

US Semiconductors

Rise of the Agents: raising CPU TAM to \$170bn

Price Objective Change

CPU reborn: Agentic AI drives 5x demand

Based on our analysis and recent industry discussions at BofA Global Tech Conference (see our [conf takeaways](#)), we raise our CY2030E server CPU TAM to \$170bn+ (from \$125bn), implying nearly 5x growth and +37% CAGR over CY25-30E (vs. +29% prior). We view the emergence of agentic AI as a powerful demand accelerant that expands the CPU opportunity and lifts both x86 incumbents and ARM challengers. Against this backdrop, we raise our **AMD** PO to \$560 from \$500 on higher CPU/GPU estimates, **ARM** to \$335 from \$245 on greater I-t chiplet potential; **INTC** is double-upgraded to Buy with a \$135 PO (see our), with higher estimates now reflecting both n-t CPU and I-t foundry upsides. **NVDA** remains our top sector pick on full-stack AI leadership, benefiting from tight CPU-GPU-networking integration, also poised to benefit from agentic CPU. **QCOM** is expected to announce new AI CPUs at its Jun-24 AI Day in NYC, but we maintain Underperform given tough competition and limited SAM.

Agentic workloads shift complexity to CPUs from XPU

Agentic AI differs from traditional genAI by shifting from a single prompt-response workflow to a multi-step system that plans, reasons, retrieves info, uses tools, executes code simultaneously. While XPUs remain critical for inference, many orchestration and decision-making functions are latency-sensitive, sequential, and I/O-intensive—making them better suited for CPUs. Our new CPU model uniquely segments the \$170bn TAM by application (1. standalone agentic CPU; 2. head/compute node CPU for AI clusters; 3. traditional/laaS), as well as by vendor (AMD, INTC, NVDA, ARM merchant, ARM custom). We expect head/compute nodes in AI clusters to use higher frequency and stronger but fewer cores, while agentic and traditional applications will rely on higher core counts.

Double-upgrading INTC, while AMD top CPU pick

Our **INTC** double-upgrade is based on higher confidence in INTC's opportunity to help address industry constraints in leading edge wafers/packaging, plus larger agentic CPU TAM. We now see CY30 EPS power \$6+ vs \$3-4 prior, though execution is key. Recent CDNS 14A node IP sign-up and Terafab engagements are additional supportive data points. For **AMD**, we raise est/PO to \$560, top CPU pick on incumbency + pipeline + upcoming AI day (Venice launch). **NVDA** remains top sector pick on full-stack leadership.

ARM/QCOM: CPU share gainers but flag val'n/competition

For **ARM**, we now use a FY31/CY30E sum-of-parts valuation, see total company value of ~\$335 (\$229 IP + \$106 chip), and we see the stock as fairly valued. **QCOM** has potential to announce new AI CPUs and lead customers at its Jun-24 AI Day, but we flag tough CPU competition amid already established incumbents.

Key debate: TAM substitution vs expansion

Debate centers on whether CPUs take share from GPUs/accelerators or contribute to overall AI TAM expansion. We see expansion dominating: CPU demand scales with system complexity, as is the case with reasoning/agentic AI. Additional debates include durability of the 30%+ AI capex CAGR and mix shifts to custom silicon (CPUs and XPUs).

BofA Securities does and seeks to do business with issuers covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision. Refer to important disclosures on page 29 to 33. Analyst Certification on page 27. Price Objective Basis/Risk on page 26.

12983146

Timestamp: 11 June 2026 06:00AM EDT

11 June 2026

Equity
United States
Semiconductors

Vivek Arya
Research Analyst
BofAS
vivek.arya@bofa.com

Duksan Jang
Research Analyst
BofAS
duksan.jang@bofa.com

Michael Mani
Research Analyst
BofAS
michael.mani@bofa.com

Liam Pharr
Research Analyst
BofAS
liam.pharr@bofa.com

Exhibit 1: PO/Rating Changes

PO/Rating Changes

PO/Rating Changes			
Ticker	OLD	NEW	RATING
AMD	\$500	\$560	BUY
ARM	\$245	\$335	NEUTRAL
INTC	\$96	\$135	BUY from U/P
NVDA	\$350	\$350	BUY
QCOM	\$165	\$165	U/P

Source: BofA Global Research

BofA GLOBAL RESEARCH

For detailed company analysis, PO/rating change and valuation discussion, see pg. 10-17.

See glossary on page 23.

Contents

AI Compute Table of Hierarchy	3
\$170bn+ Server CPU TAM	5
Market share outlook analysis	6
Traditional vs. AI Head Node (Compute) vs. Agentic Rack	7
CPUs now 8%+ of Data Center Systems TAM	9
AMD: Maintain Buy, raise PO to \$560 from \$500	10
AMD CPU Advantages	10
Upcoming catalyst: Advancing AI 2026 (July)	11
INTC: Upgrade to Buy from U/P, PO \$135	12
Fully established IDM by CY30, EPS power \$6+	12
ARM: Reit. Neutral, raise PO to \$335	14
Sum-of-Parts Valuation: IP & Chip Businesses	14
NVDA: Maintain Buy, top AI sector pick	15
Rising content, increasing opp'ty, solid supply visibility	15
Vera Rubin now in full production, lowest token cost	15
Vera CPU: head node vs. agentic AI likely 50/50 opp'ty	16
NVDA content \$/GW rising materially gen-over-gen	16
RTX Spark for personal agentic AI, for top 10% Windows	16
QCOM: U/P on competitive landscape	17
Role of CPUs in Agentic AI	18
CPUs in training	18
CPUs in inference	19
CPUs in Agentic AI	20
Agentic AI is system-led, not CPU or GPU-led	21
AI CPUs expand system TAM, not replace	21



AI Compute Table of Hierarchy


Below, we present our latest table of hierarchy for global data center systems TAM of ~\$2.1Tn, of which CPUs make up ~\$170bn of the total value.

