US Semiconductors

$210B Server CPU TAM By C2030; Gated By Memory Supply

Equities

Americas

Semiconductors

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

Seth Cho
Associate Analyst
seth.cho@ubs.com

Summary

We have grown more bullish on server CPU demand signals following a number of industry events and new supply-chain/practitioner checks over the past few weeks and see demand levels considerably higher than our prior $175B CPU TAM estimate. That said, we think memory supply will prevent the market from being a whole lot bigger than ~$210B. This note examines the issue from both the agentic CPU demand side and the memory supply side. While AMD, INTC, ARM, and others are clear beneficiaries, against current expectations, we continue to believe MU is the best way to play server CPU demand.

Server CPU demand could be >$500B in C2030, but memory will constrain the TAM

We are not huge believers that META Muse adds that much to demand (single digit percentage), but it is very important as evidence of broader proliferation of agentic workloads beyond coding. More importantly, practitioners are increasingly pointing to significant portions of enterprise and consumer workloads where the CPU-to-GPU ratio could settle at multiple CPUs per GPU, creating incremental opportunities for CPUs optimized for AI cloud-native applications and broadening the opportunity set for products such as INTC's efficiency-core (E-core)-based Clearwater Forest, currently ramping on 18A. Our own analysis suggests the "unconstrained" CPU demand could reach $500B+ (easily 3-4x our new ~$210B C2030 TAM estimate) but memory will remain the key bottleneck over the period. Our expert conversations suggest that near-term use of CXL to repurpose older DDR4 memory may provide some relief, but may still only expand the memory available for CPUs by a single-digit percentage. Net, we are raising our C2030 CPU TAM estimate from ~$175B to ~$210B, driven primarily by higher head-node and standalone AI CPU units. Even against this stronger demand backdrop, we still have a hard time seeing CPU attach increasing much beyond 1:1 in C2030 because there is simply not enough memory to support practitioner ambitions. While the stronger demand backdrop is broadly supportive for unit growth across the CPU ecosystem, companies are already citing TAM numbers in this range or even higher, though expectations for memory do not seem to appreciate the impact of server CPU demand. Net, near term, we continue to see greater EPS leverage for AMD and INTC that monetize the full processor at roughly $3-4K per CPU, compared to ARM that sees IP-based economics with ~$100+/CPU.

CPU-to-GPU attach points to significant upside, but memory ultimately governs the TAM

While AI engineers aspire to increase the CPU-to-GPU attach rate for complex agentic workloads, we see memory ultimately setting a limit on how much theoretical demand can translate into units. Our recent checks increasingly suggest the ideal CPU-to-GPU ratio for AI infrastructure could settle closer to 3:1 and potentially 4:1, albeit with meaningful variation by workload, as internet/consumer applications generally require fewer CPUs while enterprise/agentic workloads require more. Against our expectation for ~44MM accelerators in C2030, a 3-4x ratio would theoretically imply ~120MM CPUs in that "aspirational" TAM supporting AI/agentic infrastructure alone, before adding ~30MM traditional server CPUs. However, we do not think the industry can realistically support that level of CPU deployment given the corresponding memory requirement. Traditional CPUs today use ~0.5TB of memory on average, while our analysis suggests AI CPUs require ~1.5TB today and potentially 2-3TB in coming generations. Our estimates suggest this unconstrained/aspirational TAM would require 282EB of DRAM, 3.5x

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 28.**

higher compared to current industry estimates of 77EB based on UBSe bottom-up memory estimates. CXL could provide some incremental relief by enabling reuse of older-generation DRAM, but our expert conversations suggest this would add only high-single digits to available memory supply. Net, we think the CPU-to-GPU ratio increasingly supports a substantially larger theoretical CPU opportunity, but memory, rather than compute demand, becomes the binding constraint on how quickly that opportunity can be realized.

Muse is a positive CPU demand signal, but probably not the needle mover the headline math suggests

Muse created plenty of agitation when it was released, but we view it more as a validation of the CPU demand thesis than a standalone TAM needle mover. With daily active users scaling into the millions, our analysis suggests that supporting 1B active users (vs. ~2.5MM DAUs now) on multi-threaded CPUs could theoretically require as many as ~16MM CPUs if all users were active concurrently and if META provisioned infrastructure against 100% of potential demand. However, our expert conversations suggest cloud businesses typically provision for only ~5-10% of expected peak demand, which would put practical requirements closer to ~1MM units, versus ~23MM total CPUs shipped in C2025. Moreover, because many of these services are offered for free, more flexible SLAs allow workloads to be queued or deferred rather than requiring infrastructure to be provisioned for full concurrency. Net, we see Muse as another strong data point supporting the broader case for structurally higher CPU demand, but we are less convinced that Muse itself is a material needle mover for industry units.

Read-throughs

CPU demand upside would likely be positive for: 1) DRAM (DDR) demand, although supply will not catch up with theoretical demand, de-specing may occur. 2) IC substrates, specifically ABF. 3) Leading-edge foundry: TSMC. 4) PCBs. 5) MLCCs: Murata +SEMCO have ~80% server MLCC share. 6) Semicap, notably WFE and Testing.

THE LATEST DATA POINTS ON AGENTIC AI CPU DEMAND

Since publishing our initial CPU TAM analysis (here) and subsequent share analysis (here), we have conducted another round of channel checks and see the engineering community increasingly coalescing around a higher CPU-to-GPU attach ratio. In addition, META's announcement of Muse once again puts a spotlight on agentic CPU demand.

Two key takeaways from our perspective:

- CPU-to-GPU attach is moving from 0.25-to-1 to multiple CPUs per GPU
- META's Muse is putting a spotlight on consumer agentic applications, but we don't see it moving the needle from the standpoint of total CPU demand.

Multiple CPUs Per GPU Ratio

Our expert conversations with CPU practitioners lead us to believe that there is a consensus forming on two topics. First, although the exact CPU-to-GPU ratio will vary by workload and application, it appears that more use cases land around a target of 3-4 CPUs per GPU, inclusive of the 0.5-1 CPU that serves as a head node, suggesting ~2-3 CPUs could be deployed in a standalone CPU rack configuration. Beyond that, we are also picking up on enterprise agentic workloads being inherently more CPU demanding, while consumer agentic applications require relatively lower CPU-to-GPU ratios.

In a recent Dell white paper, the company analyzed CPU intensity across different workloads, providing another data point in this direction, and our follow-up conversations with practitioners suggest the ratio varies considerably by workload. According to the findings, the measured CPU-to-GPU demand could be ~1:1 for simple model serving and short tool calls and could scale to 7:1 for heavy code processing, while HPC-type applications could require even more cores.

Our discussions with experts left us to believe the majority of use cases served by consumer agentic applications would fall into the first/lighter-CPU category with a portion requiring more complex use cases. We expect little to no HPC-like workloads in serving a consumer agentic application.

Layering this framework over Dell's estimates, we arrive at a ~3:1 CPU-to-GPU ratio on the current CPU generation. Importantly, when converting the math to core counts, the ratio is closer to 400 cores:1 GPU. What we find interesting is how this core-per-GPU could scale from here - although increasing CPU cores generation over generation could support a theoretically lower total CPU attach ratio, experts we speak to suggest that increasing GPU/XPU performance will likely result in the core ratio also trailing higher over time.

Figure 1: CPU requirements vary by different workloads

Workload tested	What the agent does	Workload CPU requirements	CPU quota per agent	Sustained agents on 2x AMD 64-core Turin CPUs*	Measured CPU:GPU demand	of Consumer Agentic workload (UBSe)	Measured CPU Cores:GPU demand
Workload 1: pypi-server	Writes a Python package and serves it from a local package index	Light CPU: mostly model/code generation with short tool calls	1 core	512	1:1	65%	128:1
Workload 2: sqlite-with-gcov	Compiles SQLite with code-coverage instrumentation and fixes build errors	Bursty CPU: compiler runs substantially between model reasoning turns	1 core	238	7:1	30%	896:1
Workload 3: mcmc-sampling-stan	Writes a statistical model, compiles it, then runs computationally heavy sampling	Heavy CPU/HPC: sustained multi-core numerical computation	4 cores	60	18:1	<5%	2304:1
Blended CPU ratio				**CPU:GPU**	**4:1**		
				CPU Cores:GPU			**421:1**

Source: UBS Research, Dell "Characterizing CPU Demand for Agentic AI: Three Agentic Workloads on the Dell PowerEdge™ R7725", 14-Sep-26 *R7725 system supports 2 AMD Turin CPUs with 64 cores each

Net, while the precise ratio will ultimately depend on workload mix, the demand indicators we are picking up on suggest a meaningful increase in the standalone CPU opportunity versus our prior framework, where we assumed only ~0.33 standalone CPUs per GPU plus ~0.5 CPU for head nodes, or ~0.85 CPUs per GPU in total. This higher attach assumption is one of the primary reasons we are becoming more constructive on the longer-term CPU TAM.

So on a unit basis, we can see CPU-to-GPU ratio currently shipping in volume generation that is roughly 1-to-4, but if we were to account for core count scaling, we think that math would suggest a lower ratio over time as CPU core counts scale. In the past we have observed CPU roadmaps following ~3-4 year cycles to double CPU core counts, which would suggest that we would be approaching ~500 cores/CPU by the end of the decade. At the same time, accelerator performance does not stand still, so we ultimately think the CPU-to-GPU ratio would likely subside but not decline by half - hence, we use a 3:1 CPU-to-GPU ratio in our demand estimate.

Muse math overstated, although we view it as a positive signal

Although we are encouraged by META Muse's traction, with the application reaching ~2.5MM daily active users weeks after release, we believe some of the math behind agentic CPU demand on the back of META Muse is overstated, even under generous assumptions. Importantly, we have heard from hyperscale experts that when deploying new services, cloud vendors typically significantly oversubscribe these services with a more normal ratio of hardware capacity vs. total expected users falling within 5-10%. Our analysis assumes 2 vCPUs per META Muse sandbox and a 126-core AMD EPYC CPU, with infrastructure requirements remaining relatively modest even under elevated engagement assumptions. In our base case (3MM DAU, 25% peak concurrency, 4x oversubscription), demand reaches ~1MM concurrent sandboxes, requiring 375K physical cores or ~3K CPU units. In a high-engagement scenario (3MM DAU, 50% peak concurrency, no oversubscription), demand rises to ~2MM concurrent sandboxes, requiring ~3MM cores and ~24K CPU units. More broadly, CPU requirements scale from 595-4,762 units at 10% concurrency, 1,488-11,905 units at 25% concurrency, and 2,976-23,810 units at 50% concurrency, depending on oversubscription assumptions. Even at a theoretical 100% concurrency rate, demand ranges from 5,952 to 47,619 CPU units. Given that the bull case effectively assumes every user is active simultaneously and that no oversubscription is available, it represents an unrealistic scenario.

Figure 2: META Muse Sensitivity Table

Scenario Table	Base Case	High Engagement	Notes
Daily Active Users (DAU)	3MM	3MM	UBSe
Peak Concurrent Sandboxes (% of DAU)	25%	50%	UBSe
vCPUs per Sandbox	2	2	META
Oversubscription	4	1	Conservative vs. Firecracker/DeepSeek
Cores per CPU	126	126	AMD EPYC 9D25 (126 cores/252 threads)

Math	Base Case	High Engagement	Math
Concurrent Sandboxes	1MM	2MM	DAU x Concurrency
vCPUs Demanded	2MM	3MM	Sandboxes x vCPUs per Sandbox
Cores Needed	375K	3MM	vCPUs / Oversubscription
CPU Units (EPYC chips)	2.98K	23.81K	Cores needed / Cores per CPU

		Oversubscription			
	# of CPU Units	1x	2x	4x	8x
Concurrency	10%	4,762	2,381	1,190	595
	25%	11,905	5,952	2,976	1,488
	50%	23,810	11,905	5,952	2,976
	100%	47,619	23,810	11,905	5,952

Source: UBS research

Importantly, expert conversations and recent white papers point to consumer agentic workloads having lower CPU-to-GPU ratio requirements vs. enterprise. Some of the more CPU-hungry workloads include enterprise agentic coding and large-dataset analyses, while consumer workloads typically require tool calls and less-intensive processing, which allows engineers to use multi-threading and/or efficiency-focused CPUs, such as INTC's E-core lineup that is optimized for highly parallelized and less single-core-intensive workload.

We sized Muse's CPU footprint bottom-up from the 2-vCPU-per-sandbox spec and the AMD EPYC 9D25 (Turin, 126 cores/252 threads). In our base case, we assume 20% of daily active users have a live sandbox at peak (30% in a high engagement scenario), provisioned at 70% target utilization and oversubscribed 3x per hardware thread. While conservative relative to AWS Firecracker's tested ceiling and DeepSeek's reported DSec production density (here and here), our math implies ~2,382 CPUs for 3MM DAUs in the base case and ~4,012 CPUs in a high engagement scenario. Deployed hardware equates to 10-15% of the capacity required if every user ran on dedicated hardware, consistent with experts' guidance of 5-10% for our base case with room for slight upside for heavily oversubscribed cloud workloads.

