Not for redistribution without written consent of Morgan Stanley

Morgan Stanley | RESEARCH

IDEA

September 10, 2026 05:26 PM GMT

Asia-Pacific Technology | Asia Pacific

AI Supply Chain: Prelim 2028 AI semis vs. CSP capex growth

AI semi capex, in particular for compute chips, should outgrow overall AI capex in 2028. We add cross-checks on ASICs after AVGO's guidance and the Qualcomm/AWS AI ASIC announcement.

Will 2028 be another growth year for AI semis? Our US internet analyst Brian Nowak published the Morgan Stanley AI Guidebook: Navigating the Age of Inference, and suggests 12% CSP capex growth into 2028, raising investor questions about AI semi growth momentum. Based on our latest supply chain checks, we introduce preliminary 2028 CoWoS/CoPoS/EMIB-T capacity assumptions. In aggregate, we still see 50% capacity growth in 2028 (see Exhibit 1), implying flattish infrastructure (e.g., the shell and AI DC) spending and moderate memory pricing in 2028.

Larger chip size and more applications (CPU, networking) to drive further growth in 2028: All these packaging technologies are 2.5D (compute and memory chips on side to side) for AI/HPC chips, including AI GPU, XPU, CPU, and networking.

•CoWoS (wafer based): We still see some CoW expansion at TSMC. For oS, TSMC may need more partners than ASE. Our checks suggest both Amkor and Xintec (non-covered) will diversify oS production in the US and Taiwan, respectively.

• CoPoS (panel-based): This is a panel-level process. After converting to 12-inch CoWoS wafer equivalents, we assume 5–10kwpm in 2028. The initial target project should be Nvidia's Feynman Ultra GPU.

• EMIB-T (substrate-based): Although a substrate-based technology, we convert it to 12-inch CoWoS wafer-equivalent capacity as well. One EMIB-T substrate can generate 16 pcs of a 9x reticle-size chip, around 4x the chips a CoWoS wafer can generate. Our checks indicate ~40k–45k wafer-equivalent EMIB-T capacity in 2028, producing at least 2mn of Humufish (9x reticle) substrates. 40k–45k wafer equivalents also imply 10-15% share in 2.5D packaging.

Read-across from recent ASIC guidance at Broadcom and the Qualcomm/AWS partnership

•According to Joe Moore's report: Broadcom Inc.: Earnings solid; raising FY 28 EPS, more specific FY27 revenue guidance is now US$115 bn, and that could double again in FY28. Our CoWoS supply-chain checks already see an AVGO AI lab customer booking ~15k–30k CoWoS wafers for 2027.

•Qualcomm announced its ASIC partnership with AWS, a key Alchip ASIC customer. However, the press release points to inference, different from Alchip's Trainium4, which covers both training and inference. We therefore see limited impact to Alchip.

MORGAN STANLEY TAIWAN LIMITED+
Charlie Chan
Equity Analyst
Charlie.Chan@morganstanley.com +886 2 7730 4135

Daniel Yen, CFA
Equity Analyst
Daniel.Yen@morganstanley.com +886 2 7730 4868

MORGAN STANLEY ASIA LIMITED+
Daisy Dai, CFA
Equity Analyst
Daisy.Dai@morganstanley.com +852 2848 5310

MORGAN STANLEY TAIWAN LIMITED+
Lucas Wang
Research Associate
Lucas.Yuhsi.Wang@morganstanley.com +886 2 2730 2875

MORGAN STANLEY ASIA LIMITED+
Ethan Jia
Research Associate
Ethan.Jia@morganstanley.com +852 3963 2187

Henry Zhao
Equity Analyst
Henry.Zhao@morganstanley.com +852 2239 7731

Morgan Stanley does and seeks to do business with companies covered in Morgan Stanley Research. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of Morgan Stanley Research. Investors should consider Morgan Stanley 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.

Introducing Our 2028 CoWoS/CoPoS/EMIB Capacity

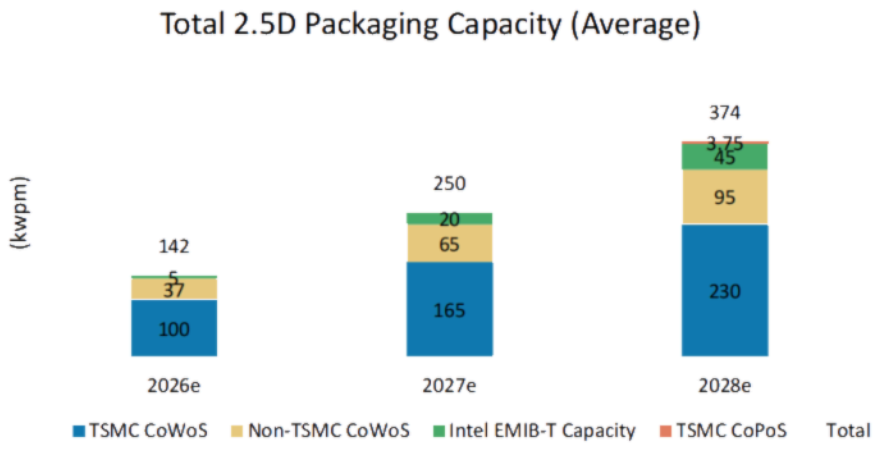
2028 preliminary CoWoS/CoPoS Capacity Build vs. our latest global CSP capex

We introduce preliminary 2028 CoWoS/CoPoS/EMIB-T capacity assumptions. In aggregate, we still see 50% capacity growth in 2028 (see Exhibit 1). All these packaging technologies are 2.5D for AI/HPC chips, including AI GPU, XPU, CPU, and networking.