Within this, we break out CPUs in three different categories (more detail in following pages):

1. **Traditional/on-prem/multi-tenant cloud CPUs:** ~\$30bn
2. **AI cluster compute/head node CPUs:** ~\$70bn
3. **AI agentic standalone node CPUs:** ~\$70bn

Exhibit 2: Of the total \$2.1Tn CY30 DC Systems TAM, we see AI CPUs representing ~\$140bn of the value (split roughly 50/50 across compute/head node and agentic AI node), non-AI CPUs ~\$30bn

Data Center Systems TAM – Breakout by: AI vs. non-AI; AI CPUs vs. non-AI CPUs



Source: BofA Global Research


BofA GLOBAL RESEARCH

The below infographic highlights how the role of the CPU evolves across different server architectures:

1. **Traditional Server:** The CPU is the primary "jack-of-all-trades," managing the OS, applications, and direct I/O for general workloads. (~\$30bn CY30E TAM)
2. **AI Server Cluster:** The CPU's role bifurcates into a Head Node for cluster orchestration and management, and a Compute Node focused on supervising and feeding data to GPUs. (~\$70bn CY30E TAM)
3. **Agentic Server:** The CPU becomes the central orchestrator for autonomous actions, managing the high-level reasoning loops, memory state, and tool integration required for AI agents. (~\$70bn CY30E TAM)

Exhibit 3: We flag three key roles of CPUs in servers: 1. Traditional, 2. AI compute/head node, 3. AI agentic node

Role of the CPU in different server architectures



Source: BofA Global Research

BofA GLOBAL RESEARCH

\$170bn+ Server CPU TAM

Given the rising role of CPUs in AI head node (compute node) and agentic AI, we now expect server CPU TAM to reach ~\$170bn by CY30E from ~\$35bn in CY25 (+37% CAGR), versus prior ~\$125bn outlook (+29% CAGR).

Notably, we break out the market into 4 different vendor/types: **1) INTC, 2) AMD, 3) ARM (merchant), and 4) ARM (custom/ASIC).**

- For **ARM merchant CPUs**, we now assume ASPs that are ~1.25x higher than average AMD CPUs, given ARM merchant almost only caters to hyperscale/AI applications, versus AMD mostly hyperscale with a bit of enterprise exposure.
- For **ARM custom CPUs**, we now assume ASPs that are ~75% of typical AMD solutions, given continued mix shift toward hyperscale solutions (higher ASPs), but merchant server CPUs typically carry 40-60% GM (i.e. 40-60% is COGS or the cost that AWS/Google would theoretically be paying for their custom CPUs)

Exhibit 4: We see total server CPU TAM reaching \$170bn+ by CY30 from \$35bn in CY25, growing at +37% CAGR

Server CPU TAM (by revenue, units, ASP)

Server CPU TAM	2023	2024	2025	2026E	2027E	2028E	2029E	2030E	CY25-30 CAGR
Revenue (\$bn)									
Total	\$19.5	\$22.5	\$35.2	\$59.6	\$89.8	\$110.4	\$138.2	\$170.1	37.1%
INTC	\$13.2	\$13.6	\$14.4	\$19.8	\$24.7	\$29.6	\$35.1	\$41.2	23.4%
AMD	\$5.6	\$7.1	\$9.6	\$16.4	\$22.2	\$28.6	\$35.7	\$43.6	35.4%
ARM	\$0.7	\$1.8	\$11.2	\$23.4	\$42.9	\$52.2	\$67.5	\$85.4	50.0%
ARM Merchant	\$0.0	\$0.7	\$10.3	\$19.5	\$34.8	\$38.2	\$47.9	\$60.6	42.6%
NVDA		\$0.7	\$10.3	\$19.5	\$33.7	\$35.4	\$39.0	\$42.2	32.7%
ARM					\$0.7	\$1.8	\$5.9	\$12.3	
QCOM					\$0.4	\$0.9	\$3.0	\$6.1	
ARM Custom (Graviton, Axion, Cobalt, etc.)	\$0.7	\$1.0	\$1.0	\$3.9	\$8.1	\$14.1	\$19.6	\$24.8	91.0%
Units (mn)									
Total	24.5	24.2	29.2	39.3	49.0	57.1	66.0	74.7	20.7%
INTC	18.5	16.9	17.9	21.1	23.2	24.9	26.7	28.6	9.7%
AMD	4.8	5.5	7.2	10.4	12.5	14.5	16.3	18.1	20.2%
ARM	1.2	1.8	4.0	7.9	13.3	17.7	22.9	28.0	47.4%
ARM Merchant	0.0	0.2	2.6	4.6	7.2	8.2	10.9	14.2	40.9%
NVDA	0.0	0.2	2.6	4.6	6.7	7.1	7.7	8.1	25.9%
ARM	0.0	0.0	0.0	0.0	0.3	0.7	2.2	4.1	
QCOM	0.0	0.0	0.0	0.0	0.2	0.4	1.1	2.0	
ARM Custom (Graviton, Axion, Cobalt, etc.)	1.2	1.6	1.5	3.3	6.1	9.5	12.0	13.8	56.6%
ASP (\$)									
Total	\$797	\$931	\$1,204	\$1,517	\$1,833	\$1,934	\$2,095	\$2,277	13.6%
INTC	\$716	\$804	\$801	\$941	\$1,065	\$1,186	\$1,311	\$1,442	12.5%
AMD	\$1,164	\$1,299	\$1,322	\$1,581	\$1,778	\$1,976	\$2,183	\$2,401	12.7%
ARM	\$589	\$998	\$2,790	\$2,972	\$3,225	\$2,956	\$2,950	\$3,049	1.8%
ARM Merchant		\$4,000	\$4,000	\$4,240	\$4,820	\$4,658	\$4,393	\$4,255	1.2%
NVDA		\$4,000	\$4,000	\$4,240	\$5,000	\$5,000	\$5,100	\$5,202	5.4%
ARM					\$2,249	\$2,476	\$2,729	\$3,002	
QCOM					\$2,249	\$2,476	\$2,729	\$3,002	
ARM Custom (Graviton, Axion, Cobalt, etc.)	\$589	\$649	\$666	\$1,191	\$1,336	\$1,484	\$1,637	\$1,801	22.0%
Value Share (%)									
Total	100%	100%	100%	100%	100%	100%	100%	100%	
INTC	68%	61%	41%	33%	28%	27%	25%	24%	
AMD	29%	32%	27%	28%	25%	26%	26%	26%	
ARM	4%	8%	32%	39%	48%	47%	49%	50%	
ARM Merchant	0%	3%	29%	33%	39%	35%	35%	36%	
NVDA		3%	29%	33%	38%	32%	28%	25%	
ARM					1%	2%	4%	7%	
QCOM					0%	1%	2%	4%	
ARM Custom (Graviton, Axion, Cobalt, etc.)	4%	5%	3%	7%	9%	13%	14%	15%	

Source: BofA Global Research estimates, Mercury Research, IDC




Market share outlook analysis

From a market share perspective, we continue to expect ARM to be the fastest gainer through CY30 as multiple new merchant (i.e. ARM AGI, NVDA Vera, QCOM CPU) and custom (i.e., AWS Graviton 5, Google Axion, Microsoft Cobalt) programs ramp.

Exhibit 5: We expect INTC/AMD to each represent ~25% value share, ARM merchant ~35%, ARM custom ~15% by CY30

Server CPU TAM (\$mn) and Value Market Share Outlook



Source: BofA Global Research estimates, Mercury Research, IDC

BofA GLOBAL RESEARCH

- **INTC (~25% value share):** We see share declining to just ~24% of total server CPU value by CY30 from ~41% in CY25, though we expect INTC's continued relative strength in traditional enterprise workloads.