Figure 3: Assuming 5-10% Target Utilization

Scenario Table	Base Case	High Engagement	Notes
Daily Active Users (DAUs)	3,000,000	3,000,000	Unique users who open Muse on a given day
(x) vCPUs per Sandbox	2	2	(x) Virtual CPUs allocated to each user's virtual machine (VM)
Total Capacity Needed (Threads)	6,000,000	6,000,000	Hardware needed if every DAU ran at once on dedicated hardware at one thread per vCPU
Daily Active Users (DAUs)	3,000,000	3,000,000	Unique users who open Muse on a given day
(x) Peak Concurrent Sandboxes (% of DAUs)	20%	30%	(x) UBSe share of DAU with a live sandbox (a running VM, whether active or idle) during peak hours
Concurrent sandboxes	600,000	900,000	# of VMs alive at peak
(/) Target Utilization	70%	65%	(/) UBSe how fully the provisioned capacity runs at peak with the remainder a headroom for demand spikes etc.
Sandboxes Provisioned	857,143	1,384,615	Sandbox slots built, including headroom
(x) vCPUs per Sandbox	2	2	(x) Virtual CPUs allocated to each user's virtual machine (VM)
vCPUs Provisioned	1,714,286	2,769,231	vCPUs across all provisioned slots
(/) Oversubscription	3	3	(/) UBSe on how many vCPUs share one physical hardware thread (benchmarked versus DeepSeek and AWS)
Threads Provisioned	571,429	923,077	Physical hardware threads needed after sharing
(/) Total capacity needed (threads)	6,000,000	6,000,000	(/) Hardware needed if every DAU ran at once on dedicated hardware at one thread per vCPU
Hardware Deployed (% of total capacity)	10%	15%	UBSe that heavily oversubscribed cloud workloads deploy 5-10% of dedicated equivalent hardware in base case
Threads Provisioned	571,429	923,077	Physical hardware threads needed after sharing
(/) Usable Threads per CPU	240	230	(/) Hardware threads available for sandboxes after host OS and hypervisor take their share
# of CPUs being used	2,382	4,012	# of CPUs required

Source: UBS research

BASED ON DEMAND SIGNALS ONE COULD DREAM OF A $500B+ TAM, BUT...

Assuming unconstrained demand and using our framework for estimating the CPU TAM as a combination of traditional CPU units, AI head-nodes, and standalone AI CPUs, our demand analysis suggests CPUs could reach 150MM units by C2030, which we think could represent a $500B+ TAM, assuming ASPs for head-node and agentic AI CPUs increase from ~$2-2.5K/CPU for traditional CPU applications to $4K+ for AI CPUs.

Figure 4: Unconstrained Estimates Suggest a $500B+ CPU TAM by C2030

Unconstrained Demand Ests	2025	2027E	2030E	Comment
Accelerators (NVDA, AMD, TPU, Trn)	9	29	~44	Assumes units stay the same but TDP/performance increases
Head node CPU, MM	3	11	~22	UBSe based announced systems
Head Node CPU-to-GPU attach rate	0.29	0.36	0.50	
Head Node ASP	$2,430	$3,160	$4,480	Higher ASP given higher portion of NVDA CPUs
Standalone (outside of rack)	0	~20	~95	Additional CPU units outside of rack
Standalone CPU-to-GPU attach rate		0.69	2.16	3CPU:1GPU target ratio ex-head nodes and 50% of trad CPUs for tool calls
Standalone CPU ASP		$2,850	$4,223	Assumes higher share of x86 even over time for multithread
AI CPU Market (Bottom-up Ests), MM	3	30	117	
AI CPU ASP, K	$2,430	$2,957	$4,271	Higher core counts and clock speeds increasing ASPs
AI CPU Market (Bottom-up Ests), $B	7	90	500	
CPU-to-GPU Attach	0.29	1.05	2.66	
Implied Total CPUs (AI + Traditional), MM	23	56	147	
Implied Total CPU TAM, $B	31	125	548	
Implied DRAM Memory, EB				

Source: UBS estimates

But the key factor that should be considered is DRAM availability. As we previously cautioned, AI CPUs require multiple times more DRAM compared to traditional CPU use cases. For example, we estimate that the average traditional CPU requires ~0.5TB of memory today, likely scaling closer to ~1TB by the end of the decade. At the same time, AI head-node requirements already exceed 1.5TB and will likely trend toward ~2TB over time. AI standalone CPUs, which typically scale more parallel agents, may see memory requirements reach ~3TB over the next few generations

Figure 5: Memory Demand For Agentic Workloads is 3-4x Higher vs Traditional CPUs

DRAM / CPU	C2025/26	C2026/27	C2030E	Growth
Traditional CPU	~0.25-0.5TB	~0.5-0.75TB	~1.0TB	~1.5-2x
AI head-node CPU	~1-1.5TB	~1-1.5TB	~1.5-2TB	~1.5x
AI standalone / agentic CPU	~1.5-2TB	~1.5-2TB	~2-3TB	~1.5x
Blended AI CPU	**~1.5TB**	**~1.5TB**	**~2-2.5TB**	**~1.5x**

Source: UBS Research

With memory demand per CPU increasing ~4x between traditional and AI CPUs, we think that the current view on DRAM ex-HBM industry suggests the memory market increasing ~2x from C2027 to C2030 would not offer sufficient upside.

Figure 6: UBSe DRAM Capacity by Vendor

EB	C2025	C2027E	C2030E
Samsung	14	19	31
SK Hynix	10	16	22
Micron	9	13	26
CXMT	3	6	10
Total	**36**	**55**	**89**

Source: UBS Research

Figure 7: Our Current Base Case Memory Model Suggests 77EB for Server DRAM in C2030

ex-HBM DRAM Market (EB)	C2025	C2027E	C2030E
MU DRAM ex-HBM	9	13	23
Total Mkt (MU share, UBS Ests)	36	54	99
Total non-HBM DRAM Mkt	36	55	89
Non-Server DRAM	26	27	30
Server DRAM Memory	11	27	69
Old DDR Reuse via CXL (UBS Ests)		2	7
Total DRAM Memory for CPUs, EB	**11**	**29**	**77**

Source: UBS Research

As we argued before, we see DRAM as a key constraint on CPU deployment, greater even than the 2nm CPU wafer supply at TSMC. Our own pro forma estimate for the DRAM available for CPU deployment in C2030 suggests ~77EB of capacity added by key vendors, based on our MU estimates and assumed market share.

At recent industry events, we have seen data points from META suggesting that CXL enables the effective reuse of older generation RDIMMs. At the AI Infra Summit, we saw META present, in one of the workshops, a 1TB server solution using a 3:1 ratio of DDR5 RDIMMs and CXL-attached DDR4 while actively masking the CXL-related latency. Across four distinct META CPU workloads, META suggested it could use the proposed design in 25-33% of the respective fleet. This would suggest the solution could offer an incremental 6-8% of memory capacity to the market via reusing older RDIMMs.

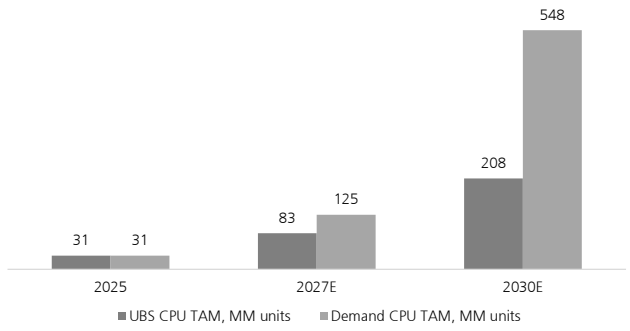
Figure 8: We See The 3:1 CPU:GPU Attach Ratio Resulting in DRAM Demand ~4x Above Our Current Estimates

	C2025		C2027E		C2030E - Base Case		C2030E - Unconstrained	
DRAM / CPU	CPU, MM	DRAM, EB	CPU, MM	DRAM, EB	CPU, MM	DRAM, EB	CPU, MM	DRAM, EB
Traditional CPU	20	10	26	12	30	15	30	15
AI head-node CPU	3	3	11	11	22	33	22	33
AI standalone / agentic CPU	0	0	5	7	15	29	117	234
Total CPU, MM	**23**		**41**		**67**		**169**	
Implied DRAM, EB		**14**		**29**		**77**		**282**

Source: UBS estimates

We point to a meaningful disconnect between the memory market and CPU expectations.

Figure 9: Bottom Up TAM Difference



Source: UBS estimates

UPDATED TAM FRAMEWORK POINTS TO A ~$210B TAM

We update our CPU market expectations taking into consideration both demand and supply, including wafer and memory availability. Our base case assumes 37MM AI CPUs by 2030 in addition to 30MM traditional CPU deployments, requiring ~77EB of DRAM.

Figure 10: Our Estimates Suggest a ~$210B TAM by C2030

Estimate	2025	2027E	2030E	Comment
Traditional CPU Market				
Traditional CPU Market, MM	20	26	30	Grows from 2025 at 8% CAGR like C2005-2020
Traditional CPU ASP, K	$1,200	$1,348	$1,606	Grows from 2025 at 1.5x 4% CAGR like C2005-2020
Traditional Market, $B	24	35	48	Implied by units and ASPs
AI CPU (Bottom-Up)				
AI Accelerator Units (Bottom-Up), MM	9	29	44	Largest programs (NVDA, AMD, TPU, Trn), ex-China
AI CPU Market (Bottom-Up Ests), MM	3	16	37	
Head Node	3	11	~22	UBSe based announced systems
Standalone	0	5	15	Additional CPU units outside of rack
AI CPU ASP, K	$2,430	$3,058	$4,378	Higher core count and clock speed increasing ASPs
AI CPU Market (Bottom-up Ests), $B	7	48	160	
AI CPU (Top-Down)				
AI CPU Market (Top-Down Ests), MM			80	TAM based on 3-4T AI Spend, and 1-to-1 CPU/GPU attach
AI CPU ASP, K			$3,000	UBS ASP ests
AI CPU Market, $B			240	Implied by units and ASPs
Total Market				
Bottom-up Units, MM	23	41	67	Traditional CPU + AI CPU Bottom-up Ests
Bottom-up TAM, $B	31	83	208	
Top-Down CPU Market, MM			110	Traditional CPU + AI CPU Top-Down Ests
Top-Down CPU Market, $B			288	

Source: UBS estimates

Interestingly, that implies an AI CPU attach ratio of 0.83:1 CPU-to-GPU in C2030, which is unchanged from our prior estimate.

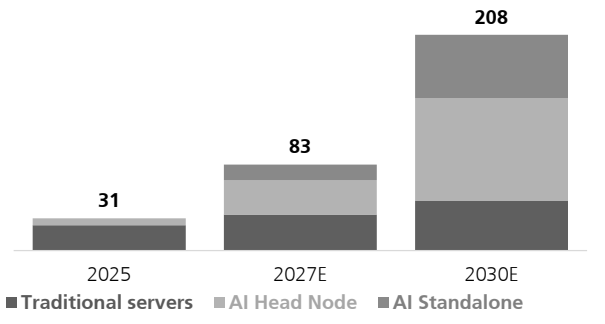
Figure 11: AI CPU Bottom-Up TAM Estimate

AI CPU Bottom-Up TAM Estimate	2025	2027E	2030E	Comment
Accelerators (NVDA, AMD, TPU, Trn)	9	29	~44	Assumes units stay the same but TDP/performance increases
Head Node CPUs, MM	3	11	~22	UBSe based announced systems
Head Node CPU per GPU attach rate	0.29	0.36	0.50	
Head Node ASP	$2,430	$3,160	$4,480	Higher ASP given higher portion of NVDA CPUs
Standalone (outside of rack)	0	5	~15	Additional CPU units outside of rack
Standalone CPU per GPU attach rate		0.18	0.33	
Standalone CPU ASP		$2,850	$4,223	Assumes higher share of x86 even over time for multithread
AI CPU Market (Bottom-up Ests), MM	3	16	37	
AI CPU ASP, K	$2,430	$3,058	$4,378	Higher core count and clock speed increasing ASPs
AI CPU Market (Bottom-up Ests), $B	7	48	160	
CPU per GPU attach	0.29	0.54	0.83	

Source: UBS estimates

We estimate ASPs of ~$4.4K (vs. prior $3.8K/CPU). Here we think this may prove to be a more conservative estimate. Should ASPs per core remain stable and core counts double, we could see $4-5K/CPU ASPs in the outer years (vs. $2-2.5K in C2025), as we expect the supply-constrained environment to contribute to ASP/core resiliency.

Figure 12: CPU Market To Reach ~$210B By C2030E



Source: UBS Research, Mercury Research

Across CPU vendors we are increasingly hearing ~$200B/yr market sizing. AMD estimates the TAM at $220B, while NVDA suggests a $200B TAM, without specifying a year.

