US Semiconductors and Semi Equipment

SemiBytes: Read Throughs from UBS Taiwan Summit, MATCH Act, MLPerf v6.1, Stock Action

Equities

Americas

Semiconductors

Timothy Arcuri
Analyst
timothy.arcuri@ubs.com
+1-415-352 5676

Natalia Winkler, CFA
Analyst
natalia.winkler@ubs.com
+1-415-352 4626

Gianmarco Vella
Associate Analyst
gianmarco.vella@ubs.com
+1-415-352 4555

Aaryan Wadhwa
Associate Analyst
aaryan.wadhwa@ubs.com
+1-212-821 6481

Seth Cho
Associate Analyst
seth.cho@ubs.com

Stacy Che
Associate Analyst
stacy.che@ubs.com
+1-212-713 1857

UBS Taiwan Summit - US Semis Takes

UBS Taiwan Summit took place last week (9/14-16) and we offer our takes across compute, memory, and semicaps. For **compute**: 1) Similar to our takeaways from the AI Infra Summit (here), optimization is increasingly happening at the system level as compute performance scales ~3x every two years versus only 1.6x for memory bandwidth and 1.4x for I/O, making memory and data movement increasingly important for extracting more performance out of each accelerator; 2) Networking is becoming a much more important part of AI compute as clusters scale to 100K+ GPUs/XPUs, which reads positively for AVGO/MRVL/ALAB as networking increasingly becomes part of the compute architecture itself, with more processing offloaded from CPUs to NICs/SmartNICs/DPUs, expanding so-called "XPU attach" opportunities and placing further focus on switching features; 3) Advanced packaging is more important than ever, but customer roadmaps are starting to diverge with TSMC targeting CoPoS 310x310mm process/equipment selection by mid-C2027 (mass production in C2028), while INTC's EMIB-T roadmap is coming to the fore as yields (for silicon bridge) are targeted to cross over 50% in 1H:27 and it supports >2m substrates for MediaTek TPU v9 demand; and 4) Advanced packaging supply continues to improve, with industry CoWoS capacity now expected to reach 260kwpm exiting C2027 versus 160kwpm exiting C2026, including 180kwpm at TSMC, 60kwpm at ASE and 20kwpm at Amkor; we understand AMD and NVDA are well positioned to capture the incremental capacity upside, supporting continued accelerator ramps into 2027. For **memory**: AI demand is driving a multi-year DRAM and NAND undersupply that is constraining server shipments, sharply increasing hardware BOMs, and accelerating capacity investment. DRAM fulfillment has fallen to ~60% of unconstrained demand (our industry check), while enterprise SSD fulfillment is only ~30%, making memory availability, rather than customer demand, the key constraint on server and storage shipments into 2027. For **semicaps**: This memory dynamic, together with advanced-node and packaging expansion, supports a broader semicap upcycle, in our view, with global WFE projected to reach ~$226B in 2027 and ~$275B in 2028 (our WFE update) and strong demand extending across lithography, wet processing, metrology, probing, and test.

MATCH Act – more bark than bite

The MATCH Act (Multilateral Alignment of Technology Controls on Hardware) has recently received more headlines surrounding its potential inclusion in broader legislation such as the National Defense Authorization Act (NDAA) and we have been a little surprised at some of our investor inbounds that reflect some hand-wringing around the issue (see UBS' note on ASML here). On 8/10, Senate sponsors pointed to the National Defense Authorization Act (NDAA, in short, the "Defense Budget") as an expected legislative vehicle for attaching the MATCH Act. First, based on our discussions with DC experts, it seems unlikely that the MATCH Act will be included in legislation and passed by the current Congress, though the topic could be taken up by the next Congress. The central objective of the MATCH Act has been to "level the playing field" between US semicap suppliers and foreign competitors (ASML/TEL/etc) - as ASML/TEL still retain scope to sell certain tools to customers such as YTMC where US companies face broad limitations. That said, our recent industry conversations suggest that both ASML and TEL have been by and large complying with restrictions implemented under existing US-Dutch/Japanese arrangements (under the Biden Administration). Additionally, our expert conversations pointed to the fact that putting additional restrictions into law could make it harder for the US government to negotiate with allies in the future. Taken together, and consistent with our 8/31 note, we continue to believe

This report has been prepared by UBS Securities LLC. **ANALYST CERTIFICATION AND REQUIRED DISCLOSURES, INCLUDING INFORMATION ON THE QUANTITATIVE RESEARCH REVIEW PUBLISHED BY UBS, BEGIN ON PAGE 5.**

there is a recurring pattern in headline risks for Semis ahead of Trump-Xi meetings, as both sides seek to enter negotiations from a position of strength. From the perspective of the US SPE suppliers, however, we think what can be done has been done – China is already pushing to source 50% of its tools domestically and local suppliers such as SMEE continue to make some progress, though the yields they can achieve remain a major question. Net, we think this is more "bark" than "bite" for US semicaps.