However, we see continued \$ revenue growth of +20% CAGR between CY26-30, particularly as INTC navigates through the industry-wide supply shortages, benefits from ASP expansion. By CY28, next-next-gen Coral Rapids should ramp and could begin to narrow the performance gap vs. competition.

- **AMD (~25% value share):** We see share of ~25-27% generally through CY30, given AMD's continued performance lead at the high-end (high-frequency, high-core-count AI head node and agentic node), offset by relative stronger ramps of ARM-based processors launching beginning CY27E.

Particularly, we flag AMD's leading core count portfolio, given higher core density within a fixed power envelope increases overall aggregate service throughput. 6th gen EPYC Venice will top out at 256 cores, while the current 5th gen Turin tops out at 192 cores, vs. NVDA Vera at 88 cores and INTC's current Granite Rapids topping at 128 cores and upcoming Diamond Rapids likely maxing at 192 cores.

Similar to INTC, we expect AMD to continue benefiting from the x86 ecosystem's advantages in security and RAS features, particularly important in agentic AI workloads.


- **ARM (merchant; ~35% value share):** We see share rising to ~36% by CY30 from ~29% in CY25, led by a continued ramp of NVDA CPUs (in both compute/agentic nodes), as well as new ARM/QCOM merchant solutions ramping in CY27-28.
- **ARM (custom; ~15% value share):** We see significant share gain toward 15% by CY30 from <5% in CY23-25 as new custom ARM-based CPU solutions launch and ramp at multiple hyperscalers. Every year, the ARM content/ASP grows as core counts and performance per core grow.



On a unit share basis, we expect INTC to remain the largest at ~38% share by CY30.

Exhibit 6: We expect INTC to represent ~38% unit share, AMD ~24%, ARM merchant ~19%, and ARM custom ~18% by CY30

Server CPU TAM (mn units) and Unit Market Share Outlook



Source: BofA Global Research estimates, Mercury Research, IDC

BofA GLOBAL RESEARCH

Traditional vs. AI Head Node (Compute) vs. Agentic Rack

In the following Exhibit, we also present the \$170bn CY30 server CPU TAM in 3 different key workloads: **1) traditional/infrastructure-as-a-service**, **2) AI compute node or head node**, and **3) AI agentic standalone CPU node**.

Particularly, we expect the two AI applications to represent ~50%/50% of the total CPU demand in AI data centers (~\$140bn) by CY30, as NVDA recently highlighted.

- **Traditional/laaS:** This includes the most traditional and conventional CPU use cases for enterprises on-prem and on public cloud.

This segment represented the vast majority of the server CPU market in 2010-2020, growing ~11-12% CAGR during that period. For the next five years, we expect this growth rate to continue, at ~11-12% sales CAGR range, or units growing ~7-9% CAGR, with ASPs growing ~3-5% CAGR.

- **AI compute node, head node:** This includes CPUs that are attached to GPUs/XPUs and are acting as supporting roles within the compute rack (most pronouncedly NVDA's NVL72 rack, AMD's MI455 'Helios' rack, TPUv8 racks). In other words, these CPUs are used as part of the scale-up domain, facilitating XPU compute.

While the vast majority of initial AI CPU use cases (~2022-2026) were compute/head nodes, we anticipate a mix shift toward agentic AI standalone CPU racks over time (2026-2030), eventually the two AI use cases likely representing ~50/50 of all AI workloads.

- **AI agentic standalone CPU node/rack:** This includes the new CPU-only standalone racks, such as the NVDA Vera CPX rack which is part of the overall 40-rack Vera Rubin SuperPOD. These CPU-only racks are highly optimized for agentic AI workloads which require sequential processing in tandem with GPU/XPU's parallel compute processing, and are responsible for orchestration, memory control, I/O, etc.



Exhibit 7: Of the \$170bn CY30 CPU TAM, we see ~\$30bn for traditional/laaS, ~\$70bn for AI compute/head node, ~\$70bn for agentic AI CPU racks
 Server CPU TAM Breakout by Workload