Figure 13: AMD + ARM Quotes

	Quote	Timeline
AMD	"First of all, we think the overall server CPU TAM is going to grow, let's call it, strong double digits in 2026, just given the – as we said, the relationship between CPU demand and overall AI ramp."	2/3/2026
	"Based on the demand signals we are seeing today and the structural increase in CPU compute requirements driven by agentic AI, we now expect the server CPU TAM to grow at greater than 35% annually, reaching over $120 billion by 2030."	5/5/2026
	"What we're actually seeing is agentic AI is creating an entirely new growth vector for server CPUs."	7/23/2026
	"So yes, we've updated our TAM to over 50% over the next 3 or 4 years, reaching over $200 billion. We are still very much focused on being 50% of that market."	7/23/2026
	"The important point on Venice is we think our share growth, and that's not just because the market is larger. Of course the market is larger, but we think our share grows as we're seeing the breadth of workloads that people are wanting to put AMD on."	7/23/2026
	"So if you think about the various categories, if today in the head node type configuration, the CPU to GPU ratio may be four GPUs to one CPU, we certainly expect that to tighten as we go into future generations. And then when we add agents, certainly for the new TAM, we're expecting that the CPU ratio will actually be greater than one. So maybe we get to the point where two CPUs for one GPU."	7/23/2026
	"So based on what we're seeing today, and we're going to talk more about how agentic AI is evolving with CPUs, we now expect the server CPU market to grow by over 50% to over $200 billion by 2030. So that's starting from today, a $25 billion market going to over $200 billion."	7/23/2026
	".....and we expect the server CPU market to grow more than 50% annually to approximately $220 billion by 2030."	8/4/2026
	"In server CPUs with very strong customer demand and improved supply, we now expect server revenue to grow more than 80% year-over-year in the second half of 2026 and more than 70% for the full year 2027 off a much higher base."	8/4/2026

	Quote	Timeline
ARM	"Today we are just focused on the CPUs, and there is $455 billion of non-accelerator CPU TAM that is an Arm sweet spot. I want to direct particular attention to the $100 billion-plus in the data center CPU TAM, as that is the market we are addressing today with the Arm AGI CPU."	3/30/2026
	"We expect material revenue from the Arm AGI CPU starting in FYE28, with an exponential ramp to around $15 billion in FYE31. "	3/31/2026
	"So we think we're going to play alongside our partners in this space and we also think the opportunity is very, very large for both and I'm actually confident that by the end of the decade I believe the largest market share by CPU type will be ARM."	5/6/2026
	"When we talked about $100 billion TAM on March 24 at the ARM Everywhere day, I think we were the first company to talk about numbers in that magnitude and we had a bunch of questions about it. Now it's nice to see the rest of the market catching up and going higher than the number. Could the number be $120 billion out in that timeframe? Certainly."	5/6/2026
	"We are seeing literally not only an explosion of CPU demand, but one of the areas that we're seeing growth in terms of CPU is the number of cores per CPU. Many of these agents want to run independent jobs or flows or a batch on a specific CPU core."	5/6/2026
	"Which is why when people look at $100 billion TAM in five years, or $125 billion, whatever the number is, largely that's going to be driven by the fact that these CPU chips are going to have lots of CPU cores which will drive ASP up"	5/6/2026

Source: Company reports

Figure 14: INTC + NVDA Quotes

	Quote	Timeline
INTC	"Within Intel Products, we forecast a more pronounced revenue decline in CCG than in DCAI as we continue to prioritize internal supply to our server end markets."	1/22/2026
	"Demand for traditional servers continue to be very strong, and we are focused on ramping available capacity to support the meaningful uptick we are seeing, including partnering with key customers to support their needs beyond 2026."	1/22/2026
	"Within the server market, customer feedback and our own market intelligence point to the likelihood of a strong year of growth for DCAI. "	1/22/2026
	"While the market continues to benefit from more power efficient CPUs stimulating a refresh cycle, all indicators point to the growing and essential role CPUs will play within hyperscale and enterprise AI datacenters as inference driven AI usage expands. "	1/22/2026
	"We have also made decisive changes to simplify our server roadmap, focusing resources on the 16-channel Diamond Rapids and areas to accelerate the introduction of Coral Rapids where we can. "	1/22/2026
	"The client TAM is unfolding as expected and server CPU demand continues to far outpace available supply."	7/23/2026

	Quote	Timeline
NVDA	"Vera CPU opens a brand new $200 billion TAM for NVIDIA, a market we have never addressed before. And every major hyperscale and system maker is partnering with us to get it deployed. We have visibility to nearly $20 billion in total CPU revenue this year, setting us up to become the world-leading CPU supplier."	5/20/2026
	"Vera is opening a brand-new CPU market built for agents and every enterprise is becoming an agent company, and they run on NVIDIA."	6/24/2026
	"So, we built a brand-new CPU from the ground up for agents. This is a new market. CPUs for humans were sliced and rented per core. Agents don't rent cores, but demand ultra-fast responses. There will be billions of agents, and they will need a CPU built for them. We believe Vera will be one of the most significant product launches in our company's history"	6/24/2026
	"We continue to see demand for approximately $20 billion in total server CPUs. And based on our customer demand and improving supply outlook, our preliminary expectation is for CPU revenue to more than double in fiscal '28, positioning us as one of the world's leading server CPU suppliers."	8/26/2026

Source: Company reports

WHAT DOES THAT MEAN IN GW DEPLOYED

Moving from a traditional CPU to an AI CPU nearly doubles total CPU+DIMM power from ~370W to ~690W, reducing CPUs per GW from ~2.70MM to ~1.45MM. Despite the lower CPU count, higher CPU ASPs, and ~2x DRAM content, we estimate this will drive CPU+memory capex per GW to ~$34B versus ~$18B, with memory accounting for the vast majority of the increase.

Figure 15: We estimate ~2.7MM of traditional CPUs per GW, but only 1.4MM of AI CPUs

CPUs per GW	Traditional CPU	AI CPU
CPU TDP	250W	450W
Approx DRAM / CPU	1TB	2TB
DRAM/DIMM	15W	15W
# of DIMMs populated	8	16
Total TDP per CPU	370W	690W
# of CPUs per GW	2,702,703	1,449,275
CPU ASP	$1,500	$3,500
CPU+DIMMs Capex per GW, $B	$17.6	$34.1
CPU Capex	$4.1	$5.1
Memory Capex ($12K/1TB)	$13.5	$29.0

Source: UBS Research

CPU-related power demand could nearly triple from ~8GW in 2025 to ~21GW by 2030E, with traditional CPU demand growing relatively modestly to ~11GW. The bigger inflection comes from standalone AI CPUs, which we estimate will reach ~10GW by 2030E versus essentially zero today.

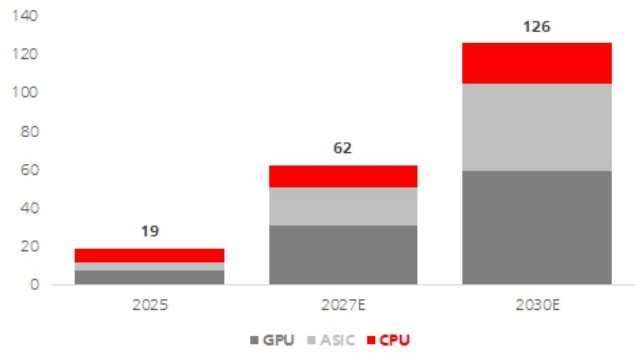
Figure 16: Our CPU TAM Estimates Suggest ~21GW of Power Required for New CPUs in C2030

CPU Power, GW	2025	2027E	2030E
Traditional CPUs	20	26	30
W per CPU + DIMMs	370	370	370
Traditional Implied GW	8	10	11
AI Head node Implied GW		captured in GW racks	
AI Standalone CPUs	0	5	15
W per CPU + DIMMs	690	690	690
AI Standalone Implied GW	0	4	10
Total CPUs, Implied GW	8	13	21

Source: UBS Research

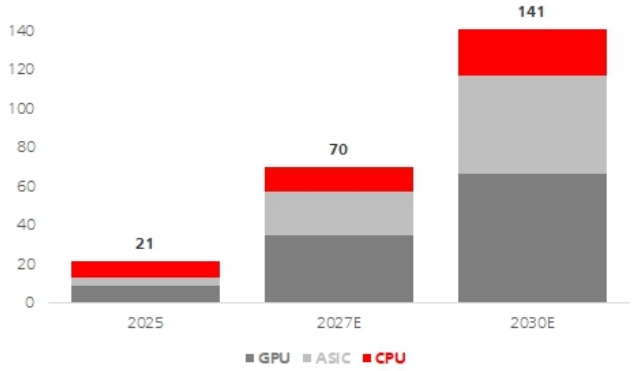
As discussed in our powered shell note (here), while accelerators will drive the bulk of incremental capacity, the rising power and memory content per CPU should support an expansion in CPU-driven infrastructure spending (~20% of 2030 annual compute capacity deployed). We see CPU deployments remaining a meaningful contributor to data center power growth through C2030, with CPU-related demand increasing materially as higher-power AI CPUs gain share alongside traditional server processors.

Figure 17: Gigawatts of Annual Compute Capacity Deployed



Source: UBS research

Figure 18: Gigawatts of Annual Data Center Capacity Deployed, Assuming PUE of 1.12



Source: UBS research

PROLIFERATION OF WORKLOADS DRIVES HETEROGENOUS HARDWARE

We are increasingly seeing a more aggressive push from outsiders to enter what was traditionally a near-monopoly market, with INTC accounting for 90%+ market share. We are also seeing strong momentum behind the ARM ecosystem as the software ecosystem has matured substantially - both ARM and QCOM are planning to introduce their own ARM-based data center CPUs over the coming years. Beyond that, we are hearing of a new crop of startups trying to leverage the RISC-V architecture and break into the market - this is where most experts we speak to remain somewhat more skeptical as RISC-V based solutions will likely face the same challenges ARM faced in 2018 when AWS introduced its Graviton Gen 1 chip - namely, a lack of software that can run on the new instruction set. Addressing this challenge may take less time now that AI coding is a helpful tool, but it will still take time.

At the same time, what we have been increasingly hearing over the past few months is that the shape of the workload dictates the hardware best suited to support it. In cases where agents are working on latency-sensitive applications, tool invocation, reasoning, or serial processing, strong single-threaded performance matters most - this is where NVDA is seeing traction for Grace and Vera, AMD for lower core-count, higher density core, and INTC for its P-core lineup.

At the same time, experts highlight that many workloads can be highly parallelized - handling many agents, tool calls, data retrievals, and background tasks. For these workloads, smaller cores and multi-threaded performance might make sense. This is where INTC's E-core line fits and we are now picking up on comments that make us incrementally more positive on INTC's E-core opportunity, although it still largely trails AMD on the performance roadmap.

However, E-core CPUs don't offer a memory advantage and a highly parallel approach often requires additional memory to support more threads working simultaneously. The solution in this case becomes heterogeneous cores deployed with flexible memory architectures and hierarchies. Here, mechanisms such as CXL offer advantages by providing additional capacity on top of local DRAM.

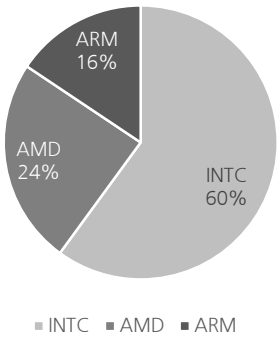
Figure 19: UBSe Server CPU Market Breakdown C2030E

(Revenues $B)	C2030E	%
AMD Server	$62	30%
INTC Server	$44	21%
ARM Server	$102	49%
of which ARM AGI	$15	7%
of which NVDA	$53	25%
of which Other	$35	17%
Total Server CPU	**$208**	**100%**

Source: UBS estimates

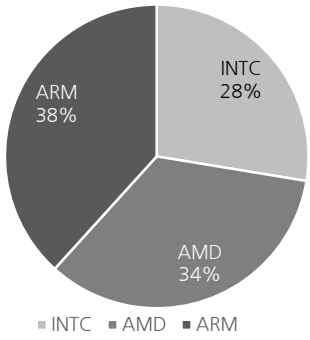
Our models and estimates are already fully pricing in our new TAM math.