Our US internet analyst Brian Nowak published Morgan Stanley AI Guidebook: Navigating the Age of Inference, and suggests 12% CSP capex growth into 2028, raising investor questions about AI semi growth momentum.

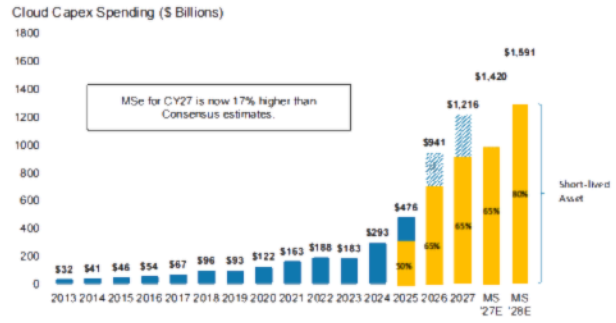
On ASIC industry, we continue to see a growing chance for MediaTek to win a second AI CSP or AI lab customer via its Nvidia NVLink Fusion partnership. The decision could be made in 4Q26, which could serve as a potential significant catalyst for MediaTek.

Exhibit 1: Total 2.5D packaging capacity from 2026-2028e (excluding EMIB-M which is mainly for CPU production)



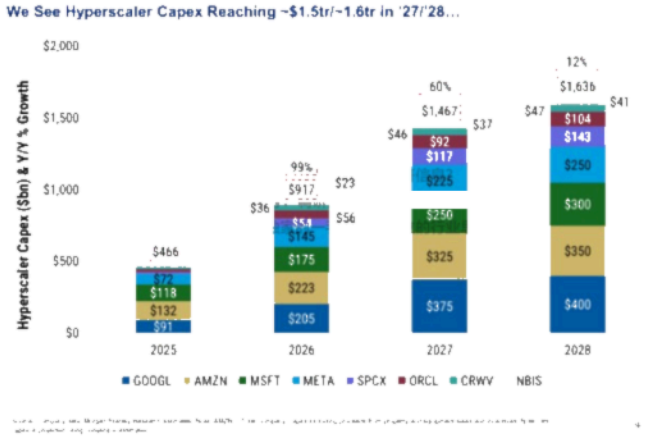
Source: Morgan Stanley Research estimates

Exhibit 2: Morgan Stanley cloud capex tracker: we expect the AI sever mix will outgrow infrastructure in 2028



Source: Morgan Stanley Research estimates

Exhibit 3: Hyperscaler Capex forecast from Morgan Stanley US research



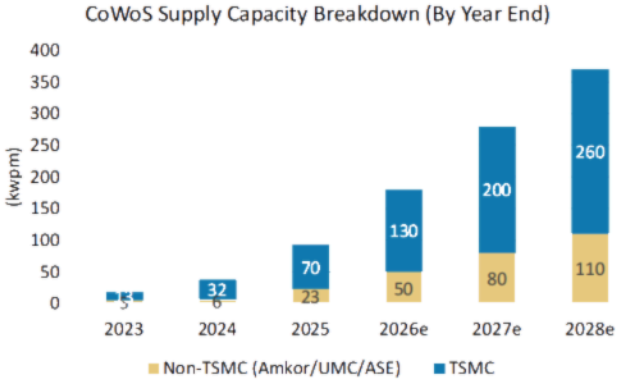
Source: Company Data, Morgan Stanley Research estimates

TSMC CoWoS and CoPoS

With continual booming demand in AI, we expect TSMC and its OSAT partners to expedite capacity expansion plans for CoWoS. We expect TSMC's CoWoS capacity to further build to 260kwpm in 2028 from 130/200kwpm in 2026/2027, with 2028's expansion mainly focusing in Arizona's AP9/10 and potentially in AP7. TSMC is now conducting trial production for CoPoS in Taoyuan. As long as the progress goes well, we expect TSMC to start building its CoPoS capacity, with the location also likely to be in AP9/10 or AP7P3.

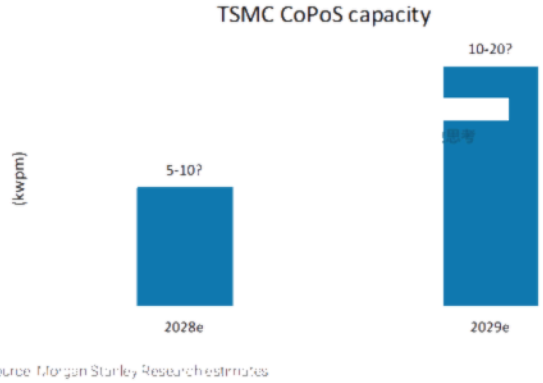
We expect the non-TSMC camp to expand CoWoS capacity to 110kwpm by 2028 year-end from 80kwpm by end-2027. For ASE/SPIL, we expect FoCoS+CoWoS capacity to rise from 30kwpm at 2027-end to 50-60kwpm at 2028-end, with expansion focused on CoWoS-L and CoWoS-R. For Amkor, we expect CoWoS capacity to increase from 30kwpm at 2027-end to 55-60kwpm at 2028-end, also focused on CoWoS-L and CoWoS-R.