		2022	2023	2024	2025	2026	2027	2028	2029E	2030E	CAGR '20-'25	CAGR '25-'30	CAGR '10-'20	Notes
CPU Sales (\$mn)	Total	\$23.9	\$19.5	\$22.5	\$35.2	\$59.6	\$89.8	\$110.4	\$138.2	\$170.1	7.3%	37.1%		Bottom-Up Build
	Intel	\$18.1	\$13.2	\$13.6	\$14.4	\$19.8	\$24.7	\$29.6	\$35.1	\$41.2	-9.0%	23.4%		
	AMD	\$5.6	\$5.6	\$7.1	\$9.6	\$16.4	\$22.2	\$28.6	\$35.7	\$43.6	44.0%	35.4%		
	Arm	\$0.2	\$0.7	\$1.8	\$11.2	\$23.4	\$42.9	\$52.2	\$67.5	\$85.4	167.8%	50.0%		
	Traditional/laaS	\$22.9	\$18.0	\$17.4	\$17.0	\$19.0	\$21.2	\$23.7	\$26.4	\$29.5	-6.9%	11.7%	11.9%	
	Intel	\$17.3	\$12.2	\$11.3	\$10.7	\$12.0	\$13.5	\$15.2	\$17.1	\$19.2	-14.0%	12.4%	11.9%	
	AMD	\$5.4	\$5.1	\$5.2	\$5.3	\$5.8	\$6.4	\$7.0	\$7.6	\$8.4	27.8%	9.8%	10.9%	
	Arm	\$0.2	\$0.7	\$0.9	\$1.0	\$1.2	\$1.3	\$1.5	\$1.7	\$1.9	66.4%	13.4%	N/A	
	AI (Compute/Agentic)	\$1.0	\$1.5	\$5.1	\$18.2	\$40.6	\$68.6	\$86.7	\$111.8	\$140.6	123.3%	50.6%		
	Intel	\$0.8	\$1.0	\$2.4	\$3.7	\$7.8	\$11.2	\$14.4	\$18.0	\$22.0	62.4%	43.2%		
	AMD	\$0.2	\$0.5	\$1.9	\$4.3	\$10.6	\$15.8	\$21.6	\$28.0	\$35.2	292.4%	52.1%		
	Arm	\$0.0	\$0.0	\$0.9	\$10.2	\$22.2	\$41.6	\$50.7	\$65.8	\$83.4		52.3%		
	Compute/Head Node	\$1.0	\$1.5	\$5.1	\$18.2	\$34.5	\$44.6	\$52.0	\$61.5	\$70.3	123.3%	31.1%		
	Intel	\$0.8	\$1.0	\$2.4	\$3.7	\$6.6	\$7.3	\$8.6	\$9.9	\$11.0	62.4%	24.7%		
	AMD	\$0.2	\$0.5	\$1.9	\$4.3	\$9.0	\$10.3	\$13.0	\$15.4	\$17.6	292.4%	32.4%		
	Arm	\$0.0	\$0.0	\$0.9	\$10.2	\$13.2	\$27.0	\$30.4	\$36.2	\$41.7		32.5%		
	Agentic AI Node	\$0.0	\$0.0	\$0.0	\$0.0	\$6.1	\$24.0	\$34.7	\$50.3	\$70.3				
	Intel	\$0.0	\$0.0	\$0.0	\$0.0	\$1.2	\$3.9	\$5.8	\$8.1	\$11.0				
	AMD	\$0.0	\$0.0	\$0.0	\$0.0	\$1.6	\$5.5	\$8.6	\$12.6	\$17.6				
	Arm	\$0.0	\$0.0	\$0.0	\$0.0	\$9.0	\$14.6	\$20.3	\$29.6	\$41.7				
CPU Units (mn)	Total	36.6	24.5	24.2	29.2	39.3	49.0	57.1	66.0	74.7	-3.9%	20.7%		Bottom-Up Build
	Intel	30.7	18.5	16.9	17.9	21.1	23.2	24.9	26.7	28.6	-11.6%	9.7%		
	AMD	5.4	4.8	5.5	7.2	10.4	12.5	14.5	16.3	18.1	27.2%	20.2%		
	Arm	0.4	1.2	1.8	4.0	7.9	13.3	17.7	22.9	28.0	77.1%	47.4%		
	Traditional/laaS	35.6	23.2	20.5	20.9	22.4	24.0	25.7	27.5	29.5	-9.9%	7.1%	9.1%	
	Intel	29.9	17.6	14.8	14.8	15.8	16.9	18.1	19.4	20.7	-14.8%	7.0%	9.3%	
	AMD	5.2	4.4	4.2	4.5	4.8	5.2	5.5	5.8	6.2	15.8%	6.6%	5.9%	
	Arm	0.4	1.2	1.5	1.6	1.8	2.0	2.1	2.3	2.5	48.2%	9.0%	N/A	
	AI (Compute/Agentic)	1.0	1.3	3.7	8.3	16.9	25.0	31.4	38.4	45.2	103.5%	40.5%		
	Intel	0.8	0.9	2.2	3.2	5.2	6.3	6.8	7.4	7.8	68.6%	19.9%		
	AMD	0.2	0.4	1.3	2.7	5.5	7.3	9.0	10.5	11.9	262.3%	34.4%		
	Arm	0.0	0.0	0.3	2.4	6.1	11.4	15.5	20.6	25.5	452.6%	60.6%		
	CPU ASP (\$k)	\$0.7	\$0.8	\$0.9	\$1.2	\$1.5	\$1.8	\$1.9	\$2.1	\$2.3	11.7%	13.6%		
	Intel	\$0.6	\$0.7	\$0.8	\$0.8	\$0.9	\$1.1	\$1.2	\$1.3	\$1.4	2.9%	12.5%		
	AMD	\$1.0	\$1.2	\$1.3	\$1.3	\$1.6	\$1.8	\$2.0	\$2.2	\$2.4	13.2%	12.7%		
	Arm	\$0.5	\$0.6	\$1.0	\$2.8	\$3.0	\$3.2	\$3.0	\$3.0	\$3.0	51.2%	1.8%		
	Traditional/laaS	\$0.6	\$0.8	\$0.8	\$0.8	\$0.8	\$0.9	\$0.9	\$1.0	\$1.0	3.3%	4.3%	2.5%	
	Intel	\$0.6	\$0.7	\$0.8	\$0.7	\$0.8	\$0.8	\$0.8	\$0.9	\$0.9	1.0%	5.0%	2.4%	
	AMD	\$1.0	\$1.1	\$1.2	\$1.2	\$1.2	\$1.2	\$1.3	\$1.3	\$1.3	10.3%	3.0%	4.7%	
	Arm	\$0.5	\$0.6	\$0.6	\$0.6	\$0.7	\$0.7	\$0.7	\$0.7	\$0.8	12.3%	4.0%	N/A	
	AI (Compute/Agentic)	\$1.0	\$1.2	\$1.4	\$2.2	\$2.4	\$2.7	\$2.8	\$2.9	\$3.1	9.7%	7.2%		
	Intel	\$1.0	\$1.1	\$1.1	\$1.2	\$1.5	\$1.8	\$2.1	\$2.4	\$2.8	-3.7%	19.4%		
	AMD	\$1.2	\$1.4	\$1.5	\$1.6	\$1.9	\$2.2	\$2.4	\$2.7	\$3.0	8.3%	13.2%		
	Arm			\$3.2	\$4.3	\$3.7	\$3.7	\$3.3	\$3.2	\$3.3		-5.2%		
CPU-XPU Ratio	AI CPU Unit TAM	1.0	1.3	3.7	8.3	16.9	25.0	31.4	38.4	45.2	103.5%	40.5%		Bottom-Up Build AI Accelerator Model
	XPU Unit TAM	4.9	6.2	11.0	16.3	23.6	28.7	28.9	32.8	35.6	39.0%	16.9%		
	AI CPU/XPU Ratio (%)	19%	20%	33%	51%	71%	87%	108%	117%	127%				

Source: BofA Global Research estimates, Mercury Research, IDC

BofA GLOBAL RESEARCH




CPUs now 8%+ of Data Center Systems TAM

At ~\$170bn+ TAM, we now believe server CPUs to make up 8%+ of overall data center systems TAM over time (vs. prior 5-6%), up modestly from a low of <7% in 2025 when the value share of AI training was at its peak.

Over time, we expect the value of CPUs in DC systems to modestly expand given their roles in AI inference and agentic AI workloads, though still at a more limited 7-8% range given other components (such as AI accelerators, HBM memory, networking, connectivity) also remain critical to the overall system.

Exhibit 8: We expect server CPUs to make up 8%+ of total DC TAM value over time, modestly expanding over time as their roles increase in agentic AI workloads

Server CPU as a % of DC Systems TAM




Source: BofA Global Research estimates, Mercury Research, IDC, LightCounting, Gartner

BofA GLOBAL RESEARCH

For AI specifically, we expect AI CPUs to represent 80%+ of total server CPUs over time.

Exhibit 9: Server CPU TAM could reach \$170bn+ by CY30E, with AI CPUs representing 80%+

Server CPU TAM and AI Server CPU as a % of mix



Source: BofA Global Research estimates, Mercury Research, IDC

BofA GLOBAL RESEARCH

AMD: Reiterate Buy, raise PO to \$560 from \$500

We flag increased server CPU market demand visibility, view AMD as the best positioned CPU vendor (unit/ASP expansion), raise estimates accordingly for CY26-28, and also raise our PO to \$560 from \$500 on unchanged 42x CY27 PE.