Figure 20: C2025 Traditional CPU Mkt Share Based on Units



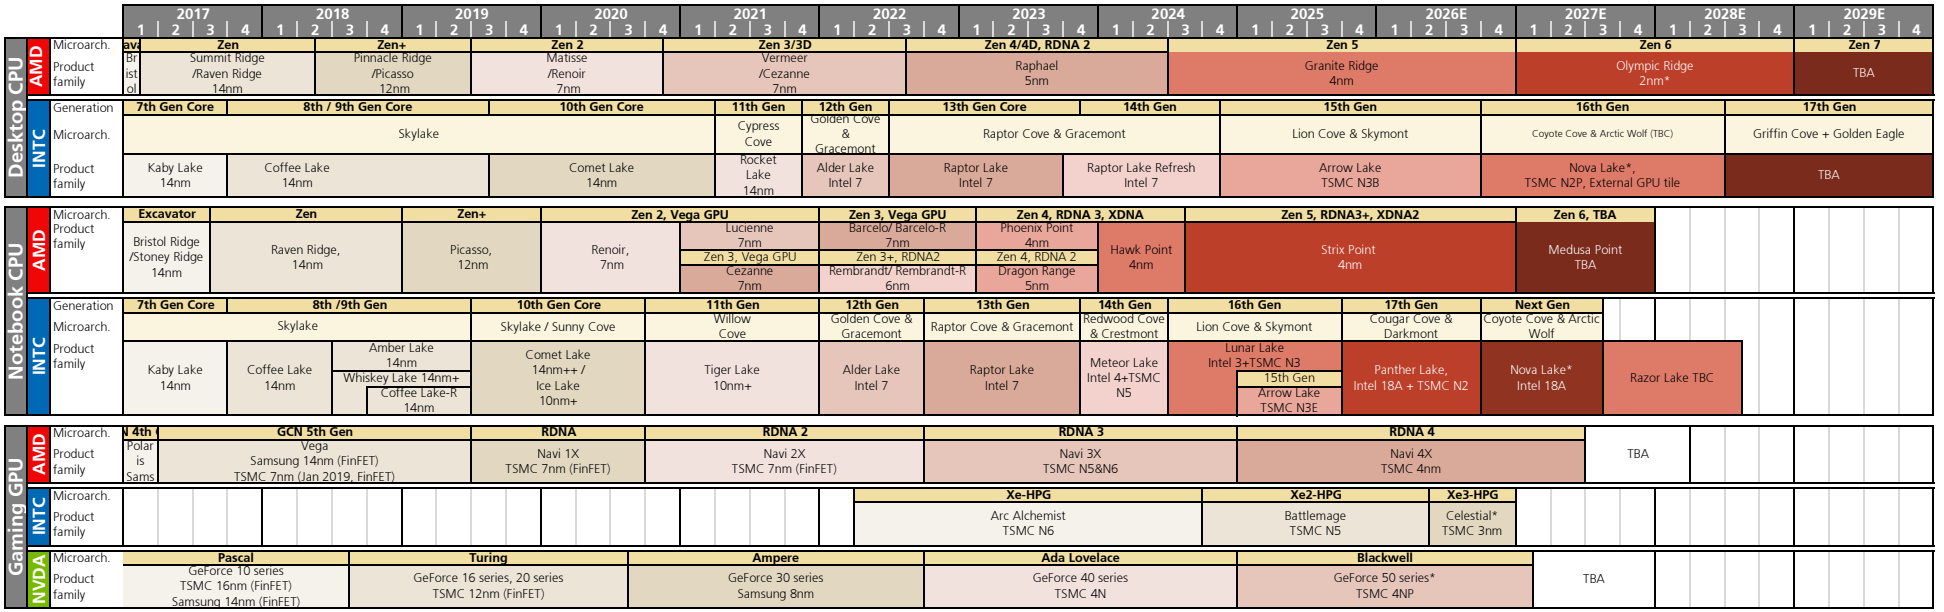
Source: UBS Research

Figure 21: UBSe C2030 Total CPU Mkt Share Based on Units



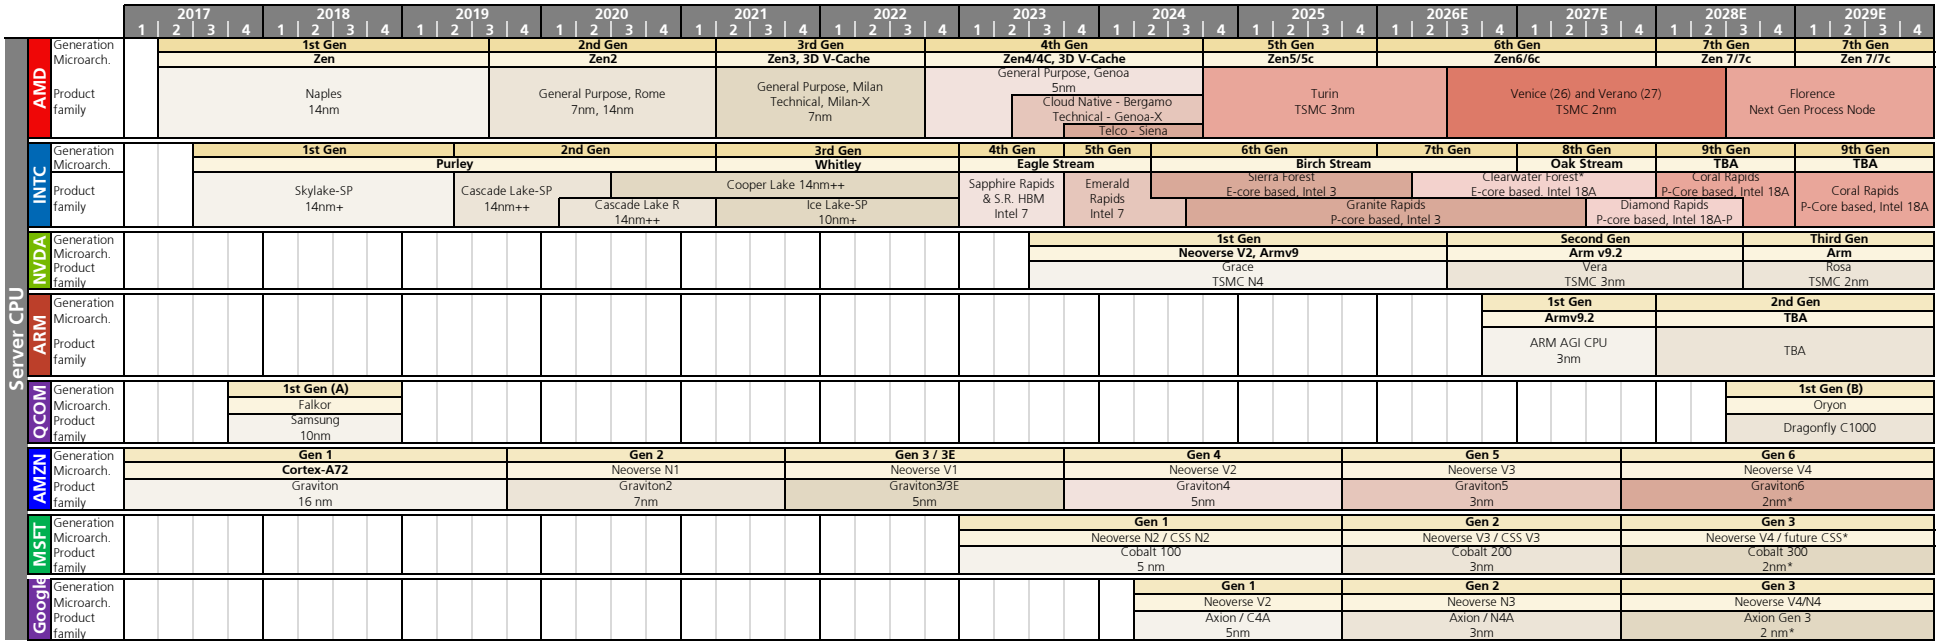
Source: UBS Research

Figure 22: Ex-Server CPU Roadmap



Source: Company reports, UBS Research

Figure 23: Server CPU Roadmap



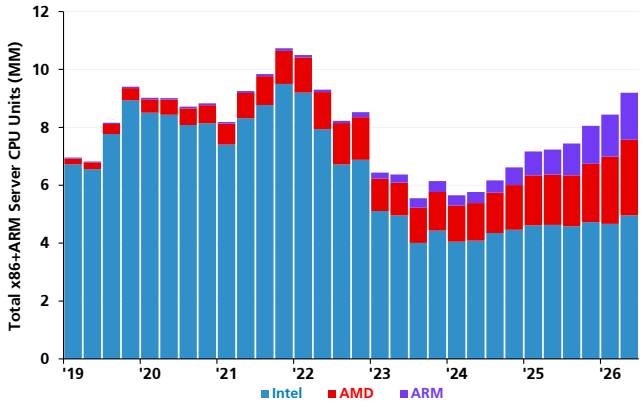
Source: Company reports, UBS research

Figure 24: CPU Specs

	AMD Venice	INTC Clearwater Forest	INTC Diamond Rapids	NVDA Vera	ARM AGI	AWS Graviton5	MSFT Cobalt 200	QCOM Dragonfly C1000	Google Axion N4A
Process	TSMC N2	Compute 18A Base INTC 3 I/O INTC 7	Cores 18A-P Base INTC 3-T Fabric hubs INTC 3	-	TSMC N3P	TSMC 3nm	TSMC 3nm	-	TSMC 3nm
Package	8 CCDs + 2 I/O dies	12 compute + 3 bases +2 I/O tiles	16 core chiplets on 4 base tiles + 2 fabric hubs	Monolithic compute die + separate memory controller and I/O dies	2 chiplets	4 chiplets	2 chiplets	Multi-chiplet	-
Base / nominal clock	2.55 GHz (9996)	2.2 GHz (6990E+, 450W)	-	-	3.2 GHz (128/136C) 3.5 GHz (64C) 2.8 GHz typical	-	-	-	-
Max boost	5.0 GHz (F-SKUs) 5.15 GHz (Venice-X)	3.2 GHz (2.8 all-core)	-	-	3.5 GHz (128/136C) 3.7 GHz (64C)	-	-	-	-
TDP	Up to 600W	300-450W	Up to 650W	250-450W	300W	-	-	-	-
Core design	Zen 6 / Zen 6c	Darkmont E-core	Panther Cove	Custom Olympus	Neoverse V3	Neoverse V3	Neoverse V3	Custom Oryon	Neoverse V3
Max cores	256 Zen 6c 128 Gen 6	288	256	88	136	192	132	250+	-
Max threads	512	288	-	176	-	-	-	-	-
Multithreading	SMT 2 threads/core	None	-	Spatial multithreading 2 threads/core	-	-	-	-	-
Core-count options	SP7: 64-256 SP8: 8-128	144 / 192 / 264 / 288	-	Single 88-core SKU	64 / 128 / 136	-	-	-	-
Compute dies x core	8 x 32 Zen 6c	12 x 24	16 x up to 16	1 x 88	2 x 70 physical	4 chiplets	2 x 66 active	Multi-chiplet	-
Total LLC	1,024 MB (9996, 4 MB/core) 1,152 MB Venice-X	576 MB	Up to 1.28 GB	164 MB	128 MB system-level cache	192 MB	192 MB	-	-
Memory type	DDR5RDIMM / MRDIMM	DDR5RDIMM	DDR5RDIMM / MRDIMM	LPDDR5X (SOCAMM2)	DDR5RDIMM	12-ch DDR5-8800	DDR5	-	DDR5
Peak memory bandwidth	1.6 TB/s	768 GB/s	Up to 1.6 TB/s	1.2 TB/s	~845 GB at DDR5-8800 ~614 GB/s at 6400	845 GB/s	-	-	-
PCIe	128 PCIe 6.0 lanes	96 PCIe 5.0 lanes	128 flexible lanes (PCIe 6.0 / CXL 3.0 / UPI 3) + 8 PCIe 4.0	88 PCIe Gen 6 lanes	96 PCIe 6.0 lanes (768 GB/s each way) + 6 PCIe 4.0 lanes	PCIe Gen 6	-	PCIe Gen 7 (>2 TB/s incl. CXL)	-
CXL	CXL 3.1	64 CXL 2.0 lanes	CXL 3.0 (on flexible lanes)	CXL 3.1	CXL 3.0 Type 3	-	-	-	-
Other links	-	6 UPI 2.0 links	UPI 3 (on flexible lanes)	NVLink-C2C (1,800 GB/s)	2 socket SMP AMBA CHI extension links	-	-	-	-