Exhibit 4: We expect TSMC's CoWoS capacity to further build to 260kwpm in 2028, with OSAT (ASE/SPIL/Amkor) together building to 110kwpm in the same year



Source: Company data, Morgan Stanley Research estimates

Exhibit 5: We expect TSMC's CoPoS capacity to further build to 5-10kwpm in 2028 and 10-20kwpm in 2029, but this might depend on results from trial production



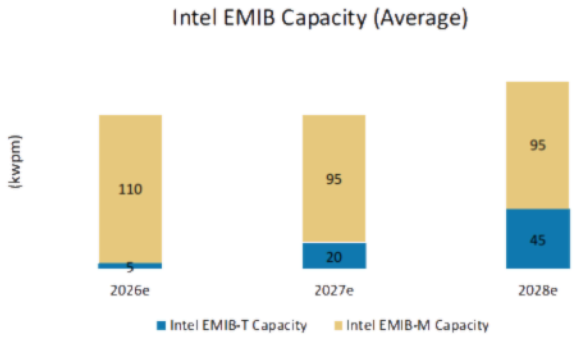
Source: Morgan Stanley Research estimates

How about Intel EMIB capacity build?

For Intel, we see current average EMIB-M capacity being 110kwpm, which would slightly decrease to 95kwpm in 2027 as it's converting some capacities toward EMIB-T. We see the major customers of EMIB-M being Trainium3, upcoming Trainium 4 and its in-house server CPU. Average EMIB-T capacity is around 5kwpm in 2026, which we expect to further scale to 15-20kwpm in 2027 and 40-45kwpm in 2028. All of the capacity mentioned above is converted for 12-inch wafer equivalent, which we believe would be much more easier to communicate and compared with TSMC CoWoS capacity.

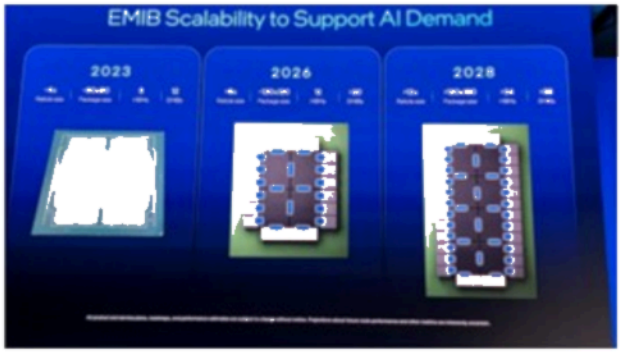
Although some small projects are ongoing, we still see MediaTek/Google Humufish consuming most of the capacities during 2027-28. We maintain our view that Intel EMIB-T could still produce close to 3mn units of Humufish from late 2027-2028, by assuming 4-5 chips per wafer of EMIB-T.

Exhibit 6: We expect Intel's EMIB-T capacity to reach 20kwpm in 2027 and 45kwpm in 2028; EMIB-M capacity is currently at 110kwpm, which would be down to 95kwpm due to conversion



Source: Morgan Stanley Research estimates

Exhibit 7: Intel's EMIB can easily support larger chips with more reticles (>12) if its supply chain can execute well

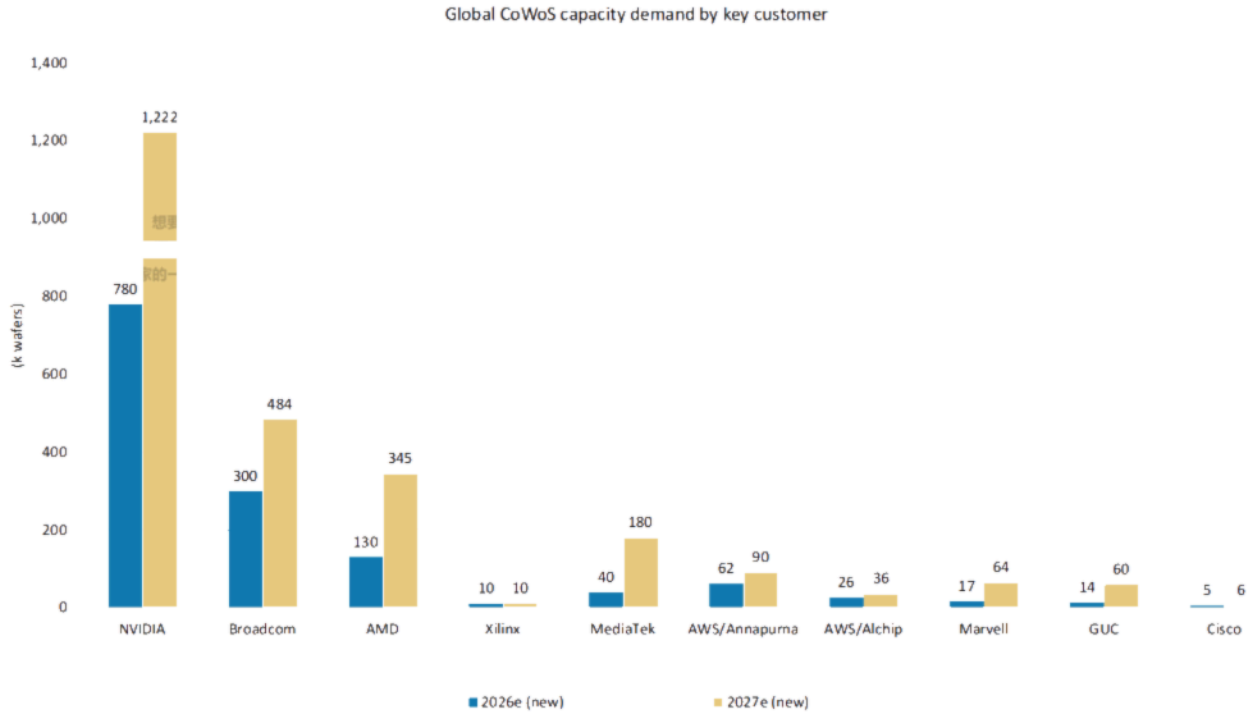


Source: Company Data, Morgan Stanley Research

2027 CoWoS allocation assumptions

Our recent industry checks suggest that the CoWoS capacity build of OSAT (both ASE and Amkor) is slower than our initial expectation and customers' forecasts. Therefore, we now see a slower ramp of AMD Venice CPU which should be the major customer of ASE and Amkor CoWoS capacity.