AMD CPU Advantages

As we mentioned in above sections, we highlight AMD's leading CPU portfolio in both **high frequency** and **core counts**.

High frequency (5GHz+)

For head node (compute node)-type workloads, CPUs with higher single-threaded performance (high-frequency) are key to overall rack performance.

NVDA's Grace (1H23) and Vera (2H26) CPUs were developed specifically for this. While NVDA provides no public disclosure of CPU frequency, third-party benchmarks for Grace point to ~3.35GHz peak perf for the 3-year old CPU. While slower than AMD Turin's (4Q24) peak 5.0GHz and INTC Granite Rapids' (3Q24) peak 4.3GHz, NVDA CPUs also benefit from the tight co-design and co-integration of CPUs and GPUs via fast NVLink C2C connect, as opposed to slower PCIe connectivity for most other CPU vendors.

Meanwhile, AMD's Turin sports up to 5.0GHz peak per-core performance with high-frequency (HF) SKUs, while also utilizing its own InfinityFabric interconnect (for GPU-to-GPU, GPU-to-CPU) in scale-up.

Once the upcoming Venice CPUs launch in 2H26, we expect this single-threaded performance lead to extend, making AMD CPUs ideal for head/compute node in AI scale-up clusters outside of the NVDA ecosystem.

Exhibit 10: AMD's Turin leads in peak frequency at 5.0GHz, with the upcoming Venice expected to lead max core count at 256-core/512-thread

Major CPU Specs, including max # of cores/threads and peak frequency

Vendor	INTC	INTC	AMD	AMD	NVDA	NVDA	ARM	QCOM	AWS	GOOGL	MSFT
Architecture	x86	x86	x86	x86	ARM	ARM	ARM	ARM	ARM	ARM	ARM
Merchant/Custom	Merchant	Merchant	Merchant	Merchant	Merchant	Merchant	Merchant	Merchant	Custom	Custom	Custom
Product	Granite Rapids	Diamond Rapids	Turin	Venice	Grace	Vera	AGI	TBD	Graviton 5	Axion	Cobalt
Process	Intel 3	Intel 18A	TSMC 4nm	TSMC 2nm	TSMC 4nm	TSMC 3nm	TSMC 3nm	TBD	TSMC 3nm	-	-
Launch	3Q24	3Q26	4Q24	3Q26	2Q23	3Q26	1Q27	2027E	2Q26	4Q24	4Q24
Peak Frequency	3.9GHz	TBD	5.0GHz	TBD	3.35GHz	TBD	3.7GHz	TBD	-	3.0GHz?	3.4GHz
Max Cores/Threads	128C/256T	192C	192C/384T	256C/512T	74C	88C/176T	136C	TBD	192C	-	-
Multi-threading	Hyper	N/A	Simultaneous	Simultaneous	N/A	Spatial	N/A	TBD	N/A	-	-

Source: BofA Global Research estimates

BofA GLOBAL RESEARCH

Core counts (up to 256 cores vs. 88-192 competition)

For agentic AI workloads and standalone CPU nodes/racks, CPUs with higher number of cores (and threads via multi-threading) are key to the overall system performance.

This is because within a fixed power envelope of a server/rack, higher core density increases overall aggregate service throughput.

AMD's upcoming 6th gen EPYC Venice (2H26) is expected to top out at 256 cores, while the current 5th gen Turin tops out at 192 cores. This is much higher than NVDA's upcoming Vera (2H26) at 88 cores, as well as INTC's current Granite Rapids topping at 128 cores and upcoming Diamond Rapids (~2H26) likely maxing at 192 cores.

As a result, AMD EPYC CPUs generally lead modeled rack-level results over others, with the EPYC 9965 ("Turin," 192C) delivering a 2.37x normalized geometric mean advantage over NVDA Vera (88 Olympus cores). This lead is expected to extend to 3.30x Vera with the EPYC "Venice" (256C). For context, INTC's Xeon 6980P ("Granite Rapids-AP," 128C) also turns in 1.46x over NVDA Vera in high-core-count agentic workloads.



Upcoming catalyst: Advancing AI 2026 (July)

Separately, we flag AMD's upcoming "Advancing AI 2026" Event in San Francisco, with CEO Dr. Lisa Su's keynote scheduled on July 23.

We anticipate AMD to formally launch its upcoming "Venice" and "Verano" CPUs, the "MI455X" GPU, as well as the "Helios" rack for scale-up compute.

- **CPUs:** We expect a great focus on the performance/core-count lead of AMD CPUs, versus not only INTC but also many other new/existing ARM-based competitors including NVDA, ARM, QCOM who have initially designed CPUs for head/compute-node purposes, not necessarily the agentic AI node/rack purposes (which benefit from having many cores/multi-threading).

We also expect AMD to play up its advantage in the x86 ecosystem with security and RAS features, particularly important in agentic AI workloads.

- **GPUs:** We flag potential for AMD to announce additional GW-scale GPU customers (as full MI455X-based "Helios" customers), on top of the current existing OpenAI and Meta lineup. For reference, AMD has been in engagement with many hyperscale customers to date.

Valuation

Our new PO of \$560 is based on an unchanged 42x CY27E PE, toward middle/upper range of historical 13x-58x, but is well supported by AMD's 50%+ annual EPS CAGR potential and its AI CPU/GPU share gain potential.

Exhibit 11: We raise AMD CY/FY27-28 sales by 6-10%, pf-EPS by 10-16% on increased CPU market outlook and potential new GW-scale GPU engagements

AMD Estimate Changes

	Non-GAAP Sales (\$mn)			EPS (Non-GAAP, ex-SBC)		
	OLD	NEW	delta	OLD	NEW	delta
1Q26	\$13,577	\$13,577	\$0	\$0.29	\$0.29	\$0.00
2Q26E	\$14,302	\$14,302	\$0	\$0.20	\$0.20	\$0.00
3Q26E	\$14,895	\$14,895	\$0	\$0.25	\$0.25	\$0.01
4Q26E	\$15,451	\$15,478	\$27	\$0.30	\$0.31	\$0.01
2026E	\$58,225	\$58,252	\$27	\$1.04	\$1.06	\$0.02
YoY%	10.2%	10.2%	0%	143.6%	148.0%	2%
1Q27E	\$15,582	\$15,664	\$82	\$0.32	\$0.34	\$0.02
2Q27E	\$16,031	\$16,170	\$139	\$0.34	\$0.37	\$0.03
3Q27E	\$16,909	\$17,082	\$173	\$0.43	\$0.47	\$0.04
4Q27E	\$17,718	\$18,018	\$300	\$0.49	\$0.54	\$0.05
2027E	\$66,240	\$66,933	\$694	\$1.58	\$1.72	\$0.14
YoY%	13.8%	14.9%	1%	52.0%	63.0%	9%
1Q28E	\$17,659	\$18,029	\$370	\$0.49	\$0.55	\$0.06
2Q28E	\$18,116	\$18,572	\$456	\$0.52	\$0.59	\$0.07
3Q28E	\$19,085	\$19,638	\$553	\$0.59	\$0.67	\$0.07
4Q28E	\$19,962	\$20,647	\$686	\$0.64	\$0.73	\$0.08
2028E	\$74,822	\$76,887	\$2,065	\$2.25	\$2.53	\$0.28
YoY%	13.0%	14.9%	3%	42.7%	47.0%	12%

Source: BofA Global Research estimates

BofA GLOBAL RESEARCH



INTC: Upgrade to Buy from U/P, PO \$135

We double-upgrade INTC to Buy from Underperform, raise ests/PO to \$135 from \$96.