MLPerf v6.1 - NVDA introduces preliminary Vera Rubin results

Results from the latest iteration of MLPerf Inference v6.1 (released September 16th) featured a record 30 participants and the first peer-reviewed results for NVDA's Vera Rubin and AMD's Instinct MI350P. NVDA highlighted results from a 288-GPU GB300 NVL72 system that achieved ~2.0MM server tokens/s on DeepSeek-R1, while its 72-GPU GB300 configuration exceeded ~1.1MM server tokens/s on GPT-OSS-120B. NVDA also introduced preliminary Vera Rubin results, with a 72-GPU VR200 system delivering ~1.2MM server tokens/s on DeepSeek-R1. In our view, the results continue to demonstrate NVDA's systems and software advantage, with strong scaling across both Blackwell and Rubin architectures. AMD's headline result was a 512-GPU MI355X cluster developed with Crusoe that achieved ~2.4MM server tokens/s on DeepSeek-R1 and ~5.4MM server tokens/s on GPT-OSS-120B, the highest aggregate throughput reported in MLPerf to date. AMD also surpassed 950k server tokens/s on GPT-OSS-120B using a more commercially relevant 72-GPU MI355X configuration. We view AMD's MI355 results as encouraging, suggesting that it could increasingly be viewed as an option for inference deployments, although scaling current generation GPUs requires customized efforts. Looking ahead, we believe MI450 Helios should improve the ease of scaling AMD solutions into relatively larger clusters. INTC's participation remained focused on Xeon 6 CPUs and Arc Pro B70 GPUs, including new end-to-end RAG benchmarks, although INTC still appears to lack a competitive data center AI accelerator roadmap versus NVDA and AMD. We note that Google's TPU was again absent from MLPerf, with Google's submissions centered on NVDA GB200 systems. We believe this likely reflects the benchmark's focus on open models such as DeepSeek-R1 and GPT-OSS-120B rather than Gemini-based workloads.

Earnings up, multiples down

Our incoming investor queries reflect a greater degree of consternation around the group - particularly about whether expectations for semis have outstripped the availability of powered-shell capacity. We have seen some of this play out in stock performance, especially during this most recent earnings season. The percentage of companies across the semis landscape guiding above Street remains high (84% during this recent earnings season, albeit down from 97% 3 months ago - Figure 2) but stock action has been muted with guide-aboves being rewarded very little (Figure 3), while guide-belows actually penalized somewhat less than in CQ1 when sentiment was more bullish. Put differently, companies are still generally guiding above Street numbers, but investors are assigning less incremental value to those beats, suggesting that upside surprises are increasingly viewed as the baseline rather than a catalyst. We think this points to a materially higher bar for expectations - in a sense a "high-water mark" that was set for many stocks during the bullish period leading up to the July correction. Typically, lower stock prices are a reflection of lower expectations, but our conversations do not suggest any reduction in this pre-July "high-water mark" for many stocks across the semis landscape (semicaps are a prime example). Ergo, we remain quite bullish on fundamentals, but we do wonder whether stocks are in a period of consolidation where the market needs to be convinced that these expectations are achievable rather than rewarding incremental progress toward achieving them.