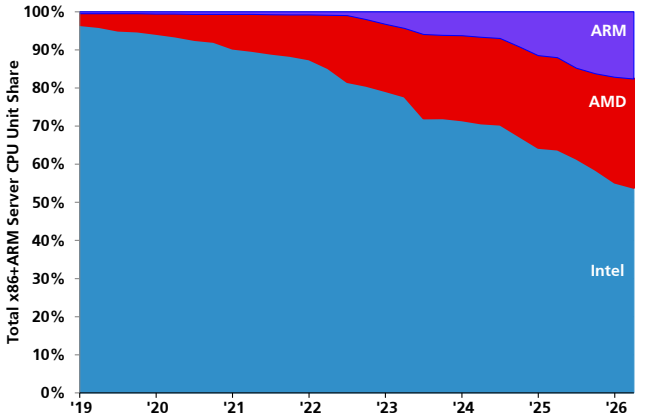
Source: Company reports, UBS research

Figure 25: Intel, AMD and ARM Server CPU Units



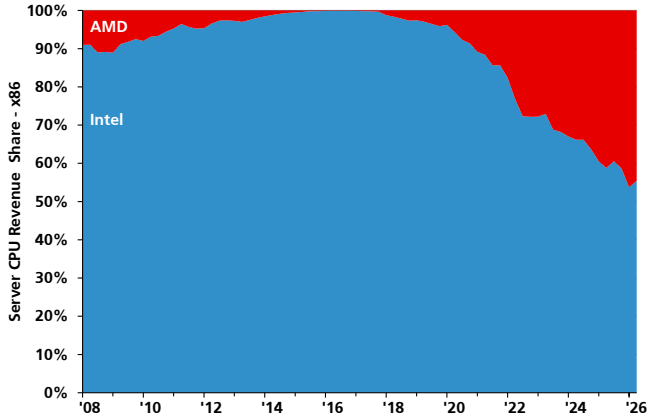
Source: Mercury Research, UBSe

Figure 26: Intel, AMD and ARM Server CPU Market Share (Units)



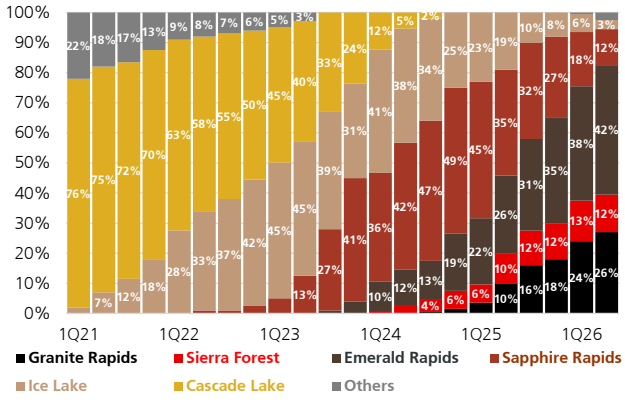
Source: Mercury Research, UBSe

Figure 27: Intel vs. AMD x86 Server CPU Market Share (Revenue)



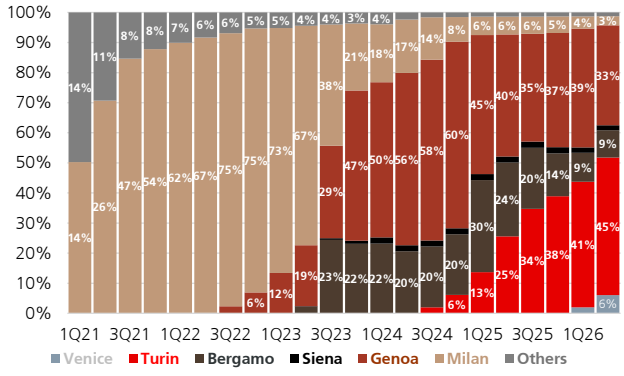
Source: Mercury Research, UBSe

Figure 28: INTC Server Units Decomposition



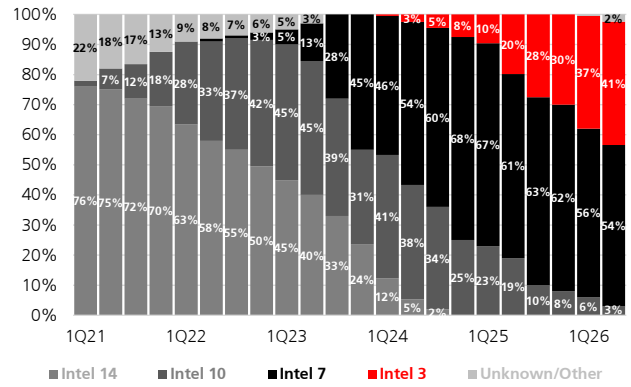
Source: Mercury Research, UBSe

Figure 29: AMD Server Units Decomposition



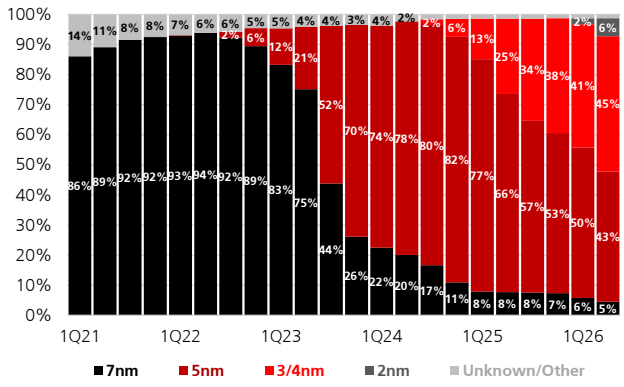
Source: Mercury Research, UBSe

Figure 30: INTC Server Units by Node



Source: Mercury Research, UBSe

Figure 31: AMD Server Units by Node



Source: Mercury Research, UBSe

Figure 32: AMD - Summary Model

Advanced Micro Devices (AMD) - Key Metrics (in $MM except EPS and Units)	C2023A	C2024A	C2025A	CQ1A	CQ2A	Estimates CQ3E	CQ4E	C2026E	CQ1E	CQ2E	CQ3E	CQ4E	C2027E	CQ1E	CQ2E	CQ3E	CQ4E	C2028E	C2029E	C2030E
Revenue	**22,680**	**25,785**	**34,639**	**10,253**	**11,536**	**13,180**	**17,122**	**52,092**	**19,063**	**20,930**	**25,508**	**32,421**	**97,922**	**32,538**	**33,841**	**35,995**	**37,618**	**139,993**	**176,924**	**221,588**
Y/Y Growth	(4%)	14%	34%	38%	50%	43%	67%	50%	86%	81%	94%	89%	88%	71%	62%	41%	16%	43%	26%	25%
Gross Profit (Non-GAAP)	11,436	13,759	18,165	5,685	6,488	7,386	9,476	29,035	10,492	11,499	14,066	17,958	54,016	18,105	18,897	20,208	21,213	78,424	101,143	129,261
Gross Margin	**50.4%**	**53.4%**	**52.4%**	**55.4%**	**56.2%**	**56.0%**	**55.3%**	**55.7%**	**55.0%**	**54.9%**	**55.1%**	**55.4%**	**55.2%**	**55.6%**	**55.8%**	**56.1%**	**56.4%**	**56.0%**	**57.2%**	**58.3%**
Ex-Semi-Custom GM (UBSe)	57%	55%	54%	56%	57%	57%	56%	56%	55%	55%	55%	56%	55%	56%	56%	57%	57%	56%	58%	59%
R&D (GAAP)	5,872	6,456	8,091	2,397	2,528	2,638	2,818	10,381	3,118	3,368	3,608	3,848	13,942	3,968	4,038	4,108	4,178	16,292	17,712	19,512
SG&A (GAAP)	2,352	2,783	4,144	1,253	1,401	1,491	1,591	5,736	1,661	1,731	1,781	1,831	7,004	1,806	1,856	1,931	2,001	7,594	8,184	8,754
Other & Non-GAAP Adjustment (GAAP)	1,835	1,586	1,223	290	284	284	284	1,142	284	284	284	284	1,136	284	284	284	284	1,136	1,136	1,136
Total Opex (Non-GAAP)	**6,616**	**7,669**	**10,397**	**3,145**	**3,394**	**3,634**	**3,914**	**14,087**	**4,284**	**4,604**	**4,894**	**5,184**	**18,966**	**5,279**	**5,399**	**5,544**	**5,684**	**21,906**	**23,916**	**26,286**
Y/Y Growth	10%	16%	36%	42%	40%	32%	30%	35%	36%	36%	35%	32%	35%	23%	17%	13%	10%	16%	9%	10%
Operating Income (Non-GAAP)	4,854	6,138	7,768	2,540	3,094	3,752	5,562	14,948	6,208	6,895	9,172	12,774	35,050	12,826	13,498	14,664	15,529	56,518	77,227	102,975
Operating Margin	**21.4%**	**23.8%**	**22.4%**	**24.8%**	**26.8%**	**28.5%**	**32.5%**	**28.7%**	**32.6%**	**32.9%**	**36.0%**	**39.4%**	**35.8%**	**39.4%**	**39.9%**	**40.7%**	**41.3%**	**40.4%**	**43.6%**	**46.5%**
Non-GAAP EPS	**$ 2.65**	**$ 3.31**	**$ 4.17**	**$ 1.37**	**$ 1.66**	**$ 2.00**	**$ 2.87**	**$ 7.91**	**$ 3.17**	**$ 3.50**	**$ 4.59**	**$ 6.23**	**$ 17.58**	**$ 6.19**	**$ 6.46**	**$ 6.97**	**$ 7.30**	**$ 26.94**	**$ 35.57**	**$ 45.18**
Free Cash Flow																				
CFFO	1,667	3,041	7,709	2,955	2,366	4,145	4,329	13,796	5,815	6,149	7,185	9,622	28,771	12,585	12,292	13,057	13,979	51,913	68,439	91,002
Capex	546	636	974	389	808	958	1,158	3,313	900	900	900	900	3,600	900	900	900	900	3,600	3,600	3,600
FCF	**1,121**	**2,405**	**6,735**	**2,566**	**1,558**	**3,187**	**3,171**	**10,483**	**4,915**	**5,249**	**6,285**	**8,722**	**25,171**	**11,685**	**11,392**	**12,157**	**13,079**	**48,313**	**64,839**	**87,402**
% revenue	5%	9%	19%	25%	14%	24%	19%	20%	26%	25%	25%	27%	26%	36%	34%	34%	35%	35%	37%	39%
Segment Revenues																				
Data Center	**6,496**	**12,579**	**16,635**	**5,775**	**6,718**	**8,394**	**12,467**	**33,355**	**14,419**	**16,400**	**20,747**	**27,432**	**78,999**	**27,556**	**28,897**	**30,359**	**31,716**	**118,528**	**153,270**	**196,924**
Data Center GPU Revenue	**555**	**5,025**	**6,754**	**2,430**	**2,750**	**3,550**	**6,400**	**15,130**	**7,600**	**8,500**	**11,250**	**15,900**	**43,250**	**15,900**	**17,000**	**18,100**	**19,200**	**70,200**	**97,800**	**131,400**
Server CPU Revenue	**5,433**	**6,794**	**8,827**	**3,065**	**3,684**	**4,539**	**5,705**	**16,993**	**6,457**	**7,537**	**9,133**	**11,167**	**34,295**	**11,261**	**11,468**	**11,793**	**12,010**	**46,531**	**52,961**	**61,987**
Server ASP ($)	$ 1,140	$ 1,351	$ 1,347	$ 1,535	$ 1,591	$ 1,647	$ 1,725	$ 1,638	$ 1,775	$ 1,850	$ 1,975	$ 2,100	$ 1,943	$ 2,107	$ 2,114	$ 2,121	$ 2,128	$ 2,118	$ 2,254	$ 2,420
Server Units (K)	4,768	5,028	6,554	1,997	2,316	2,756	3,307	10,375	3,638	4,074	4,624	5,318	17,654	5,344	5,425	5,560	5,644	21,973	23,493	25,616
Server TAM (K units)	23,307	22,432	26,046	6,506	6,640	6,727	6,939	26,812	6,683	6,823	6,913	7,126	27,545	--	--	--	--	--	--	--
AMD Share	20%	22%																		
XLNX FPGAs	**252**	**318**	**385**	**89**	**93**	**98**	**102**	**382**	**97**	**92**	**88**	**83**	**361**	**86**	**88**	**91**	**94**	**359**	**405**	**455**
Pensando	**254**	**441**	**669**	**191**	**191**	**208**	**260**	**850**	**265**	**271**	**276**	**282**	**1,094**	**310**	**341**	**375**	**412**	**1,437**	**2,104**	**3,081**
Data Center Operating Margin	20%	28%	22%	28%	31%	31%	35%	32%	36%	36%	36%	36%	36%	37%	37%	37%	37%	37%	38%	38%
Data Center Operating Profit Y/Y Drop-through				32%	65%	38%	37%		42%	39%	39%	37%		38%	38%	39%	43%			
Client	**4,651**	**7,054**	**10,640**	**2,885**	**3,062**	**2,997**	**2,957**	**11,901**	**2,900**	**2,847**	**2,987**	**3,147**	**11,881**	**2,943**	**2,969**	**3,466**	**3,605**	**12,983**	**14,080**	**14,290**
PC Revenue	**4,386**	**6,718**	**10,154**	**2,751**	**2,914**	**2,952**	**2,938**	**11,555**	**2,876**	**2,822**	**2,947**	**3,130**	**11,775**	**2,921**	**2,947**	**3,431**	**3,589**	**12,888**	**13,995**	**14,205**
Desktop Revenue	1,386	1,863	3,196	798	829	862	827	3,317	820	807	851	909	3,387	889	884	1,038	1,101	3,912	4,389	4,617
AMD DT ASP ($)	$ 174	$ 188	$ 266	$ 270	$ 275	$ 275	$ 272	$ 273	$ 278	$ 282	$ 286	$ 288	$ 284	$ 286	$ 286	$ 311	$ 311	$ 299	$ 309	$ 307
AMD DT Units (K)	7,950	9,931	12,029	2,956	3,015	3,135	3,041	12,147	2,950	2,862	2,976	3,155	11,942	3,107	3,092	3,339	3,539	13,078	14,202	15,038
DT TAM (K units)	67,010	69,021	75,737	17,683	17,652	19,987	21,236	76,558	17,685	17,667	20,019	21,277	76,647	18,215	18,197	20,619	21,915	78,947	81,315	83,754
AMD Share	12%	14%	16%	17%	17%	16%	14%	16%	17%	16%	15%	15%	16%	17%	17%	16%	16%	17%	17%	18%
Mobile Revenue	3,001	4,854	6,958	1,953	2,085	2,090	2,111	8,239	2,056	2,015	2,096	2,221	8,388	2,032	2,062	2,393	2,488	8,975	9,606	9,588
AMD Mobile ASP ($)	$ 98	$ 115	$ 157	$ 173	$ 173	$ 170	$ 170	$ 172	$ 180	$ 180	$ 180	$ 180	$ 180	$ 178	$ 178	$ 193	$ 193	$ 186	$ 193	$ 193
AMD Mobile Units (K)	30,609	42,212	44,314	11,264	12,052	12,293	12,416	48,026	11,423	11,195	11,642	12,341	46,601	11,415	11,587	12,398	12,893	48,293	49,774	49,681
Mobile TAM (K units)	160,289	186,556	180,960	41,006	42,182	43,053	43,653	169,894	41,751	42,715	43,538	43,951	171,955	43,004	43,996	44,844	45,270	177,114	182,427	187,900
AMD Share	19%	23%	24%	27%	29%	29%	28%	28%	27%	26%	27%	28%	27%	27%	26%	28%	28%	27%	27%	26%
AMD Pro Graphics Revenue	267	335	486	134	148	45	19	346	24	25	40	17	106	23	22	35	15	95	85	85
Client Operating Margin	-1%	13%	23%	16%	15%	19%	20%	18%	18%	19%	20%	21%	20%	19%	19%	20%	20%	20%	20%	20%
Client Operating Profit Y/Y Drop-through				3%	-30%	-57%	5%		351%	-35%	-279%	37%		86%	19%	20%	13%			
Gaming	**6,212**	**2,595**	**3,910**	**720**	**779**	**640**	**514**	**2,653**	**583**	**511**	**568**	**623**	**2,285**	**795**	**720**	**839**	**927**	**3,281**	**3,725**	**3,795**
PC GPU Gaming Revenue	**858**	**707**	**1,618**	**440**	**388**	**369**	**313**	**1,510**	**282**	**310**	**217**	**272**	**1,081**	**244**	**269**	**188**	**226**	**927**	**771**	**641**
Semi-Custom Revenue	**5,354**	**1,888**	**2,292**	**280**	**391**	**271**	**201**	**1,143**	**301**	**201**	**351**	**351**	**1,204**	**551**	**451**	**651**	**701**	**2,354**	**2,954**	**3,154**
Gaming Operating Margin	16%	11%	12%	11%	15%	15%	17%	14%	16%	15%	14%	15%	15%	16%	14%	13%	14%	14%	16%	16%
Embedded	**5,321**	**3,557**	**3,454**	**873**	**977**	**1,149**	**1,184**	**4,183**	**1,160**	**1,172**	**1,207**	**1,219**	**4,757**	**1,243**	**1,256**	**1,331**	**1,371**	**5,200**	**5,849**	**6,579**
Embedded Operating Margin	49%	40%	36%	39%	40%	42%	44%	41%	35%	32%	85%	176%	83%	156%	171%	198%	215%	186%	267%	375%
Q/Q Revenue Growth																				
Data Center				7%	16%	25%	49%		16%	14%	27%	32%		0%	5%	5%	4%			
Client				(7%)	6%	(2%)	(1%)		(2%)	(2%)	5%	5%		(6%)	1%	17%	4%			
Gaming				(15%)	8%	(18%)	(20%)		13%	(12%)	11%	10%		28%	(9%)	17%	10%			
Embedded				(8%)	12%	18%	3%		(2%)	1%	3%	1%		2%	1%	6%	3%			
Y/Y Revenue Growth																				
Data Center		94%	32%	57%	107%	93%	132%	101%	150%	144%	147%	120%	137%	91%	76%	46%	16%	50%	29%	28%
Client		52%	51%	26%	23%	9%	(5%)	12%	1%	(7%)	(0%)	6%	(0%)	1%	4%	16%	15%	9%	8%	1%
Gaming		(58%)	51%	11%	(31%)	(51%)	(39%)	(32%)	(19%)	(34%)	(11%)	21%	(14%)	36%	41%	48%	49%	44%	14%	2%
Embedded		(33%)	(3%)	6%	19%	34%	25%	21%	33%	20%	5%	3%	14%	7%	7%	10%	12%	9%	12%	12%