We highlight INTC's expanded opportunity in server CPUs (similar to AMD described in the AMD section) which it can manufacture itself (gives higher supply visibility), as well as increased opportunities in its external foundry business.

Fully established IDM by CY30, EPS power \$6+

Particularly, we now value INTC as a fully integrated IDM company with EPS power of ~\$6.24 by CY30, with our new PO based on 25x CY30E PE, discounted back two years.

This is a change from our prior CY28 sum-of-parts valuation methodology (more in Exhibit 13) which we now believe under-represents many of the company's CPU and foundry potentials that are further out in CY29-30+.

- **Products:** We now expect INTC server CPUs sales to reach ~\$40bn+ by CY30E (vast majority of DCAI segment is server CPUs), or about ~25% share of \$170bn TAM.
- **Foundry:** While we are not changing our base case foundry segment model until there is further clarity on specific deal details, we present the multiple deals INTC is currently engaged with, size the potential opportunities below. Opportunities span across: Apple M-Series wafers, MediaTek TPU wafers, Terafab IP/packaging, and potentially other ARM-based server CPUs, edge AI expansion into Client – all made more possible via recent CDNS 14A collab announcement.

Exhibit 12: We expect \$6+ INTC EPS power by CY30, now see total company value of \$135, versus \$96 prior based on CY28 sum-of-parts

INTC Sales Power, EPS Power, Valuation Methodology

	CY26E	CY27E	CY28E	CY29E	CY30E	CY30E Methodology	CAGR
Products	\$55.0	\$61.4	\$68.4	\$77.6	\$86.1		11.9%
CCG	\$31.6	\$32.8	\$34.4	\$38.6	\$42.4		
DCAI	\$23.4	\$28.6	\$33.9	\$39.0	\$43.7	~25% Share of \$170bn TAM	
Foundry	\$1.1	\$8.0	\$17.5	\$30.8	\$47.1		158.4%
A/P	\$0.4	\$1.9	\$3.5	\$4.8	\$6.0	30.0% Share of \$20bn 2.5D/3D TAM	
Wafer	\$0.5	\$5.7	\$11.4	\$17.0	\$23.1	~6-7% Share of \$380bn Wafer Foundry TAM	
Apple	\$0.0	\$2.5	\$5.4	\$8.0	\$10.6	40.0% Share of M-Series, 25% A-Series	
TPU/MediaTek	\$0.5	\$3.3	\$6.0	\$9.0	\$12.4	100.0% Share of TPU (MediaTek)	
Terafab	\$0.1	\$0.3	\$1.5	\$4.0	\$8.0	IP, Packaging, Wafer (partial): mid-2028 begin	
Other HPC Products	\$0.0	\$0.0	\$1.0	\$5.0	\$10.0	ARM-based server CPU opp'ty	
Other	\$2.3	\$2.5	\$2.5	\$2.6	\$2.8		
Net Sales	\$58.4	\$71.9	\$88.3	\$111.0	\$135.9		23.5%
Products EBIT	\$18.0	\$21.6	\$25.7	\$30.8	\$33.2	38.5% vs. OpM Target (38.5%)	
Foundry EBIT	-\$9.3	-\$6.6	-\$1.8	\$3.1	\$8.5	18.0% vs. OpM Target (27.5%)	
Other, Unallocated	-\$1.1	-\$1.0	-\$1.0	-\$1.0	-\$1.0		
Total EBIT (non-GAAP)	\$7.6	\$13.9	\$22.8	\$32.9	\$40.6		52.0%
OpM (%)	13.0%	19.4%	25.8%	29.6%	29.9%		
Interest Income	-\$0.2	-\$0.8	-\$0.8	-\$0.8	-\$0.8		
EBT	\$7.4	\$13.2	\$22.0	\$32.1	\$39.8		
Tax Rate	11.5%	12.0%	12.0%	12.0%	12.0%		
Net Income	\$6.5	\$11.6	\$19.4	\$28.3	\$35.1		
NCI Adj.	-\$1.1	-\$1.2	-\$1.2	-\$1.2	-\$1.3		
Net Income (non-GAAP)	\$5.4	\$10.4	\$18.2	\$27.0	\$33.8		
EPS Power	\$1.06	\$2.01	\$3.47	\$5.07	\$6.24		55.9%
Shares Outstanding	5,106	5,172	5,248	5,326	5,406		
Multiple					25x		
Price Objective (25x CY30E discounted back)			\$135	\$145	\$156		
Discount			-7%	-7%			

Source: BofA Global Research estimates



Exhibit 13: We now see ~\$3.47 CY28E EPS Power vs. ~\$2.25 prior, on greater server CPU/Foundry opportunities

INTC Valuation Methodology: New vs. Old

Prior Method	CY28E New	CY28E Old			
	Total	Total	IDM	Foundry	Apple
Products	\$68.4	\$66.3	\$66.3	\$0.0	\$0.0
CCG	\$34.4	\$34.4	\$34.4		
DCAI	\$33.9	\$31.9	\$31.9		
Foundry	\$17.5	\$11.4	\$0.0	\$6.0	\$5.4
A/P	\$3.5	\$3.5		\$3.5	
Wafer	\$11.4	\$7.9	\$0.0	\$2.5	\$5.4
Apple	\$5.4	\$5.4			\$5.4
Other Wafer	\$6.0	\$2.5		\$2.5	
Terafab	\$1.5	\$0.0			
Other HPC Products	\$1.0	\$0.0			
Other	\$2.5	\$2.5	\$2.5		
Net Sales	\$88.3	\$80.2	\$68.8	\$6.0	\$5.4
Products EBIT	\$25.7	\$15.1	\$15.1		
Foundry EBIT	-\$1.8	\$1.5		\$1.5	
Other, Unallocated	-\$1.0	-\$1.0	-\$1.0		
Total EBIT (non-GAAP)	\$22.8	\$15.6	\$14.1	\$1.5	\$0.0
OpM (%)	25.8%	45.4%	20.4%	25.0%	0.0%
Interest Income	-\$0.8	-\$0.8	-\$0.8		
EBT	\$22.0	\$14.8	\$13.3	\$1.5	\$0.0
Tax Rate	12.0%	12.0%	12.0%	12.0%	12.0%
Net Income	\$19.4	\$13.0	\$11.7	\$1.3	\$0.0
NCI Adj.	-\$1.2	-\$1.2	-\$1.2		
Net Income (non-GAAP)	\$18.2	\$11.8	\$10.5	\$1.3	\$0.0
EPS Power	\$3.47	\$2.25	\$2.00	\$0.25	\$0.00
Shares Outst.	5,248	5,246	5,246	5,246	5,246
Methodology		SOP	PE	EV/S	EV/S
Multiple		43x	37x	10x	10x
Price Power	\$135	\$96	\$74	\$11	\$10

