qualification risks remain hurdles to commercial adoption.

Ultra Dense Interconnect - next-generation interconnect opportunity for FIT. Nvidia showcased six suppliers for this next-generation interconnect, which is intended to replace Amphenol's current HD2 connector in midplane and cable cartridge applications. However, we believe adoption is still about two years away. In the meantime, the supply chain is working to improve yields and bring ASPs below US$100, with pogo-pin supply currently a bottleneck. Our checks identify FIT as the key design partner for Kyber's planned board-to-board connector, a position that bodes well for its supply share once the product ramps.

ZQD - higher flow density and an expected ASP premium. Nvidia's new ZQD targets higher flow rates per unit area and was displayed with 14 suppliers. Although the supply chain is working to improve yields and lower costs, greater design complexity should keep ASPs above current QD pricing of US$20–30.

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

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

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

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

Major Takeaways from OCP Summit 2026

ODM/OEM: Almost all the ODM/OEMs had a booth at OCP, but during our visit, we feel most of the products were largely similar to what was showcased at Computex, such as the Vera Rubin NVL72 rack, the LPX tray, the Vera standalone tray. So net, we think from a hardware perspective, there was nothing new on Nvidia or AMD. We believe the Nvidia VR200 rack is already starting ramp in small volume, with the LPX rack starting to ramp later this year or early next year, and the Vera standalone CPU rack ramping late 1Q27 or early 2Q27.

AMD: Wiwynn, Meta, and AMD each showcased the Helios rack, which was unchanged from last year's OCP display with one exception: this year's version featured the actual compute tray in place of last year's dummy tray, allowing a closer look. That said, the compute tray itself was not new to us, as we had already seen it at OCP Taiwan in August.

- Based on our observations at the OCP Summit, we believe Shennan Circuits is one of the pilot build PCB suppliers for the MI455 Helios EAM PCB.
- Based on our observations at the OCP Summit, we believe Fositek (MQD02) and Parker (UQD04) are one of the pilot build quick disconnect connector suppliers for the MI455 Helios compute tray.
- HPE displayed the Helios switch tray at the OCP Summit. We believe the other supplier of the Helios switch trays is Celestica.

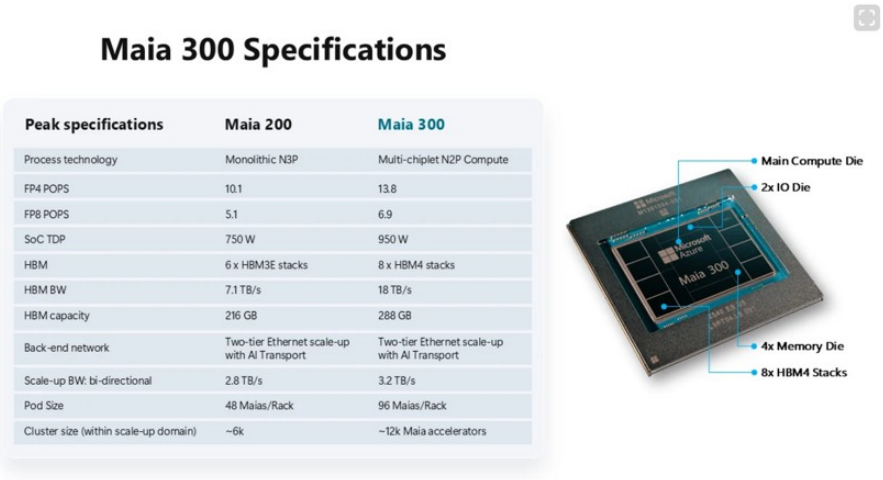
Microsoft: Maia 300 represents a major expansion of Microsoft's in-house inference infrastructure. Announced on October 12, 2026, with production deployments planned for 2027, it forms part of Microsoft's commitment to scale the Maia family to multi-gigawatt deployments and make it the primary inference platform for its products and services. The design targets long-context reasoning, conversational AI, and multi-step agentic workloads, with an emphasis on improving inference economics through coordinated silicon, software, and system design.