Source: Company data, FactSet, UBS estimates

Figure 33: INTC - Summary Model

Intel (INTC) - Key Metrics (in $MM except EPS and units)	C2024A	C2025A	CQ1A	CQ2A	Estimates CQ3E	CQ4E	C2026E	CQ1E	CQ2E	CQ3E	CQ4E	C2027E	CQ1E	CQ2E	CQ3E	CQ4E	C2028E	C2029E	C2030E
Revenue (Non GAAP)	**$53,101**	**$52,853**	**$13,577**	**$16,128**	**$16,974**	**$17,769**	**$64,448**	**$17,886**	**$18,134**	**$19,202**	**$19,970**	**$75,192**	**$19,978**	**$21,156**	**$22,673**	**$23,482**	**$87,288**	**$107,256**	**$132,058**
Y/Y Growth	(2%)	(0%)	7%	25%	24%	30%	22%	32%	12%	13%	12%	17%	12%	17%	18%	18%	16%	23%	23%
Gross Profit (Non -GAAP)	19,099	19,410	5,570	6,737	7,180	7,552	27,039	7,566	7,798	8,449	8,897	32,709	8,780	9,361	10,101	10,543	38,786	48,792	62,686
Gross Margin	**36.0%**	**36.7%**	**41.0%**	**41.8%**	**42.3%**	**42.5%**	**42.0%**	**42.3%**	**43.0%**	**44.0%**	**44.6%**	**43.5%**	**44.0%**	**44.3%**	**44.6%**	**44.9%**	**44.4%**	**45.5%**	**47.5%**
R&D (Non-GAAP)	16,546	13,774	3,375	3,368	3,775	3,860	14,378	3,840	3,820	3,870	3,930	15,460	4,050	4,170	4,290	4,410	16,920	20,140	23,050
MG&A (Non-GAAP)	5,507	4,624	1,038	1,175	1,250	1,280	4,743	1,200	1,250	1,320	1,380	5,150	1,310	1,430	1,530	1,560	5,830	6,240	6,890
Other	(2,700)	(1,904)	(511)	(576)	--	--	(1,087)	--	--	--	--	--	--	--	--	--	--	--	--
Total Opex (Non-GAAP)	**$19,353**	**$16,494**	**$3,902**	**$3,967**	**$4,475**	**$4,564**	**$16,908**	**$4,503**	**$4,526**	**$4,614**	**$4,711**	**$18,354**	**$4,761**	**$4,965**	**$5,140**	**$5,266**	**$20,131**	**$23,162**	**$25,978**
Y/Y Growth	*2%*	*(15%)*	*(9%)*	*(8%)*	*14%*	*15%*	*3%*	*15%*	*14%*	*3%*	*3%*	*9%*	*6%*	*10%*	*11%*	*12%*	*10%*	*15%*	*12%*
Operating Income (Non-GAAP)	(254)	2,916	1,668	2,770	2,705	2,988	10,131	3,062	3,272	3,835	4,186	14,355	4,020	4,396	4,961	5,278	18,655	25,630	36,708
Operating Margin	**(0.5%)**	**5.5%**	**12.3%**	**17.2%**	**15.9%**	**16.8%**	**15.7%**	**17.1%**	**18.0%**	**20.0%**	**21.0%**	**19.1%**	**20.1%**	**20.8%**	**21.9%**	**22.5%**	**21.4%**	**23.9%**	**27.8%**
Non-GAAP EPS	**($0.13)**	**$0.42**	**$0.29**	**$0.43**	**$0.39**	**$0.45**	**$1.57**	**$0.46**	**$0.50**	**$0.59**	**$0.64**	**$2.20**	**$0.64**	**$0.69**	**$0.78**	**$0.83**	**$2.94**	**$3.96**	**$5.42**
Free Cash Flow																			
CFFO	8,288	9,697	1,096	7,006	6,561	4,747	19,410	3,903	6,388	7,502	7,448	25,241	6,945	8,754	9,748	9,549	34,996	47,299	58,197
Less: Gross Capex	25,122	17,672	5,136	2,479	6,000	6,600	20,215	7,500	7,000	7,000	7,000	28,500	10,000	10,000	10,000	10,000	40,000	40,000	30,000
Add: Contribution from Partners	14,650	6,685	1,607	2,642	1,900	2,000	8,149	1,200	700	200	200	2,300	200	200	200	200	800	800	800
Net Capex	10,472	10,987	3,529	(163)	4,100	4,600	12,066	6,300	6,300	6,800	6,800	26,200	9,800	9,800	9,800	9,800	39,200	39,200	29,200
FCF	**($2,228)**	**($1,612)**	**($2,433)**	**$7,169**	**$2,461**	**$147**	**$7,344**	**($2,397)**	**$88**	**$702**	**$648**	**($959)**	**($2,855)**	**($1,046)**	**($52)**	**($251)**	**($4,204)**	**$8,099**	**$28,997**
% of revenue	*(4%)*	*(3%)*	*(18%)*	*44%*	*15%*	*1%*	*11%*	*(13%)*	*0%*	*4%*	*3%*	*(1%)*	*(14%)*	*(5%)*	*(0%)*	*(1%)*	*(5%)*	*8%*	*22%*
Y/Y Growth			*(34%)*	*(783%)*	*175%*	*(93%)*		*(1%)*	*(99%)*	*(71%)*	*342%*		*19%*	*(1,289%)*	*(107%)*	*(139%)*			
Product Revenue																			
Intel Products	**$49,471**	**$49,147**	**$12,779**	**$15,139**	**$15,932**	**$16,654**	**$60,504**	**$16,634**	**$16,916**	**$17,884**	**$18,609**	**$70,042**	**$18,219**	**$19,177**	**$20,374**	**$20,739**	**$78,508**	**$86,779**	**$93,871**
Q/Q			(1%)	18%	5%	5%		(0%)	2%	6%	4%		(2%)	5%	6%	2%			
Y/Y	4%	(1%)	9%	28%	26%	29%	23%	30%	12%	12%	12%	16%	10%	13%	14%	11%	12%	11%	8%
Operating Income	**$13,008**	**$12,739**	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
Operating Margin	26%	26%	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
CCPG	**$33,346**	**$32,228**	**$7,727**	**$8,877**	**$8,895**	**$8,756**	**$34,255**	**$8,098**	**$7,801**	**$8,148**	**$8,143**	**$32,190**	**$7,624**	**$8,014**	**$8,384**	**$8,158**	**$32,181**	**$33,196**	**$34,783**
Q/Q			(6%)	15%	0%	(2%)		(8%)	(4%)	4%	(0%)		(6%)	5%	5%	(3%)			
Y/Y		(3%)	1%	13%	4%	7%	6%	5%	(12%)	(8%)	(7%)	(6%)	(6%)	3%	3%	0%	(0%)	3%	5%
Operating Income	**$11,594**	**$9,317**	**$2,516**	**$2,343**	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
Operating Margin	35%	29%	33%	26%	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
Desktop Revenue	**$9,650**	**$8,966**	**$2,141**	**$2,310**	**$2,159**	**$2,029**	**$8,639**	**$1,844**	**$1,741**	**$1,758**	**$1,671**	**$7,014**	**$1,503**	**$1,503**	**$1,549**	**$1,626**	**$6,182**	**$6,078**	**$5,977**
Q/Q			(7%)	8%	(7%)	(6%)		(9%)	(6%)	1%	(5%)		(10%)	--	3%	5%			
Y/Y	(5%)	(7%)	1%	7%	(10%)	(12%)	(4%)	(14%)	(25%)	(19%)	(18%)	(19%)	(18%)	(14%)	(12%)	(3%)	(12%)	(2%)	(2%)
Desktop Units ('000s)	58,316	52,506	10,886	11,267	10,478	9,849	42,479	8,864	8,244	8,244	7,832	33,184	7,048	7,048	7,260	7,623	28,980	27,785	26,257
Desktop ASP	$165	$171	$197	$205	$206	$206	$203	$208	$211	$213	$213	$211	$213	$213	$213	$213	$213	$219	$228
Notebook Revenue	**$19,091**	**$18,657**	**$4,449**	**$5,381**	**$5,537**	**$5,648**	**$21,015**	**$5,196**	**$5,013**	**$5,290**	**$5,340**	**$20,840**	**$5,011**	**$5,412**	**$5,683**	**$5,344**	**$21,450**	**$22,345**	**$23,798**
Q/Q			(8%)	21%	3%	2%		(8%)	(4%)	6%	1%		(6%)	8%	5%	(6%)			
Y/Y	12%	(2%)	1%	21%	12%	17%	13%	17%	(7%)	(4%)	(5%)	(1%)	(4%)	8%	7%	0%	3%	4%	6%
Notebook Units ('000s)	150,490	148,178	30,748	33,207	32,543	32,543	129,041	29,940	28,742	30,179	31,085	119,946	28,598	30,886	32,430	30,808	122,722	125,598	128,541
Notebook ASP	$127	$126	$145	$162	$170	$174	$163	$174	$174	$175	$172	$174	$175	$175	$175	$173	$175	$178	$185
Other	**$1,549**	**$1,365**	**$353**	**$355**	**$357**	**$359**	**$1,424**	**$364**	**$369**	**$374**	**$379**	**$1,485**	**$384**	**$390**	**$396**	**$402**	**$1,572**	**$1,666**	**$1,766**
NEX in CCG	**$3,056**	**$3,239**	**$784**	**$831**	**$848**	**$848**	**$3,310**	**$848**	**$848**	**$890**	**$917**	**$3,503**	**$899**	**$890**	**$934**	**$962**	**$3,685**	**$3,866**	**$4,057**
DCAI	**$16,125**	**$16,919**	**$5,052**	**$6,262**	**$7,037**	**$7,898**	**$26,357**	**$8,536**	**$9,115**	**$9,736**	**$10,466**	**$37,853**	**$10,595**	**$11,163**	**$11,989**	**$12,580**	**$46,328**	**$53,583**	**$59,088**
Q/Q			7%	24%	12%	12%		8%	7%	7%	7%		1%	5%	7%	5%			
Y/Y		5%	22%	59%	71%	67%	56%	69%	46%	38%	33%	44%	24%	22%	23%	20%	22%	16%	10%
Operating Income	**$1,414**	**$3,422**	**$1,542**	**$2,956**	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
Operating Margin	9%	20%	31%	47%	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
Edge (ex-NEX in DCAI, incl Custom until D	**$3,308**	**$3,250**	**$766**	**$687**	**$694**	**$701**	**$2,849**	**$687**	**$701**	**$743**	**$773**	**$2,904**	**$757**	**$750**	**$787**	**$811**	**$3,105**	**$3,258**	**$3,418**
Custom ASIC	**--**	**--**	**$370**	**$500**	**$650**	**$800**	**$2,320**	**$1,000**	**$1,150**	**$1,288**	**$1,443**	**$4,881**	**$1,587**	**$1,745**	**$1,920**	**$2,112**	**$7,364**	**$9,558**	**$11,062**
Server Revenues	**$12,817**	**$13,419**	**$3,916**	**$5,075**	**$5,693**	**$6,397**	**$21,080**	**$6,784**	**$7,197**	**$7,635**	**$8,177**	**$29,793**	**$8,173**	**$8,587**	**$9,196**	**$9,568**	**$35,523**	**$40,359**	**$44,112**
Server Units ('000s)	15,315	16,709	4,354	4,397	4,661	4,941	18,353	5,039	5,191	5,346	5,560	21,137	5,393	5,501	5,721	5,779	22,395	22,682	22,975
Server ASP	$824	$791	$887	$1,142	$1,221	$1,295	$1,143	$1,334	$1,374	$1,415	$1,457	$1,396	$1,501	$1,546	$1,592	$1,640	$1,571	$1,761	$1,898
Accelerators	**$200**	**$200**	**$53**	**$55**	--	--	**$108**	**$64**	**$67**	**$71**	**$74**	**$276**	**$78**	**$82**	**$86**	**$90**	**$335**	**$408**	**$495**
Intel Foundry	**$17,317**	**$17,826**	**$5,421**	**$5,765**	**$6,365**	**$7,037**	**$24,588**	**$6,945**	**$7,060**	**$7,422**	**$7,598**	**$29,025**	**$7,600**	**$8,190**	**$8,861**	**$9,301**	**$33,950**	**$45,310**	**$62,637**
Q/Q			20%	6%	10%	11%		(1%)	2%	5%	2%		0%	8%	8%	5%			
Y/Y	(8%)	3%	16%	31%	50%	56%	38%	28%	22%	17%	8%	18%	9%	16%	19%	22%	17%	33%	38%
Internal Foundry Revenue	**$17,288**	**$17,683**	**$5,251**	**$5,477**	**$6,025**	**$6,627**	**$23,380**	**$6,495**	**$6,560**	**$6,822**	**$6,958**	**$26,835**	**$6,750**	**$7,020**	**$7,371**	**$7,371**	**$28,510**	**$28,510**	**$28,510**
External Foundry Revenue	**$29**	**$143**	**$170**	**$288**	**$340**	**$410**	**$1,208**	**$450**	**$500**	**$600**	**$640**	**$2,190**	**$850**	**$1,170**	**$1,490**	**$1,930**	**$5,440**	**$16,800**	**$34,127**
Operating Income	**($13,291)**	**($10,318)**	**($2,320)**	**($2,320)**	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
Operating Margin	(77%)	(58%)	(43%)	(40%)	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
All Other	**$3,601**	**$3,563**	**$628**	**$701**	**$702**	**$705**	**$2,736**	**$802**	**$717**	**$719**	**$722**	**$2,960**	**$909**	**$808**	**$810**	**$813**	**$3,340**	**$3,677**	**$4,061**
Q/Q			9%	12%	0%	0%		14%	(11%)	0%	0%		26%	(11%)	0%	0%			
Y/Y		(1%)	(33%)	(33%)	(29%)	23%	(23%)	28%	2%	2%	2%	8%	13%	13%	13%	13%	13%	10%	10%
Operating Income	**($57)**	**$264**	**$103**	**($139)**	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
Operating Margin	(2%)	7%	16%	(20%)	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--