Exhibit 8: Global CoWoS demand breakdown: 2026e vs. 2027e



Source: Company data, Morgan Stanley Research (e) estimates, note estimates are compiled using our supply chain checks

Exhibit 9: Global CoWoS demand Y/Y growth profile

Y/Y	2023	2024	2025	2026e	2027e
NVIDIA	119%	280%	113%	84%	57%
Broadcom	56%	191%	25%	253%	61%
AMD	485%	470%	50%	117%	165%
Xilinx	63%	242%	0%	0%	0%
AWS/Annapurna				3%	45%
AWS/Alchip		71%	(69%)	420%	38%
Marvell	(22%)	1438%	(17%)	13%	276%
GUC		(15%)	300%	600%	329%
Cisco			36%	67%	10%
Others	23%	(49%)	50%	(33%)	20%
Total demand	**95%**	**218%**	**85%**	**102%**	**80%**

Source: Company data, Morgan Stanley Research (e) estimates, note estimates are compiled using our supply chain checks

CoWoS update (Exhibit 10):

- AMD's total 2027 CoWoS consumption increases 165% Y/Y to 345k:
  - We expect AI GPU-related production to be all at TSMC CoWoS, while MI455/450 would be the key focus for 2027, with minor production for MI500 series (chip code name Arcadia) in late 2027. Both MI455 and MI450 could reach around 500-650k units respectively. We estimate total TSMC CoWoS-L booking increases 200% Y/Y in 2027, to 210k.
  - We see the non-TSMC camp, including ASE/SPIL and Amkor, being mainly in charge of Venice CPU, high-end PC CPU, and gaming GPU production. Powertech was also planning for the production of Venice, but we now expect it to focus on gaming GPU only, given the lower yield rate of panel level packaging.
  - However, given the slower production ramp at OSAT, we believe TSMC would

also need to support around 80k CoWoS-L for Venice production, partially squeezing out the capacity of AI GPU build.
    - We now expect Venice CPU's CoWoS booking to increase from 50k in 2026 to 210k in 2027, implying 4.4mn units of CPU production, lower than our initial forecast of 5.6mn.
- Microsoft's Maia200 CoWoS booking increases from 1k in 2026e to 5k in 2027, we estimate:
    - We have been seeing upward revision demand of Microsoft Maia 200 into 2027; we believe the demand could reach around 5k of CoWoS in 2027, suggesting around 100-150k chips.
- US AI Lab CoWoS booking could reach 15-30k in 2027:
    - With the recent announcement of AI ASIC from one US AI lab, we expect its CoWoS booking at TSMC to reach 15-30k in 2027, with support from Broadcom.

Exhibit 10: Global CoWoS demand breakdown

(k wafer)	2023	2024	2025	2026e	2027e	2023	2024	2025e	2026e	2027e
NVIDIA	53	200	425	780	1,222	45%	54%	62%	56%	49%
TSMC			390	680	1,090					
CoWoS-L			390	650	910					
CoWoS-S			0	20	50					
CoWoS-R			0	10	130					
Non-TSMC			35	100	132					
Amkor			35	100	132					
CoWoS-S			20	20	12					
CoWoS-R			15	80	120					
ASE/SPIL			0	0	0					
CoWoS-S			0	0	0					
CoWoS-R			0	0	0					
Broadcom	23	68	85	300	484	20%	18%	12%	22%	19%
TSMC			83	260	420					
CoWoS-L			0	15	60					
CoWoS-S			83	245	360					
ASE/SPIL			2	30	40					
CoWoS-L			0	0	0					
CoWoS-S			2	30	40					
Amkor				10	24					
CoWoS-S				10	24					
AMD	7	40	60	130	345	6%	11%	9%	9%	14%
TSMC			60	80	215					
CoWoS-L			0	70	210					
CoWoS-S			60	10	5					
ASE/SPIL			0	50	110					
CoWoS-L			0	50	110					
Amkor				0	20					
CoWoS-L				0	20					
Xilinx	3	10	10	10	10	3%	3%	1%	1%	0%
MediaTek				40	180				3%	7%
TSMC				40	180					
CoWoS-S				40	180					
AWS/Annapurna			60	62	90				4%	4%
TSMC			60	32	54					
CoWoS-R			60	32	54					
ASE/SPIL				30	36					
CoWoS-R				30	36					
AWS/Alchip	9	16	5	26	36	8%	4%	1%	2%	1%
Intel Habana	0	7	9	0	0	0%	2%	1%	0%	0%
Marvell	1	18	15	17	64	1%	5%	2%	1%	3%
TSMC				5	50					
CoWoS-L				5	50					
CoWoS-R			15	0	0					
ASE/SPIL				12	14					
CoWoS-R				12	14					
GUC	1	1	2	14	60	1%	0%	0%	1%	2%
Cisco		2	3	5	6		1%	0%	0%	0%
Others	20	10	15	10	12	17%	3%	2%		
Total demand	**117**	**372**	**689**	**1,394**	**2,509**	**100%**	**100%**	**100%**		

Source: Company data, Morgan Stanley Research (e) estimates; note, estimates are compiled using our supply chain checks