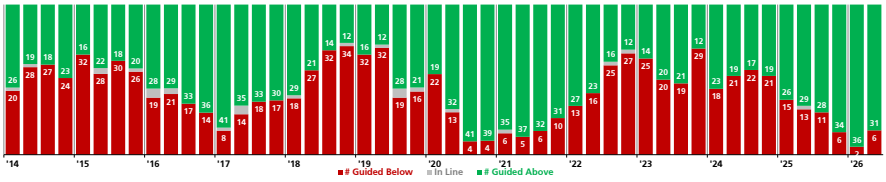


Figure 1: MLPerf v6.1 Table

Participants	Accelerator	# of Accelerators	DeepSeek R-1 Server '000 tokens/s	GPT-OSS-120B Server '000 tokens/s
AMD	AMD Instinct MI355X 288GB HBM3e	8	-	113
AMD	AMD Instinct MI355X 288GB HBM3e	72	-	964
Azure	NVIDIA GB200	288	1,599	-
Azure	NVIDIA GB200	72	427	-
Azure	NVIDIA GB300	72	589	-
CoreWeave	NVIDIA GB200	72	420	901
CoreWeave	NVIDIA GB300	72	-	1,160
Crusoe	AMD Instinct MI355X 288GB HBM3e	512	2,405	-
Crusoe	AMD Instinct MI355X 288GB HBM3e	512	-	5,393
Google	NVIDIA GB200	72	389	-
Nebius	NVIDIA GB300	72	603	1,122
Nebius	NVIDIA GB300	8	-	128
Nebius	NVIDIA VR200	36	591	-
NVIDIA	NVIDIA B300-SXM-270GB	8	60	101
NVIDIA	NVIDIA GB200	72	396	899
NVIDIA	NVIDIA GB300	288	2,028	-
NVIDIA	NVIDIA GB300	72	597	1,121
NVIDIA	NVIDIA GB300	8	68	128
NVIDIA	NVIDIA VR200	72	1,176	-
Oracle	NVIDIA GB300	72	587	1,093

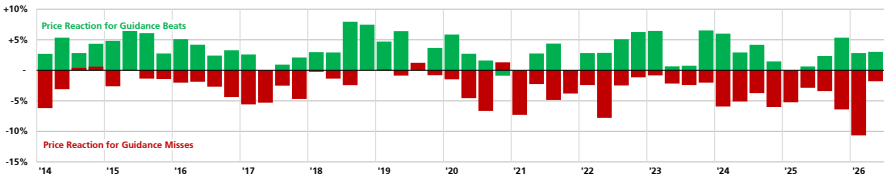
Source: MLCommons, UBS research

Figure 2: Companies' Guidance Above/Below Consensus



Source: Factset

Figure 3: Price Reaction Post Earnings



Source: Factset

Valuation Method and Risk Statement

We use various valuation techniques such as P/E, EV/FCF for valuing the companies in this report. Risk factors include but are not limited to macroeconomic factors such as a downturn in the economy, a disruption of international trade, technological disruption due to new inventions, or business model innovation whereby structural changes in the industry alter the future course of unit sales, ASPs, and revenues.

Advanced Micro Devices Inc:
We value AMD based on a NTM P/E methodology. Upside risks to our AMD thesis include higher-than-anticipated penetration into cloud/data center servers and better traction in AI. Downside risks include margin pressure in the form of aggressive price competition from Intel and an inability to effectively shift production to TSMC in a timely manner.
Astera Labs Inc:
We value ALAB based on a NTM P/E methodology. Upside risks to our ALAB thesis include a faster-than-expected production ramp of Scorpio X, faster-than-expected growth in retimers, and a slower-than-expected industry transition to CPO. Downside risks include a slower-than-expected production ramp of Scorpio X, slower-than-expected growth in retimers, and a faster-than-expected industry transition to CPO.
Broadcom Inc.:
AVGO: We use an EV/FCF methodology for valuation of AVGO. The major risks to our price target come primarily from the fact that AVGO might continue to do M&A and thereby affect FCF and dividend payout ratios that form the basis of our valuation. Should market conditions tighten they could impact the dividend yield and create more downside in our valuation.
Risk to upside and downside for our universe of semiconductor and semiconductor equipment stock include macroeconomic and geopolitical environments, M&A, market share loss/gain, currency fluctuations, tax policies, among others.
Marvell Technology Group:
MRVL: We use a PE multiple methodology for valuation. Cloud customers are very keen for another source beyond AVGO/CSCO, but it could take time for the next window to open. In storage, we may be under-estimating the internal controller efforts and capabilities at SSD suppliers.
NVIDIA Corp:
NVDA: Our price target is based on applying a P/E multiple. NVIDIA company-specific risks include: 1) competition from AMD in graphics processors (GPUs) and professional visualization products, which historically have driven the bulk of the company's gross profit, 2) intense competition in ARM-based applications processors, where rivals seek to leapfrog NVIDIA's early lead in dual-core technology, 3) emerging competition from Intel with its many-integrated core (MIC) processor family vs. the company's professional GPU computing products, where Intel's approach could be more flexible and easier to develop software. In addition, NVIDIA faces the semiconductor sector risk where revenue trends have historically been linked to corporate profitability, which could be challenged for robust growth, as the economic recovery has been more subdued vs. prior deep downturns.