Source: Company data, FactSet, UBS estimates

Figure 34: ARM - Summary Model

Arm Holdings Plc. (ARM) - Key Metrics (in $MM except EPS and units)	C2022A	C2023A	C2024A	C2025A	FQ4:26A CQ1A	FQ1A CQ2A	Estimates FQ2E CQ3E	FQ3E CQ4E	C2026E	FQ4:27E CQ1E	FQ1E CQ2E	FQ2E CQ3E	FQ3E CQ4E	C2027E	FQ4:28E CQ1E	FQ1E CQ2E	FQ2E CQ3E	FQ3E CQ4E	C2028E	C2029E	C2030E	FISCAL YEAR SUMMARY F2023A	F2024A	F2025A	F2026E	F2027E
Revenue	**$2,703**	**$2,938**	**$3,694**	**$4,671**	**$1,490**	**$1,289**	**$1,379**	**$1,552**	**$5,711**	**$1,923**	**$1,710**	**$1,855**	**$2,162**	**$7,650**	**$2,505**	**$2,554**	**$2,692**	**$2,893**	**$10,644**	**$14,648**	**$24,698**	**$2,679**	**$3,233**	**$4,007**	**$4,920**	**$6,143**
Y/Y Growth		8.7%	25.7%	26.4%	20.1%	22.4%	21.5%	25.0%	22.3%	29.0%	32.6%	34.5%	39.3%	34.0%	30.3%	49.4%	45.1%	33.8%	39.1%	37.6%	68.6%	-0.9%	20.7%	23.9%	22.8%	24.9%
Gross Profit (Non-GAAP)	$2,592	$2,837	$3,607	$4,588	$1,465	$1,264	$1,354	$1,527	$5,611	$1,818	$1,628	$1,730	$1,951	$7,126	$2,212	$2,223	$2,358	$2,501	$9,294	$12,070	$17,305	$2,580	$3,131	$3,926	$4,832	$5,964
Gross Margin	**95.9%**	**96.6%**	**97.6%**	**98.2%**	**98.3%**	**98.1%**	**98.2%**	**98.4%**	**98.2%**	**94.6%**	**95.2%**	**93.2%**	**90.2%**	**93.2%**	**88.3%**	**87.0%**	**87.6%**	**86.4%**	**87.3%**	**82.4%**	**70.1%**	**96.3%**	**96.8%**	**98.0%**	**98.2%**	**97.1%**
Opex (Non-GAAP)	**$1,839**	**$1,821**	**$1,994**	**$2,549**	**$734**	**$733**	**$780**	**$822**	**$3,069**	**$866**	**$900**	**$931**	**$966**	**$3,662**	**$999**	**$1,038**	**$1,077**	**$1,117**	**$4,232**	**$4,909**	**$5,564**	**$1,797**	**$1,723**	**$2,049**	**$2,717**	**$3,202**
Y/Y Growth			9.5%	27.8%	29.7%	18.4%	20.4%	14.8%	20.4%	18.0%	22.7%	19.3%	17.4%	19.3%	15.4%	15.4%	15.7%	15.7%	15.5%	16.0%	13.3%	-2.9%	-4.1%	18.9%	32.6%	17.8%
Operating Income (Non-GAAP)	$753	$1,016	$1,613	$2,039	$731	$531	$574	$705	$2,541	$952	$728	$799	$986	$3,464	$1,213	$1,185	$1,281	$1,384	$5,062	$7,161	$11,741	$783	$1,408	$1,877	$2,115	$2,762
Operating Margin	**27.9%**	**34.6%**	**43.7%**	**43.7%**	**49.1%**	**41.2%**	**41.6%**	**45.4%**	**44.5%**	**49.5%**	**42.6%**	**43.1%**	**45.6%**	**45.3%**	**48.4%**	**46.4%**	**47.6%**	**47.8%**	**47.6%**	**48.9%**	**47.5%**	**29.2%**	**43.6%**	**46.8%**	**43.0%**	**45.0%**
Non-GAAP EPS	**$ 0.62**	**$ 0.94**	**$ 1.45**	**$ 1.72**	**$ 0.60**	**$ 0.45**	**$ 0.48**	**$ 0.58**	**$ 2.10**	**$ 0.78**	**$ 0.60**	**$ 0.66**	**$ 0.81**	**$ 2.85**	**$ 0.99**	**$ 0.97**	**$ 1.05**	**$ 1.13**	**$ 4.14**	**$ 5.80**	**$ 9.31**	**$ 0.64**	**$ 1.28**	**$ 1.64**	**$ 1.77**	**$ 2.28**
Free Cash Flow																										
CFFO	$256	$906	$806	$1,522	$260	$902	$510	$536	$2,208	$538	$1,013	$686	$684	$2,921	$859	$1,156	$1,159	$1,200	$4,374	$5,644	$8,910	$739	$1,090	$397	$1,524	$2,486
Capex and intangibles	($104)	($182)	($233)	($629)	($108)	($237)	($177)	($200)	($722)	($247)	($203)	($220)	($256)	($926)	($297)	($277)	($292)	($314)	($1,180)	($1,474)	($2,239)	($133)	($183)	($298)	($642)	($861)
FCF	**$152**	**$724**	**$573**	**$893**	**$152**	**$665**	**$332**	**$337**	**$1,486**	**$291**	**$810**	**$466**	**$427**	**$1,995**	**$562**	**$879**	**$867**	**$886**	**$3,194**	**$4,170**	**$6,671**	**$606**	**$907**	**$99**	**$882**	**$1,625**
% of revenue	(6%)	-24.6%	-15.5%	-19.1%	-10.2%	-51.6%	-24.1%	-21.7%	-26.0%	-15.1%	-47.4%	-25.1%	-19.8%	-26.1%	-22.4%	-34.4%	-32.2%	-30.6%	-30.0%	-28.5%	-27.0%	22.6%	28.1%	2.5%	17.9%	26.5%
Y/Y Growth		376.3%	-20.9%	55.8%	-6.7%	343.3%	-19.1%	99.3%	66.4%	91.5%	21.9%	40.3%	26.9%	34.2%	93.0%	8.5%	85.9%	107.2%	60.1%	30.6%	60.0%	75.1%	49.7%	-89.1%	790.9%	84.3%
Revenue Reporting																										
Royalty revenue	**$1,725**	**$1,662**	**$2,075**	**$2,549**	**$671**	**$715**	**$708**	**$807**	**$2,901**	**$831**	**$856**	**$899**	**$975**	**$3,561**	**$1,019**	**$1,029**	**$1,112**	**$1,201**	**$4,361**	**$5,332**	**$6,377**	**$1,675**	**$1,802**	**$2,168**	**$2,613**	**$3,061**
% of revenue	64%	56.6%	56.2%	54.6%	45.0%	55.5%	51.3%	52.0%	50.8%	43.2%	50.1%	48.5%	45.1%	46.6%	40.7%	40.3%	41.3%	41.5%	41.0%	36.4%	25.8%	62.5%	55.7%	54.1%	53.1%	49.8%
Y/Y growth		-3.7%	24.8%	22.8%	10.5%	22.2%	14.2%	9.5%	13.8%	23.9%	19.7%	27.0%	20.9%	22.8%	22.6%	20.2%	23.7%	23.1%	22.4%	22.3%	19.6%	7.2%	7.6%	20.3%	20.5%	17.1%
Q/Q growth					-9.0%	6.6%	-1.0%	14.0%		3.0%	3.0%	5.0%	8.5%		4.5%	1.0%	8.0%	8.0%								
of which ARM v7			90.0%	25.0%																		43.0%	36.0%	28.6%	25.0%	
of which ARM v8			135.0%	44.0%																		49.0%	46.5%	44.7%	15.0%	
of which ARM v9		25.0%	95.0%	31.0%																		8.0%	17.5%	26.7%	60.0%	
License and other revenue	**$978**	**$1,276**	**$1,619**	**$2,122**	**$819**	**$574**	**$672**	**$745**	**$2,810**	**$991**	**$754**	**$776**	**$807**	**$3,328**	**$936**	**$908**	**$890**	**$899**	**$3,634**	**$3,659**	**$3,693**	**$1,004**	**$1,431**	**$1,839**	**$2,307**	**$2,982**
% of Total revenue	36%	43.4%	43.8%	45.4%	55.0%	44.5%	48.7%	48.0%	49.2%	51.6%	44.1%	41.8%	37.3%	43.5%	37.4%	35.6%	33.1%	31.1%	34.1%	25.0%	15.0%	37.5%	44.3%	45.9%	46.9%	48.5%
Y/Y growth		30.5%	26.9%	31.1%	29.2%	22.6%	30.4%	47.6%	32.4%	21.1%	31.3%	15.6%	8.3%	18.4%	-5.6%	20.5%	14.7%	11.4%	9.2%	0.7%	0.9%	-12.0%	42.5%	28.5%	25.4%	29.3%
Q/Q growth					62.2%	-29.9%	17.0%	11.0%		33.0%	-24.0%	3.0%	4.0%		16.0%	-3.0%	-2.0%	1.0%								
Annualized Contract Value (ACV)	$4,171	$4,346	$4,898	$6,113																		$1,030	$1,182	$1,365		
% Y/Y		4.2%	12.7%	24.8%																		-	14.8%	15.5%		
% Q/Q																										
Remaining Performance Obligation (RPO)	$6,993	$8,177	$9,362	$8,852																		$1,712	$2,484	$2,226		
% Y/Y		16.9%	14.5%	-5.4%																		-2.2%	45.1%	-10.4%		
% Q/Q																										
implied orders	$408	$1,528	$1,162	$1,043	$518	$935	$1,020	$1,132	$3,604	$1,505	$1,151	$1,186	$1,233	$5,076	$1,431	$1,397	$1,369	$1,383	$5,580	$5,655	$5,746	$530	$1,823	$1,163	$1,143	$4,592
book-to-bill	2.4x	0.8x	1.4x	2.0x	0.9x	1.1x	1.0x	1.0x	0.8x	1.0x	1.0x	1.0x	1.0x	0.7x	1.0x	1.0x	1.0x	1.0x	0.7x	0.6x	0.6x	0.9x	1.7x	0.8x	0.9x	1.0x