Maia 300 rack integrates 96 accelerators within a 167 kW power envelope, providing approximately 27.6 TB of aggregate HBM capacity. There are 4 accelerators per tray, implying 24 compute trays per rack, including six high-radix switch ASICs that are integrated into three liquid-cooled switch trays, providing the rack-local connectivity required by all 96 accelerators. Microsoft projects 40% more tokens per megawatt than leading GPU-based systems and up to 7× Maia 200's rack-level token throughput, though this takes into account that the Maia 300 rack includes twice as many accelerators per rack.

- The Official Microsoft Blog A two-tier Ethernet scale-up network connects up to 128 racks, or 12,288 accelerators, while liquid cooling removes more than 80% of rack heat. Each compute tray includes 8x fans (Exhibit 3), and based on our observations at the OCP Summit, we believe AVC is one of the pilot build fan suppliers for the sample compute tray.
- Based on our observations at the OCP Summit, we believe ISU Petasys is one of the pilot build PCB suppliers for Microsoft's Maia 300 AI accelerator.
- Based on our observations at the OCP Summit, we believe Fositek and Staubli are pilot build MQD03 (quick disconnect connector) suppliers for the compute tray.

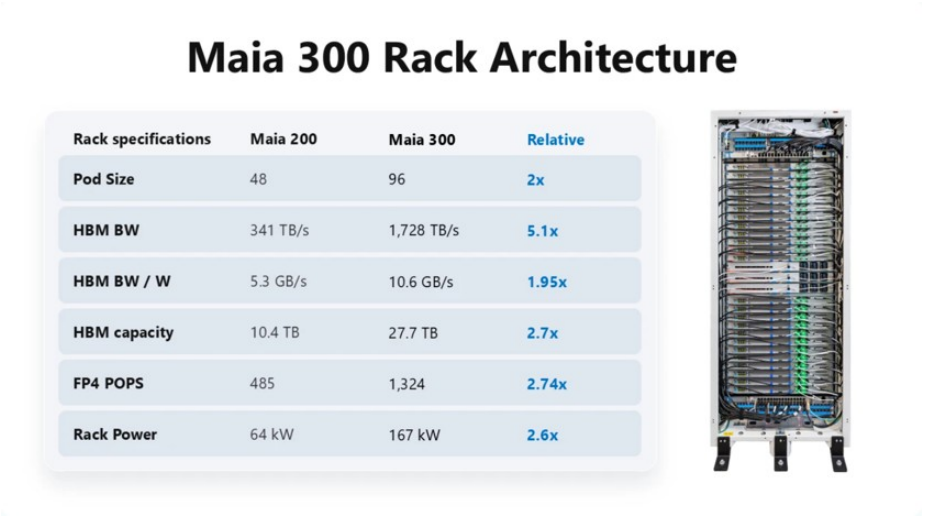
- We believe Hon Hai and Wiwynn could be ODM partners for these racks.
- With each rack being 167 kW of power, one gigawatt of deployment would imply 5-5.5K racks, assuming a PUE of 1.1-1.2.

Exhibit 1: Maia 300 Specifications



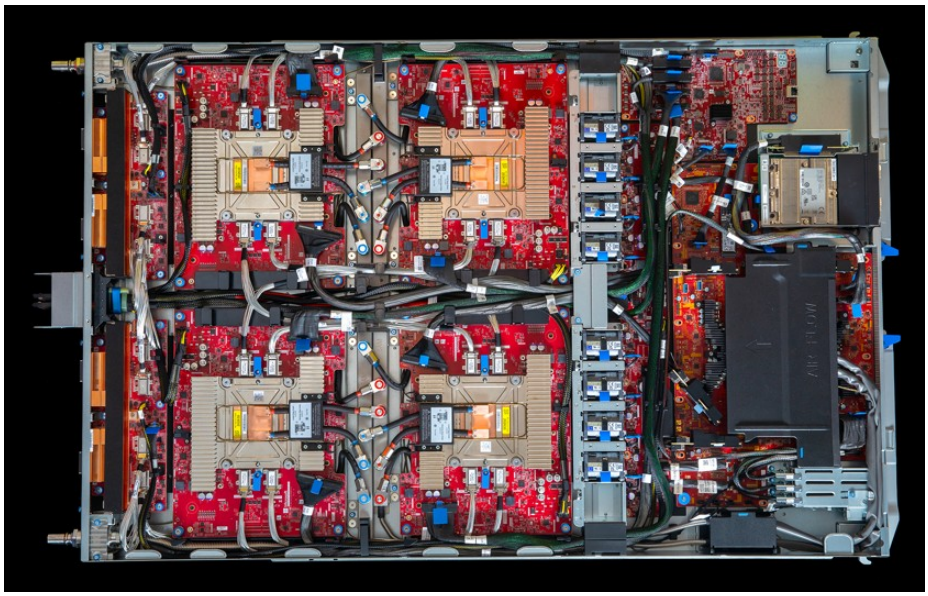
Source: Microsoft.