Source: BofA Global Research estimates

BofA GLOBAL RESEARCH

Exhibit 14: We raise INTC CY27-28E sales by 1-3%, pf-EPS by 9-12%

INTC Estimate Changes

Non-GAAP Sales (\$mn)				EPS (Non-GAAP, ex-SBC)		
	OLD	NEW	delta	OLD	NEW	delta
2026E	\$58,225	\$58,252	\$27	\$1.04	\$1.06	\$0.02
YoY%	10.2%	10.2%	0%	143.6%	148.0%	2%
1Q27E	\$15,582	\$15,664	\$82	\$0.32	\$0.34	\$0.02
2Q27E	\$16,031	\$16,170	\$139	\$0.34	\$0.37	\$0.03
3Q27E	\$16,909	\$17,082	\$173	\$0.43	\$0.47	\$0.04
4Q27E	\$17,718	\$18,018	\$300	\$0.49	\$0.54	\$0.05
2027E	\$66,240	\$66,933	\$694	\$1.58	\$1.72	\$0.14
YoY%	13.8%	14.9%	1%	52.0%	63.0%	9%
1Q28E	\$17,659	\$18,029	\$370	\$0.49	\$0.55	\$0.06
2Q28E	\$18,116	\$18,572	\$456	\$0.52	\$0.59	\$0.07
3Q28E	\$19,085	\$19,638	\$553	\$0.59	\$0.67	\$0.07
4Q28E	\$19,962	\$20,647	\$686	\$0.64	\$0.73	\$0.08
2028E	\$74,822	\$76,887	\$2,065	\$2.25	\$2.53	\$0.28
YoY%	13.0%	14.9%	3%	42.7%	47.0%	12%

Source: BofA Global Research estimates

BofA GLOBAL RESEARCH



ARM: Reit. Neutral, raise PO to \$335

We see ARM as one of the most prominent beneficiaries of the rising server CPU tide, as the fastest overall ecosystem share gainer through CY30.

As we've highlighted in above sections, we see ARM's collective server CPU value share to reach ~50% by CY30 across: ~35% in merchant CPUs (NVDA, ARM, QCOM, etc.) and ~15% in custom CPUs (hyperscale ASIC such as AWS Graviton, GOOGL Axion, MSFT Cobalt).

However, at ~\$300+, we view the stock as fairly valued, with our new sum-of-parts valuation methodology yielding a PO of \$335, versus our previous total company valuation of \$245 based on 69x CY28E PE.

Sum-of-Parts Valuation: IP & Chip Businesses

We flag the immense opportunities ahead in standalone chip business (vs. historical only IP business), and we now value the company on a sum-of-parts basis.

- **IP business:** For the historical IP business, we continue to value based on ~2.0x PEG framework, in-line with high-growth semis and IP peers.

This is based on +20% Royalty sales CAGR and +12% Licensing sales CAGR, which at 65% combined OpM would yield EPS power of \$5+ by FY31, or +25% EPS CAGR.
- **Chip business:** For the new chiplet business, we value based on ~35x PE, in-line with high-growth AI compute vendor average of ~35x.

We see potential for \$15bn chip sales power by FY31 (~CY30), in-line with ARM mgmt's outlook of 15% share of ~\$100bn CPU TAM. We use 30% OpM target that ARM outlined (in-line with ~30% OpM for more established vendors such as AMD).
- **Total company:** For the total company valuation, we discount back two years and combine – 1) **\$229 worth of IP business**, and 2) **\$106 worth of chip business** – totaling **~\$335 of total company value**.

Exhibit 15: We see total company value of \$335, versus \$245 prior, now based on CY30 sum-of-parts for IP and Chip businesses, discounted back 2 years

ARM Sales Power, EPS Power, Valuation Methodology

	FY26	FY31E - IP Only	FY31E - Chip Only	FY31E - Total Company
IP Sales	\$4,920	\$10,569		\$10,569
Royalty Sales	\$2,613	\$6,503 20% CAGR		
Licensing Sales	\$2,307	\$4,066 12% CAGR		
Chip Sales	\$0		\$15,000	\$15,000
CPU TAM			\$100,000 15% TAM & Share	
Total Sales Power	\$4,920	\$10,569 17% CAGR	\$15,000 25% CAGR	\$25,569 39% CAGR
IP EBIT	\$2,115	\$6,870 65% OpM		\$6,870 65% OpM
Chip EBIT	\$0		\$4,500 30% OpM	\$4,500 30% OpM
Total EBIT	\$2,115	\$6,870	\$4,500	\$11,370
Interest Income	\$107	\$50	\$50	\$100
EBT	\$2,222	\$6,920	\$4,550	\$11,470
Tax Rate	15.0%	15.0%	15.0%	15.0%
Net Income	\$1,889	\$5,882	\$3,868	\$9,749
EPS Power	\$1.77	\$5.35 25% CAGR	\$3.52 15% CAGR	\$8.86 38% CAGR
Shares Outst.	1,068	1,100	1,100	1,100
~CY30 Valuation		\$265 2.0x PEG	\$123 35x PE	\$388 44x Implied PE
~CY29 Valuation		\$246	\$114	\$361 -7% Discount
~CY28 Valuation		\$229	\$106	\$335 -7% Discount

Source: BofA Global Research estimates

BofA GLOBAL RESEARCH



NVDA: Maintain Buy, top AI sector pick

Rising content, increasing opp'ty, solid supply visibility

Following 'GTC Taipei' CEO keynote on Jun-4 (see [our takeaways](#)), we hosted NVDA CFO Colette Kress for a well-attended keynote, and Sr. Director, IR Stewart Stecker for a bus tour investor meeting.