Exhibit 11: AI HBM consumption: up to 44bn Gb in 2027

AI chip vendor	Product name	CoWoS capacity allocation (k wafers)	Chips per CoWoS wafer	Implied shipments (k)	HBM chip density (GB)	HBM chip units	Total HBM size (GB)	HBM generation	HBM vendor	Total HBM demand (k GB)
AI GPU (2027e)										
NVIDIA	B300	40	14	560	36	8	288	HBM3e 12hi	Hynix/Micron	161,280
	Vera CPU	250	23	5,750						
	Rubin R200	740	8	5,920	36	8	288	HBM4 12hi	Hynix/Micron/Samsung	1,704,960
	Rubin Ultra	130	8	1,040	48	8	384	HBM4e 12hi	Hynix/Micron/Samsung	399,360
AMD	MI350 series	5	12	60	36	8	288	HBM3e 12hi	Samsung/Micron	17,280
	MI455	75	7	525	36	12	432	HBM4 12hi	Samsung/Micron	226,800
	MI450 (for Meta)	45	14	630	36	6	216	HBM4 12hi	Samsung/Micron	136,080
	MI500 (Arcadia)	10	4	40	48	16	768	HBM4e 12hi	Samsung/Micron	30,720
	Venice CPU	210	21	4,410						
AI ASIC (2027e)										
Google	TPU v7p (Ironwood; AVGO)	13	16	208	24	8	192	HBM3e 8hi	Hynix/Samsung	39,936
	TPU v8i (Sunfish; AVGO)	330	12	3,960	36	8	288	HBM3e 12hi	Samsung/Hynix/Micron	1,140,480
	TPU v8t (Zebrafish; MediaTek)	180	20	3,600	36	6	216	HBM3e 12hi	Samsung/Hynix/Micron	777,600
	TPU v9 (Humufish; MediaTek)			400	48	12	576	HBM4e 12hi	Samsung/Hynix/Micron	230,400
AWS	Trainium 3	140	17	2,380	36	4	144	HBM3e 12hi	Hynix/Samsung/Micron	342,720
Microsoft	Maia 200	5	29	145	16	4	64	HBM3	Samsung	9,280
	Maia 300	50	11	550	36	8	288	HBM4 12hi	Samsung	158,400
Meta	MTIA 3 (Iris)	55	10	550	36	8	288	HBM3e 12hi	Samsung/Hynix/Micron	158,400
US AI Lab	Nexus (Titan 1)	25	13	325	36	6	216	HBM4 12hi	Samsung/Hynix/Micron	70,200
	Titan 2	5	8	40	36	8	288	HBM4 12hi	Samsung/Hynix/Micron	11,520
Others		18								
Total		**2,454**								**5,615,416**
								Total HBM demand (mn Gb)		**44,923**

Source: Company data, Morgan Stanley Research (e) estimates. Note: Est mates are compiled using our Asian supply chain checks

Exhibit 12: AI wafer consumption: at least US$55bn in 2027

AI chip vendor	Product name	CoWoS capacity allocation (k wafers)	Chips per CoWoS wafer	Implied shipments (k)	Compute die size	Geometry	Compute die units	Wafer consumption (k wafers)	Wafer price (US$)	Wafer revenue TAM (US$ mn)
AI GPU (2027e)										
NVIDIA	B300	40	14	560	850	4nm	2	44	23,042	1,024
	Vera CPU	250	23	5,750	850	3nm	1	228	27,300	6,229
	Rubin R200	740	8	5,920	850	3nm	2	470	27,300	12,827
	Rubin Ultra	130	8	1,040	850	3nm	2	83	27,300	2,253
AMD	MI350 series	5	12	60	110	3nm	8	2	27,300	48
	MI455	75	7	525	850	2nm	2	56	30,000	1,667
	MI450 (for Meta)	45	14	630	850	2nm	1	33	30,000	1,000
	MI500 (Arcadia)	10	4	40	850	2nm	2	4	30,000	127
	Venice CPU	210	21	4,410	155	2nm	8	176	30,000	5,292
AI ASIC (2027e)										
Google	TPU v7p (Ironwood; AVGO)	13	16	208	700	3nm	2	14	27,300	371
	TPU v8i (Sunfish; AVGO)	330	12	3,960	800	3nm	2	296	27,300	8,075
	TPU v8t (Zebrafish; MediaTek)	180	20	3,600	800	3nm	2	269	27,300	7,341
	TPU v9 (Humufish; MediaTek)			400	850	2nm	4	63	30,000	1,905
AWS	Trainium 3	140	17	2,380	700	3nm	2	127	27,300	3,465
Microsoft	Maia 200	5	29	145	700	3nm	1	3.8	27,300	103
	Maia 300	50	11	550	850	2nm	1	29.1	30,000	873
Meta	MTIA 3 (Iris)	55	10	550	850	3nm	2	44	27,300	1,192
US AI Lab	Nexus (Titan 1)	25	13	325	750	3nm	4	45.5	26,000	1,183
	Titan 2	5	8	40	850	2nm	4	8.5	30,000	254
Others		18								
Total		**2,454**						**1,996**		**55,230**

Source: Company data, Morgan Stanley Research (e) estimates. Note: Est mates are compiled using our Asian supply chain checks