Exhibit 2: Maia 300 Rack Architecture



Source: Microsoft.

Exhibit 3: Maia 300 server compute tray



Source: Microsoft.

Nvidia MGX Ecosystem

PTFE: We previously wrote that our supply chain checks suggest that Nvidia had initially planned to adopt PTFE CCL for the very high-layer-count midplane in its Kyber architecture. With Kyber now delayed, and Rubin Ultra no longer expected to use that architecture, as previously discussed, the focus has shifted back to the Oberon architecture. Nvidia is again evaluating PTFE for Rubin Ultra on Oberon, and at this year's OCP Summit, Nvidia showcased its NVLink PTFE PCB, with a total of 7 suppliers, including VGT, TTM, Taiflex, SYTECH, Shennan Circuits, Kinwong, and Flexium (Exhibit 4 & Exhibit 5). However, our checks indicate that actual adoption for Rubin Ultra remains uncertain.

However, we find the supplier list puzzling. Our checks suggest that WUS PCB and Unimicron are also working on this solution, and that Flexium, which has collaborated with other PCB suppliers, may handle only a few process steps. Our checks also indicate that Shennan Circuits, WUS PCB and VGT are now the leading PCB makers for this solution, with relatively higher yields.

Advantages of PTFE

- **Best signal performance:** the lowest Dk and Df of the mainstream options, below even standard M9-class CCLs. This matters most where insertion-loss budgets are tightest, i.e. the highest speeds and highest layer counts.
- **Stable electricals:** the low Dk/Df hold up across frequency and temperature, and PTFE absorbs very little moisture, so performance is stable in use.
- **Less dependence on Q glass:** because PTFE's intrinsic dielectric constant is already very low, reaching the target Dk requires less (and in some hybrid constructions, none) of the expensive, supply-constrained Q glass. That is both a cost lever and a supply-security benefit.

Disadvantages of PTFE

- **Mechanically soft:** low modulus makes it difficult to handle and to hold to tolerance.
- **Difficult lamination:** PTFE does not flow like a thermoset, so bonding needs either very high temperatures or lower-melting fluoropolymer (such as PFA (polyfluoroalkoxy))/thermoset bond layers; its high, non-linear CTE stresses copper adhesion and complicates layer-to-layer registration, and pressing typically requires vacuum or an inert atmosphere with longer cycles.
- **Difficult drilling:** the soft, gummy resin tends to smear and can tear, and any ceramic filler is abrasive on tooling. Also, the hole wall will not accept electroless copper without an added activation step (sodium etch or plasma) on top of normal desmear.
- **Higher cost and qualification risk:** materials and fabrication are both more expensive, the qualified supply base is narrower, and yields at very high layer counts carry risk.