Source: Company data, FactSet, UBS estimates



ESTIMATE CHANGES

Figure 35: INTC - New vs Old Estimates

	C2026E			C2027E			C2028E		
	New	Old	Change	New	Old	Change	New	Old	Change
Revenue ($MM)	$64,448	$64,448	-	$75,192	$75,192	-	$87,382	$87,382	-
Y/Y Growth	21.9%	21.9%		16.7%	16.7%		16.2%	16.2%	
Gross Margin	42.0%	42.0%	-	43.5%	43.5%	-	44.4%	44.4%	-
Operating Margin	15.7%	15.6%	+0.2%	19.1%	19.1%	-	21.4%	21.4%	-
Non-GAAP EPS	$ 1.57	$ 1.55	+1.1%	$ 2.20	$ 2.19	+0.0%	$ 2.94	$ 2.94	+0.0%
FCF	$7,344	$7,244	+1.4%	($959)	($962)	-0.3%	($4,077)	($4,080)	-0.1%
Gross Capex	$20,215	$20,215	-	$28,500	$28,500	-	$40,000	$40,000	-
Segment Revenues ($MM)									
Intel Products	$60,504	$60,504	-	$70,042	$70,042	-	$78,602	$78,602	-
Y/Y Growth	23.1%	23.1%		15.8%	15.8%		12.2%	12.2%	
Client Computing (CCPG)	$34,255	$34,255	-	$32,190	$32,190	-	$32,181	$32,181	-
Y/Y Growth	6.3%	6.3%		-6.0%	-6.0%		0.0%	0.0%	
Datacenter and AI (DCAI)	$26,357	$26,357	-	$37,853	$37,853	-	$46,421	$46,421	-
Y/Y Growth	55.8%	55.8%		43.6%	43.6%		22.6%	22.6%	
Intel Foundry	$24,588	$24,588	-	$29,025	$29,025	-	$33,950	$33,950	-
Y/Y Growth	37.9%	37.9%		18.0%	18.0%		17.0%	17.0%	
of which external revenue	$1,208	$1,208		$2,190	$2,190		$5,440	$5,440	
All Other	$2,736	$2,736	-	$2,960	$2,960	-	$3,340	$3,340	-
Y/Y Growth	-23.2%	-23.2%		8.2%	8.2%		12.8%	12.8%	
Intersegment Eliminations	($23,380)	($23,380)		($26,835)	($26,835)		($28,510)	($28,510)	

Source: Company reports, UBS estimates.

Figure 36: UBSe vs. Street

	Actual 2025	Actual 2026/1F	Actual 2026/2F	Estimate 2026/3F	Estimate 2026/4F	Estimate 2026	Estimate 2027/1F	Estimate 2027/2F	Estimate 2027/3F	Estimate 2027/4F	Estimate 2027	Estimate 2028/1F	Estimate 2028/2F	Estimate 2028/3F	Estimate 2028/4F	Estimate 2028
Revenues ($MM)																
UBS New	52,853	13,577	16,128	16,974	17,769	64,448	17,886	18,134	19,202	19,970	75,192	19,978	21,156	22,673	23,482	87,288
UBS Old				16,974	17,769	64,448	17,886	18,134	19,202	19,970	75,192	19,899	21,155	22,757	23,571	87,382
Consensus				16,518	17,119	63,341	16,928	17,858	18,799	19,584	73,169	19,507	20,357	21,513	22,612	83,990
Delta				2.8%	3.8%	1.7%	5.7%	1.5%	2.1%	2.0%	2.8%	2.4%	3.9%	5.4%	3.8%	3.9%
Non-GAAP EPS																
UBS New	$ 0.42	$ 0.29	$ 0.43	$ 0.39	$ 0.45	$ 1.57	$ 0.46	$ 0.50	$ 0.59	$ 0.64	$ 2.20	$ 0.64	$ 0.69	$ 0.78	$ 0.83	$ 2.94
UBS Old				$ 0.39	$ 0.45	$ 1.57	$ 0.46	$ 0.50	$ 0.59	$ 0.64	$ 2.20	$ 0.63	$ 0.69	$ 0.79	$ 0.83	$ 2.94
Consensus				$ 0.40	$ 0.44	$ 1.56	$ 0.41	$ 0.49	$ 0.58	$ 0.65	$ 2.13	$ 0.67	$ 0.75	$ 0.85	$ 0.94	$ 3.22
Delta				(3.5%)	2.4%	0.8%	13.7%	1.8%	0.8%	(1.6%)	2.8%	(5.4%)	(7.5%)	(8.3%)	(12.3%)	(8.7%)

Source: FactSet, Company reports, UBS estimates

Valuation and Multiple Justification

We lower our PT from $103 to $100 based on our SOTP valuation framework. We value each business separately. We value INTC's product business assuming a normalized EBIT margin of 25%. We apply a 27x EV/ EBIT multiple and net off the non-controlling interest (basically debt represented by the SCIP agreements for Arizona and Ireland). Our framework looks at an independent valuation for Intel's manufacturing assets based on how other global foundry comps are valued per installed wafer of capacity. We value INTC's Foundry business based on an EV/k-wpm basis of $8/sh. We believe this to be the best metric to assign a market-implied valuation to the company's Foundry aspirations, but apply a ~45% discount vs. TSMC to reflect node and yield gaps, customer stickiness, and utilization risk.

Figure 37: Intel SOTP Valuation (New)

	Total $MM	Per Share	Comments
Intel Products			
Revenue	70,042		
Normalized EBIT margin	25%		
Normalized EBIT	17,511		
Peer EV/EBIT	27x		
Products SoTP Value	**465,783**	**$86**	
Intel Foundry			
Intel Capacity (k-wsm)	364		
Peer Value per k-wsm	160		
Gross Foundry Value	**58,393**	**$11**	
Noncontrolling Interest	(14,200)	($3)	Brookfield only, at historical cost
Net Foundry SoTP Value	**44,193**	**$8**	
Other			
Market Cap	14,310		Based on valuation and market cap.
Intel Stake	61%		
Other SoTP Value	**8,736**	**$2**	
TOTAL			
SoTP Enterprise Value	**518,712**	**$95**	
Net Debt	27,184	$5	
SoTP Equity Value	**545,896**	**$100**	

Source: UBS estimates, FactSet.

Figure 38: Intel SOTP Valuation (Old)

	Total $MM	Per Share	Comments
Intel Products			
Revenue	70,042		
Normalized EBIT margin	25%		
Normalized EBIT	17,511		
Peer EV/EBIT	27x		
Products SoTP Value	**476,289**	**$88**	
Intel Foundry			
Intel Capacity (k-wsm)	364		
Peer Value per k-wsm	166		
Gross Foundry Value	**60,445**	**$11**	
Noncontrolling Interest	(14,200)	($3)	Brookfield only, at historical cost
Net Foundry SoTP Value	**46,245**	**$9**	
Other			
Market Cap	14,888		Based on valuation and market cap.
Intel Stake	62%		
Other SoTP Value	**9,198**	**$2**	
TOTAL			
SoTP Enterprise Value	**531,732**	**$98**	
Net Debt	27,184	$5	
SoTP Equity Value	**558,916**	**$103**	

Source: UBS estimates, FactSet.

Figure 39: Intel Foundry Comps

Ticker	Company Name	Capacity	EV	EV/k-wsm
2330-TW	TSMC	1,642	1,802,505	604
005930-KR	Samsung Foundry	1,685	969,422	178
2303-TW	UMC	440	45,618	104
981-HK	SMIC	507	65,426	129
GFS-US	GlobalFoundries	292	23,736	81
1347-HK	Hua Hong	247	18,747	76
688249-CN	Vanguard (200mm)	146	8,752	60
6770-TW	PSMC (200mm)	163	8,379	51
	Simple Average	640	367,823	**160**
	Median	366	34,677	**92**

Source: UBS research, FactSet. Note: Samsung EV adjusted for Foundry capacity, TSMC 45% haircut to EV reflect node and yield gaps, customer stickiness, and utilization risk.

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.

AMD: 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.

ARM: Because Arm is growing far faster than many of its competitors, we derive our price target based on a PEG methodology which effectively normalizes P/E multiples for growth. We use a 1.8x target PEG, a modest premium to semiconductor peers owing to Arm's restricted float but in line with Software and far below near-peer semiconductor EDA firms. Our LT EPS CAGR for Arm is 42.2%; applying our 1.8x PEG multiple results in a 78x P/E multiple. Then, 78x NTM P/E times 2028E EPS of $4.14 in turn results in a $320 price target.

INTC: Our PT is based on a SOTP analysis of the five contributing segments (Intel Products, Intel Foundry, Altera, Mobileye and All Other) each one valued using a tailored P/S multiple applied to F2026E revenues. NVDA has built a formidable moat for new compute-intensive workloads in the data center and could ultimately leverage its GPU architecture to more broadly displace INTC. New client and server CPUs from AMD also present a threat that we could be underestimating.

MU: The memory industry is highly cyclical and our PT is based solely on a P/E framework. A weaker macro environment could drive demand destruction and pressure memory pricing. More severe-than-expected DRAM and NAND ASP declines, slower technology migrations, yield challenges, or process-scaling difficulties could create downside to estimates. In AI, advances in model efficiency and memory optimization techniques (e.g., KV-cache compression, quantization, sparsity, and distillation) could reduce memory intensity and temper long-term demand growth. Conversely, delays in supply additions, technology transitions, or yield ramps could constrain supply growth, support stronger ASPs, and drive upside to our estimates and shares.