Exhibit 13: AI HBM consumption: up to 29bn Gb in 2026

AI chip vendor	Product name	CoWoS capacity allocation (k wafers)	Chips per CoWoS wafer	Implied shipments (k)	HBM chip density (GB)	HBM chip units	Total HBM size (GB)	HBM generation	HBM vendor	Total HBM demand (k GB)
AI GPU (2026e)										
	B300	390	14	5,460	36	8	288	HBM3e 12hi	Hynix/Micron/Samsung	1,572,480
	Vera CPU	90	23	2,070						
NVIDIA	Spectrum/CPX	60		-						
	Rubin R200	260	8	2,080	36	8	288	HBM4 12hi	Hynix/Micron/Samsung?	599,040
	H200	20	27	540	24	6	141	HBM3e 8hi	Hynix	76,140
	MI300	3	12	36	24	8	192	HBM3e 12hi	Samsung	6,912
AMD	MI350/375	7	12	84	36	8	288	HBM3e 12hi	Samsung/Micron	24,192
	MI455	65	7	455	36	12	432	HBM4 12hi	Samsung/Micron	196,560
	Venice	50	21	1,050						
AI ASIC (2026e)	独家的一手信息，独立的行业思考						独家的一手信息，独立的行业思考			
	TPU v7p (Ironwood; AVGO)	145	16	2,320	24	8	192	HBM3e 8hi	Hynix/Samsung	445,440
Google	TPU v8i (Sunfish; AVGO)	80	12	960	36	8	288	HBM3e 12hi	Hynix/Samsung/Micron	276,480
	TPU v8t (Zebrafish; MediaTek)	40	20	800	36	6	216	HBM3e 12hi	Hynix/Micron	172,800
	Tomhawk 5/6	35								
AWS	Trainium 2.5	20	16	320	24	4	96	HBM3e 8hi	Hynix/Samsung/Micron	30,720
	Trainium 3	100	17	1,700	36	4	144	HBM3e 12hi	Hynix/Samsung/Micron	244,800
GUC's customers		10	20	200	24	6	144	HBM3e 8hi	Samsung?	1,152
Microsoft	Maia 200	1	29	29	16	4	64	HBM3	Samsung	1,856
	Maia 300	5	11	55	36	8	288	HBM4 12hi	Samsung	15,840
Meta	MTIA 3 (Iris)	15	10	150	36	8	288	HBM3e 12hi	Hynix/Samsung	43,200
Open AI	Nexus (Titan 1)	10	13	130	36	6	216	HBM4	Hynix	28,080
Others		15								
Total		**1,421**								**3,735,692**
									Total HBM demand (mn Gb)	**29,886**

Source: Company data, Morgan Stanley Research (e) estimates. *Note: Estimates are compiled using our Asian supply chain checks

Exhibit 14: AI wafer consumption: at least US$27bn in 2026

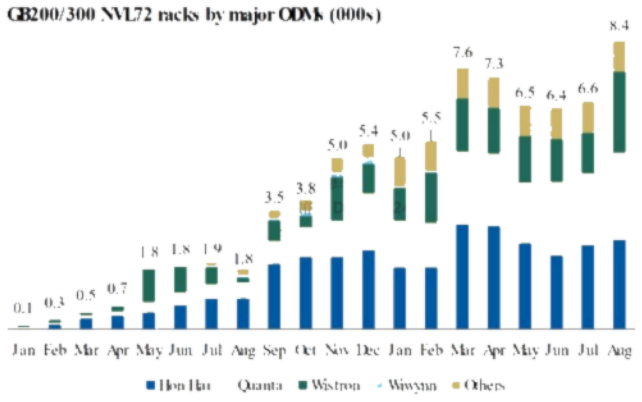
AI chip vendor	Product name	CoWoS capacity allocation (k wafers)	Chips per CoWoS wafer	Implied shipments (k)	Compute die size	Geometry	Compute die units	Wafer consumption (k wafers)	Wafer price (US$)	Wafer revenue TAM (US$ mn)
AI GPU (2026e)										
	B300	390	14	5,460	850	4nm	2	433	21,945	9,510
	Vera CPU	90	23	2,070		3nm				
NVIDIA	Spectrum/CPX	60		-						
	Rubin R200	260	8	2,080	850	3nm	2	165	26,000	4,292
	H200	20	27	540	814	4nm	1	15	21,945	331
	MI300	3	12	36	110	5nm	8	1	18,000	19
AMD	MI350/375	7	12	84	110	3nm	8	2	26,000	64
	MI400	65	10	650	110	2nm	8	32	28,125	886
	Venice	50	25	1,250						
AI ASIC (2026e)										
	TPU v7p (Ironwood; AVGO)	145	16	2,320	700	3nm	2	152	26,000	3,942
Google	TPU v8i (Sunfish; AVGO)	80	12	960	800	3nm	2	72	26,000	1,864
	TPU v8t (Zebrafish; MediaTek)	40	20	800	800	3nm	2	60	26,000	1,554
AWS	Trainium 3	100	17	1,700	700	3nm	2	91	26,000	2,357
GUC's customers		10	20	200	645	4nm	1	5	21,945	107
Microsoft	Maia 200	4	29	116	700	3nm	1	3.0	26,000	79
	Maia 300	5	11	55	850	2nm	1	2.9	28,125	82
Meta	MTIA 3 (Iris)	15	10	150	850	3nm	2	11.9	26,000	310
Total		**1,424**						**1,075**		**26,107**

Source: Company data, Morgan Stanley Research (e) estimates. *Note: Estimates are compiled using our Asian supply chain checks

Connecting Nvidia GPU shipments and rack shipments:

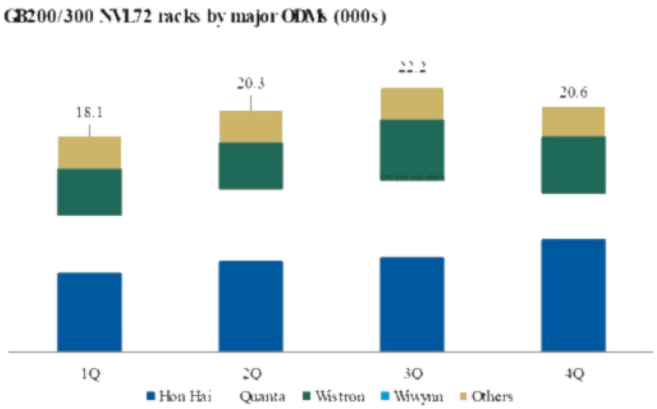
In Howard Kao's rack shipment update, Greater China Technology Hardware: GB200 and GB300 NVL72 Racks in August 2026 (9 Sep 2026), August racks recovered, implying more Blackwell GPU consumption. His CY26 full-year rack forecast is now at 80-85K racks. This is similar to our observation that KYEC continues final testing for meaningful volumes of B-series GPUs in 3Q26.

