20 Sep 2026 17:16:31 ET | 12 pages

Taiwan Semiconductors

US marketing feedback- AI hardware growth beyond GPU/ Accelerator

Laura (Chia Yi) Chen [AC]

Jack Chen

Nicholas Lai

CITI'S TAKE

In our recent US marketing, investors focused on the AI infrastructure cycle's longevity beyond 2028 and the next value-upgrade bottlenecks. Taiwan's supply chain offers strong visibility, with AI growth broadening from GPUs into networking, custom ASICs, CPO, advanced packaging, and power infrastructure. Nvidia retains an edge via its integrated ecosystem, while CPO emerges as the key architectural shift addressing bandwidth and power bottlenecks. Core beneficiaries include TSMC (2nm/3nm and advanced packaging), ASEH (outsourcing), MediaTek (TPU content) and Hon Hai (high-value Rubin racks and CSP ASICs). The main long-term risk remains whether AI monetization can justify hyperscalers' massive capex.

AI Semiconductor Outlook: Logic, Networking and CPU Could Drive the Next Leg of Growth — We believe custom ASIC units should grow rapidly from a low base, but Nvidia could continue to capture a major share of system value. Looking beyond accelerators, the next major growth vectors should increasingly include networking and CPUs. TSMC's commonly cited AI growth metrics primarily capture GPU and ASIC-related revenue and do not fully reflect networking or the potential emergence of agentic-AI-driven CPU demand. As these additional categories scale, logic semiconductor growth could regain investor attention after a period when memory and other higher-beta AI beneficiaries dominated share price performance. The key question for the broader cycle is no longer simply whether GPU demand remains strong, but whether hyperscaler capex and AI model monetization can sustain the enormous infrastructure commitments now being discussed. Supply-chain visibility remains strong, although suppliers themselves acknowledge that long-term demand ultimately depends on hyperscalers achieving sufficient returns from AI investment.

See Appendix A-1 for Analyst Certification, Important Disclosures and Research Analyst Affiliations



TSMC: AI growth broadens beyond GPU/ASIC — TSMC remains central to nearly every major AI hardware transition. The most important point is that its AI opportunity is broader than the GPU/ASIC revenue definition investors typically focus on. Networking silicon, AI CPUs, silicon photonics and advanced packaging provide incremental growth vectors, while the underlying compute roadmap continues migrating rapidly toward leading-edge nodes. **We expect stronger 2027 growth for TSMC thanks to full-year contribution from 2nm, increased 3nm AI-chip volumes and potential pricing increases, creating another powerful combination of volume, mix and ASP expansion.** In terms of competitive landscape, we think Intel still needs substantial leading-edge capacity for its own CPUs and appears constrained in its ability to simultaneously support internal products and large external AI customers. In addition, Intel's 14A ecosystem and PDK/customer engagement process would require considerable qualification and design-cycle time before translating into meaningful external production. For critical AI products, customers also need predictable multiyear capacity commitments; TSMC's manufacturing scale and supply-chain coordination remain important competitive advantages. Nevertheless, we see potential opportunities for Intel and UMC's collaboration at 12nm, that could support products such as connectivity, display-driver or consumer-related chips, while its advanced packaging technology could also win specialized programs. However, those opportunities should be distinguished from a broad challenge to TSMC's leading-edge foundry position. Strategically, we think TSMC is becoming more selective about where it deploys capital and engineering resources. We believe front-end wafer manufacturing, SoIC and differentiated technologies such as silicon photonics remain priorities, while more conventional back-end operations are increasingly suitable for outsourcing. This suggests that continued AI growth can benefit TSMC directly through leading-edge wafer revenue while simultaneously creating substantial opportunities for OSAT partners.

TSMC outsourcing creates a structural opportunity for ASEH- Given the tight supply for both advanced node wafer and CoWoS, we think leading-edge front-end manufacturing generates the highest strategic value, and we believe TSMC has an incentive to prioritize wafer fabrication and differentiated packaging technologies such as SoIC while outsourcing larger portions of more back-end processing. CoWoS capacity can continue expanding, but TSMC's own back-end investment intensity may eventually moderate as qualified OSAT partners assume more steps. ASE/SPIL are already participating in substrate-related outsourced work and should gain additional opportunities as the ecosystem matures. Other Taiwanese packaging companies can also participate, while interposer production itself could become more distributed across external partners. Importantly, **ASEH** is not dependent solely on TSMC outsourcing. It is developing its own full-process advanced-packaging capabilities, including FoCoS, providing a more direct route to capture high-value AI packaging revenue. AMD is an important customer/catalyst for ASEH, in our view, particularly as CPU and heterogeneous-compute architectures require increasingly sophisticated packaging. ASEH's management's expectation that leading-edge advanced-packaging revenue can roughly double underscores the potential operating leverage. As chiplets, HBM, silicon photonics and increasingly complex interconnects proliferate, packaging content per device rises and the technical barriers to entry increase, potentially supporting structurally better margins for ASEH.