Exhibit 4: Nvidia's NVLink PTFE Board

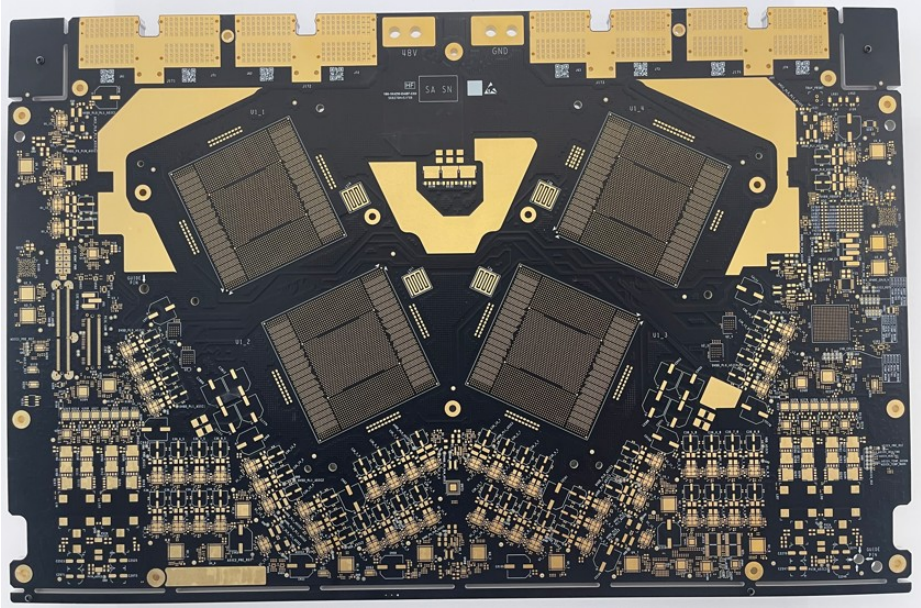
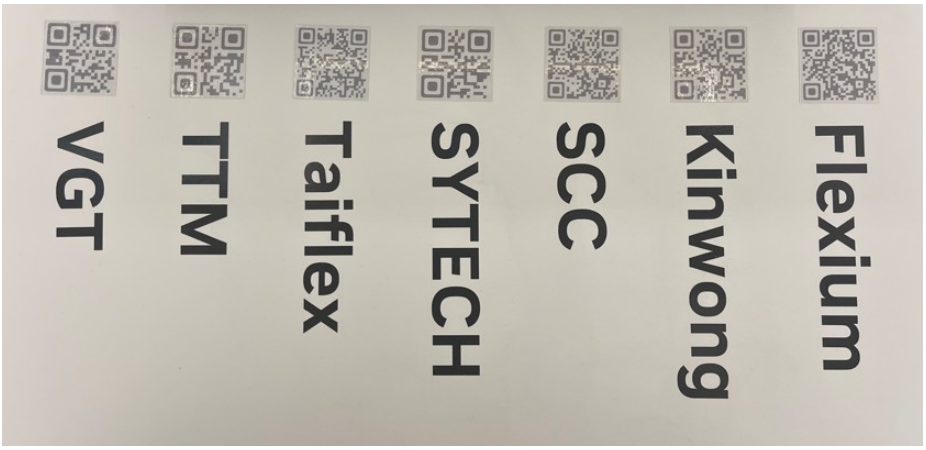


Exhibit 5: Nvidia PTFE PCB Supplier List



Ultra Dense Interconnect: Currently, Nvidia's NVL72 racks are using Amphenol's Paladin HD2 connectors for VR200 NVL72 midplanes and cable cartridges. Our previous checks done around Computex 2026 in June, indicate ASP could reach US$70-90+ per set. For the next-generation Kyber racks, our checks have previously indicated that a new midplane connector called "Voronoi" will be introduced, and FIT is the key design partner for this new board-to-board connector. At this year's OCP Summit, Nvidia showcased a new connector it calls the "Ultra Dense Interconnect", with six suppliers in total. This was our first in-person look at the design. We believe adoption is still two years out, and the supply chain is currently focused on improving yields and optimizing costs to bring the ASP below US$100. Our checks suggest that pogo pin supply is currently a bottleneck for this connector.

Exhibit 6: Nvidia next-generation Ultra Dense Interconnect



ZQD: At COMPUTEX in Taipei earlier this year, we noted that our checks indicated Nvidia was planning to introduce a new quick disconnect (QD) design for next-generation server racks to achieve higher flow rates per unit area as cooling requirements rise. At this year's OCP Summit, Nvidia showcased that design, called ZQD, with 14 suppliers in total. Our checks indicate the supply chain is also working to improve yields and optimize costs to bring ZQD ASPs down. However, given the greater complexity of the design, we expect ZQD ASPs to remain above current QD pricing of US$20–30 per set.

Exhibit 7: Nvidia's new ZQD liquid cooling quick disconnect connectors



人在草木间