Exhibit 15: Industry-wide GB200/300 NVL72-equivalent monthly rack output, by major ODMs (2025-2026)



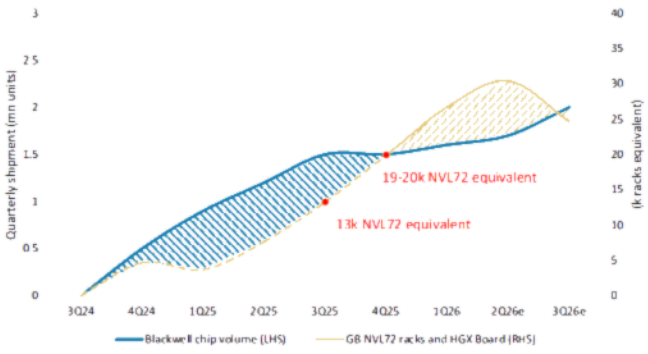
Source: Company data, Morgan Stanley Research estimates.

Exhibit 16: Industry-wide GB200/300 NVL72-equivalent monthly rack output, by major ODMs (2026)



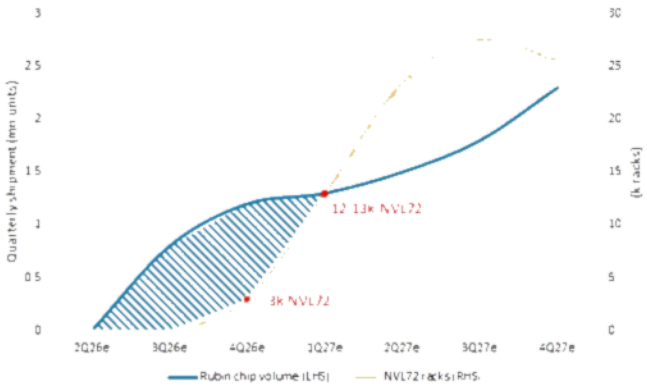
Source: Company data, Morgan Stanley Research estimates

Exhibit 17: Blackwell chip vs. rack shipment



Source: Morgan Stanley Research estimates

Exhibit 18: Rubin chip vs. rack shipment



Source: Morgan Stanley Research estimates

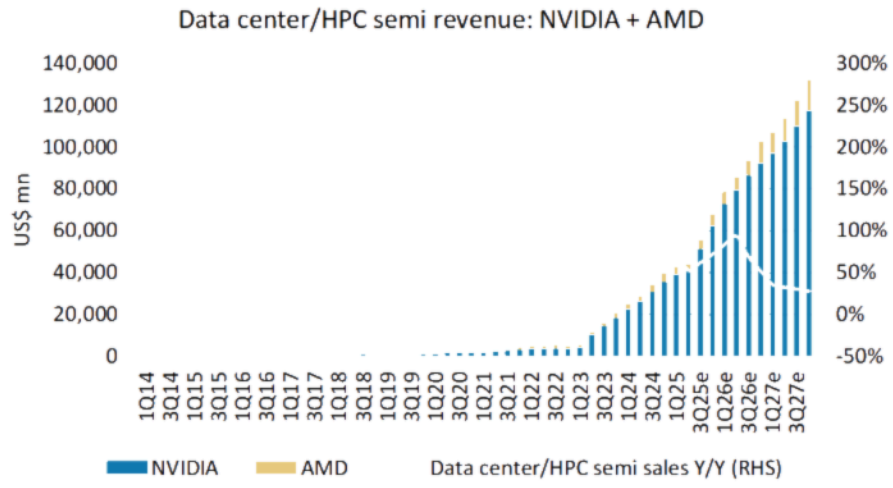
AI semis – stock implications, P/E multiples, revenue exposure

Exhibit 19: P/E multiple trend of AI semis



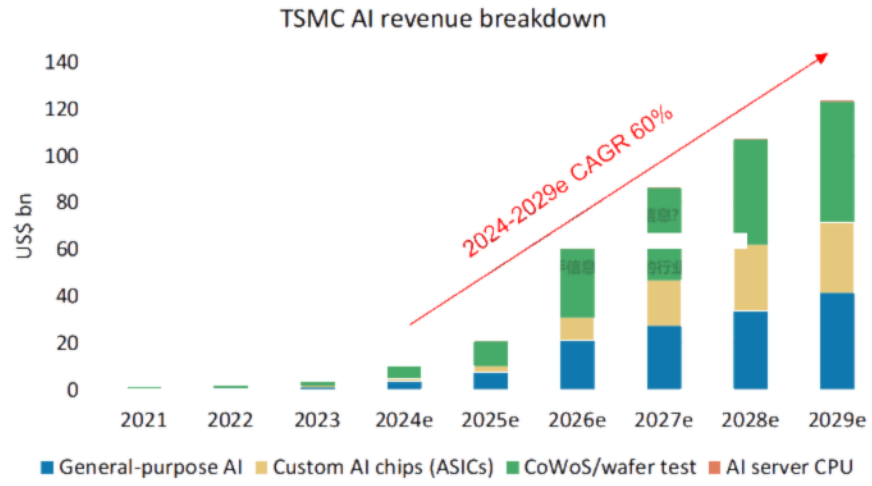
Note: Alternative AI semis group: AMD, Alchip, Andes, Marvell, Broadcom. AI semi enablers group: TSMC, Synopsys, Cadence, ASML, BESI, Ibiden, KYEC, Advantest. Source: Company data, Morgan Stanley Research