CPO is emerging as the next major AI infrastructure transition - Co-packaged optics (CPO) appears to be one of the clearest structural transitions in AI hardware because it addresses two increasingly severe constraints simultaneously: data transmission bandwidth and power consumption**. We expect Nvidia should be the initial catalyst, first through Spectrum-X CPO switches and subsequently**

through CPO integration associated with next-generation NVLink scale-up architectures and Rubin Ultra. Importantly, there are two different CPO adoption timelines that are often confused. CPO embedded directly alongside GPU/compute packages could remain two to three years away, but switch-level CPO is considerably closer and should begin scaling earlier. Moving optical connectivity closer to the switching silicon can materially reduce electrical I/O distance and power while supporting higher aggregate bandwidth. During Semicon Taiwan in early Sep, CPO coupling, optical alignment, fiber-array assembly and testing were among the most prominent themes. **TSMC** is developing its COUPE silicon-photonics/packaging platform, while a broader ecosystem is forming around optical engines, lasers, fiber-array units (FAUs), receptacles, precision alignment and test equipment. The most visible near-term bottlenecks appear to be external laser sources and fiber/FAU capacity, helping explain aggressive expansion by suppliers such as Lumentum and Corning. The transition could also redistribute semiconductor supply-chain value as CPO moves more value into advanced semiconductor packaging, silicon photonics, optical coupling and testing. We believe **Largan's** high precision automation capability will be able to catch CPO opportunity particularly in the most challenging optical fiber array alignment process.

Custom ASIC: MediaTek and GUC riding with Google - The Google TPU ecosystem illustrates both the opportunity and the execution complexity of custom AI silicon. We believe **MediaTek's** role should increase substantially over coming generations, MediaTek' s target market share is 15–20% initially but we see more upside ahead. With current TPU v8 generation, MediaTek's responsibility is concentrated more heavily in I/O, SerDes and wafer-handling-related activities; in the subsequent generation, its scope expands toward portions of the compute implementation as well. The SerDes configuration itself is also expected to increase, while advanced-node compute content materially lifts MediaTek's dollar value per device. Consequently, revenue can grow much faster than TPU units or market share alone would suggest. One execution risk is advanced packaging. A future Google/MediaTek program is expected to utilize Intel's EMIB-T technology, which is less proven at very high volume than TSMC's mature CoWoS ecosystem. Much of the difficult integration work, however, is distributed across substrate suppliers and OSAT partners. We think the initial yields and volumes could therefore be constrained, but the program has strategic importance for Google, Intel and MediaTek and is likely to proceed despite an early learning curve. And we believe its existing TPU generations should also continue shipping during the transition, cushioning the revenue impact. We believe Google's broader engagement with AMD and Marvell should not automatically be interpreted as additional TPU suppliers, in our view, Marvell has important networking and storage relationships with Google, while AMD opportunities appear more closely associated with CPUs. More broadly, the ASIC ecosystem remains structurally more fragmented than Nvidia's platform. Google is gradually internalizing more silicon expertise, but it still requires external design-service and implementation partners. We also see substantial growth of Google's CPU at GUC, which is fully COT model with back-end only service.

AI systems and power: Hon Hai benefits from Rubin, while 800VDC adoption should be gradual- At the system level, **Hon Hai** remains a major beneficiary of Nvidia's rack-scale transition. This year's growth has been driven largely by higher AI rack shipments, but next year adds a much higher ASP as the supply chain transitions from GB300 toward Vera Rubin. Higher rack value means revenue and profit can rise substantially even if unit growth moderates, provided margins remain broadly stable. Beyond Nvidia, Hon Hai is also positioning itself to manufacture custom-ASIC and hyperscaler systems, including CPU- and TPU-

related racks, which could begin contributing more meaningfully later next year. **Power architecture is another major debate, particularly around 800VDC**. The key distinction is that 800VDC is strategically attractive but not necessarily mandatory for Rubin or even every Rubin Ultra deployment. Rack power rises materially from the GB300 generation toward Rubin, requiring changes such as sidecar power shelves and more sophisticated power distribution. New greenfield AI data centers and neocloud deployments have stronger incentives to design around 800VDC because they can optimize the entire facility for future high-density racks. Existing facilities, however, can continue using alternative architectures, making adoption gradual rather than immediate. We believe **Delta** remains structurally well positioned in high-power data-center infrastructure, but expectations around the speed of 800VDC adoption may be slower than expected. The industry also faces a significant engineering learning curve as rack power density rises, including thermal management and potential overheating issues. Separately, shortages of power-management ICs and capacitors remain important constraints regardless of whether a data center uses 800VDC. Analog and power-semiconductor supply is supported by both AI-driven demand and relatively tight capacity at IDM suppliers and specialty foundries such as UMC and VIS.

Companies Mentioned:

Advanced Micro Devices (AMD.O; US$559.82; 1; 18 Sep 26; 16:00) | Alphabet Inc (GOOGL.O; US$349.54; 1; 18 Sep 26; 16:00) | ASE Technology Holding (3711.TW; NT$638.0; 1; 18 Sep 26; 13:30) | Corning Incorporated (GLW.N; US$150.13; 1; 18 Sep 26; 16:00) | Delta Electronics (2308.TW; NT$1735.0; 1; 18 Sep 26; 13:30) | GUC (3443.TW; NT$7150.0; 1H; 18 Sep 26; 13:30) | Hon Hai Precision (2317.TW; NT$250.5; 1; 18 Sep 26; 13:30) | Intel Corp (INTC.O; US$108.6; 1; 18 Sep 26; 16:00) | Largan Precision (3008.TW; NT$6430.0; 1; 18 Sep 26; 13:30) | Lumentum Holdings Inc (LITE.O; US$930.91; 1; 18 Sep 26; 16:00) | Marvell Technology Inc (MRVL.O; US$244.25; 1; 18 Sep 26; 16:00) | Mediatek (2454.TW; NT$4710.0; 1; 18 Sep 26; 13:30) | NVIDIA Corp (NVDA.O; US$222.27; 1; 18 Sep 26; 16:00) | TSMC (2330.TW; NT$2460.0; 1; 18 Sep 26; 13:30) | UMC (2303.TW; NT$156.0; 1; 18 Sep 26; 13:30) | Vanguard International Semi. (5347.TWO; NT$172.0; 1; 18 Sep 26; 15:00)