Exhibit 20: We still expect AI chip revenue to rise QoQ



Source: Company data, Refinitiv, Morgan Stanley Research (e) estimates

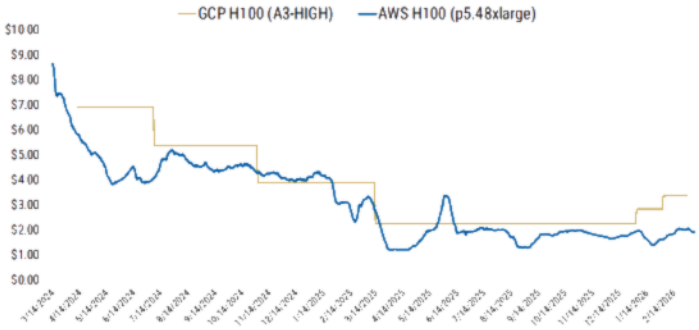
Exhibit 21: TSMC's AI-related revenue 2024-29 CAGR could reach 60%



Source: Company data, Morgan Stanley Research (e) estimates

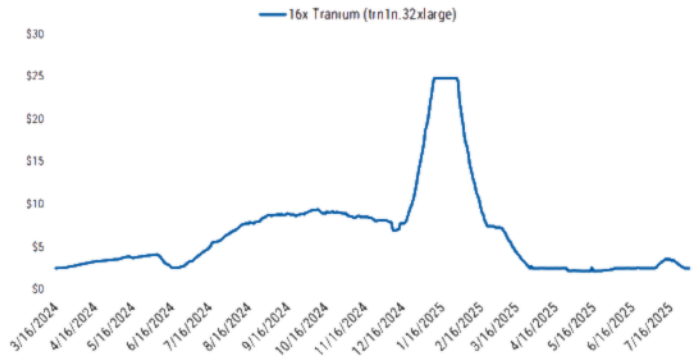
AI GPU and ASIC rental price tracker

Exhibit 22: AI GPU H100 per GPU per hour as of end-March



Source: Company data, Morgan Stanley Research

Exhibit 23: AI ASIC equivalent computing power – 16x Inferentia 2 per hour



Source: Company data, Morgan Stanley Research

Key Featured Reports on the AI Supply Chain

Asia-Pacific Technology: AI Supply Chain: Google ASIC, who gets what for design services? (24 Aug 2026)

Asia-Pacific Technology: AI Supply Chain: GPU Architecture Optimization Given Memory Shortage and Supernode Challenge (10 Aug 2026)

Asia-Pacific Technology: AI Supply Chain: Further Updates on 2027 CoWoS Allocation and ASIC Dynamics (8 Jul 2026)

Asia-Pacific Technology: AI Supply Chain: Preliminary 2027 TSMC CoWoS Allocation (23 Jun 2026)

Asia-Pacific Technology: AI Supply Chain: Cloud Capex Strength; More CoWoS Allocation for TPU (7 May 2026)

Asia-Pacific Technology: AI Supply Chain: Addressing Questions on Nvidia GPU/LPU and Google TPU (9 Apr 2026)

Greater China Semiconductors: AI Supply Chain Tracker: Key Investor Feedback from GTC/OFC (23 Mar 2026)

Asia Technology: AI Supply Chain: CPO and ASIC Dynamic Update (3 Mar 2026)

Asia-Pacific Technology: AI Supply Chain: CES implications, ASIC production, China AI chips (6 Jan 2026)

Asia-Pacific Technology: AI Supply Chain: TSMC CoWoS Expansion; ASIC Dynamics; Asia Field Trip (1 Dec 2025)

Foundation

Technology: Rise of the AI Agent – Global Implications (19 Apr 2026)

Global Technology: Supply-chain Reorientation (24 Jul 2025)

Global Technology: China – AI: The Sleeping Giant Awakens (13 May 2025)

Global Technology: AI Cloud Capex in the Spotlight (26 Feb 2025)

Global Semiconductors: AI ASIC 2.0: Potential winners (15 Dec 2024)

Global Technology: Global Technology – Dawn of the AI Smartphone Era: Edge AI – Apple Intelligence Fuels Innovation – More Charts, Fewer Words (17 Jul 2024)

Global Technology: AI PCs To Usher In The Next Leg Of PC Market Growth (21 May, 2024)

Key upstream AI supply chain companies

Tracking China's Semi Localization: Raising China AI GPU TAM on recent geopolitical dynamics (22 Jun 2026)

Greater China Technology Semiconductors: MCU: Supply tightness may drive prices up further (21 Jun 2026)

Greater China Semiconductors: Power Semis – Supply Driven Upcycle (18 Jun 2026)

Greater China Technology Semiconductors: Thoughts on Investor Concerns Regarding CPO (10 Jun 2026)

WinWay Technology Co Ltd: May sales came in better than feared; OW (9 Jun 2026)

MPI Corporation: Strong ASIC Ramp-up Drives May Sales; OW (9 Jun 2026)