Key takeaways: **1)** Vera CPU opportunity is roughly split 50/50 between head node (compute rack) and agentic AI (standalone CPU rack), **2)** NVDA's \$200bn CPU TAM only includes agentic AI, RL-type workloads, not general purpose traditional CPU workloads, **3)** NVDA content to rise meaningfully every system gen, at ~\$40bn/GW for Blackwell, ~\$60-80bn/GW for Rubin/Rubin Ultra, and potentially ~\$100bn/GW for Feynman, as NVDA consistently adds new chips/racks to the system; **4)** NVDA's full-stack architecture and ecosystem incumbency provides advantages in both hyperscale and ACIE; **5)** GPU useful life may be longer than depreciable life, as 3-4 year old Hoppers are still fully utilized; and **6)** NVDA's \$124bn purchase commitment suggests robust demand/supply visibility ahead for AI infra.


Vera Rubin now in full production, lowest token cost

NVDA announced Vera Rubin is now ramping into full production, across its five-rack system spanning Vera Rubin NVL72 compute, Vera CPU, Groq 3 LPX, Spectrum-6 SPX Ethernet, and Vera BlueField-4 STX storage racks.

Per NVDA, the full system delivers up to 10x higher inference perf per watt and 10x lower cost per token, and up to 35x higher throughput per watt for trillion-parameter models when paired with Groq 3 LPX. In a world of limited power, throughput per watt directly translates to revenue, and NVDA's extreme co-design and system-level solution can provide the lowest token cost by up to x factors. NVDA also noted the Vera Rubin platform is being scaled by hundreds of supply-chain ecosystem partners, with the supply chain 2x as large as Grace Blackwell's.

Exhibit 16: NVDA will begin offering standalone Vera CPU racks, alongside the Vera Rubin Platform beginning in 2H26, more directly addressing previously unaddressed or overlooked workloads in AI inference

Nvidia Vera Rubin SuperPOD with 5 different types of racks making up one giant 40-rack pods



Source: Nvidia

Vera CPU: head node vs. agentic AI likely 50/50 opp'ty

Traditional CPUs are increasingly becoming bottlenecks to agentic GPU utilization, directly affecting token throughput and latency. Vera CPU is designed to address this issue, with its 88x custom NVDA Olympus (Arm) cores all working together in tandem with the greatest cross-sectional bandwidth (3x bandwidth per core, 2x total bandwidth vs. x86), vs. traditional CPUs being rented core by core.

While Vera was initially designed as a head/compute node and falls short in raw core counts vs. competition (at 100-200+ cores per CPU), we flag NVDA's extreme codesign of full end-to-end system can help offset some of the disadvantages in even agentic workloads (where core/thread counts matter). Notably, the Vera CPU supports NVDA's **Spatial Multi-threading (SMT)**, vs. many other x86 products having offered multi-threading of some sort (INTC Hyperthreading, AMD simultaneous multithreading).

NVDA expects ~\$20bn of Vera CPU sales in F2H27, roughly split evenly between head/compute node sockets (traditional GPU support role in NVL racks) and standalone agentic AI sockets (reinforcement learning, orchestration roles in CPU racks).

NVDA has shipped 2.5mn Grace CPUs to date, and the follow-on Vera shipment could already reach ~4-5mn units in just the first two quarters of launch, at an estimated ASP of \$4-5k per CPU. NVDA also noted that the Vera builds on Grace's existing system-level software and optimization on NVDA tech stack, a key differentiation vs. first-gen Grace. Overall, the total CPU unit demand could potentially be higher than GPU over time, reinforcing the CPU-GPU ratio argument of 1-to-1 across both racks (vs. 1-to-2 and 1-to-4 prior).

NVDA content \$/GW rising materially gen-over-gen

NVDA believes every new generation of its systems can increase the company's addressable market materially, from the current ~\$40bn/GW for **Blackwell Ultra**, toward \$60-80bn/GW for **Vera Rubin and Rubin Ultra**, and toward ~\$100bn/GW for **Feynman** maybe. While exact figures may vary, NVDA continues to address more parts of the AI system each gen, with the current Vera Rubin solution comprised of 7 different chips across 5 different racks (vs. Blackwell just 1 main compute rack), and Rubin Ultra (2027) expected to add another rack on top. Beginning the Feynman generation (2028), NVDA also plans to offer optional optical scale-up (via CPO) for customers wanting a larger NVLink domain (vs. regular copper up to 576-die or 144 GPU packages).

"King of diversity": ACIE as important as hyperscale/cloud

NVDA's sales mix between ACIE (AI Clouds, Industrial, Enterprise) and hyperscale is ~50/50 today, with ACIE likely the faster growing segment over time. NVDA is the "king of diversity", as it supplies to every hyperscale, every major model maker, and has significant incumbency and moat in ACIE-type workloads (we view custom XPU's have limited opportunities here). NVDA's continuous upgrades to system software (i.e. Hopper/Blackwell) and new model rollouts are also key competitive advantages.

RTX Spark for personal agentic AI, for top 10% Windows

NVDA and Microsoft unveiled RTX Spark, a new superchip for Windows PCs purpose-built for personal AI agents and expanding NVDA's platform from cloud AI factories into personal AI endpoints. Built on TSMC 3nm, the chip offers 1 petaflop of AI perf., enabling 120bn-parameter LLMs with up to 1mn-token context locally. RTX Spark combines a Blackwell RTX GPU with 6,144 CUDA cores with a 20-core Grace CPU over NVLink-C2C, with MediaTek collaborating on the custom CPU design. Critically, we flag the chip compares to the Apple Mac Mini (10-12 CPU cores) or MacBook Pro (15-18 CPU cores), and likely targets the top 10% of the Windows market at ~25-30mn units/yr

We maintain Buy, \$350 PO, and a top AI pick, given attractive 16x CY27 PE or 0.4x PEG vs. Mag-7 avg of 1.9x.



QCOM: U/P on competitive landscape


We reiterate Underperform on QCOM. In our view, the core handset franchise faces multi-year headwinds as rising memory prices further soften the weak handset market, while Apple modem in-sourcing removes high-margin revenue stream, and China OEM concentration heightens pricing and geopolitical risks.

Against that backdrop, the bull case is increasingly driven by QCOM's diversification into auto/IoT and data center business potentials. While we await further detail at upcoming Jun-24 AI-focused Investor Day, we see compute initiatives years away from being financially material, carry meaningful execution risk, and will require material investment to compete head-on with far better-resourced incumbents INTC, AMD, ARM, and NVDA.

While QCOM may only need to capture a small portion of the ~\$170bn+TAM to lift the stock, we see sustained share gain momentum as difficult amid intensifying competition. QCOM's ability to leverage their existing high-volume agreements with TSMC and Samsung for new data center products may work in their favor, however.

Exhibit 17: QCOM only expected to grow at a 4% CAGR between FY25-FY29E due to handset market softness and Apple roll-off

Summary of QCOM segment revenue forecast




Source: BofA Global Research Estimates

BofA GLOBAL RESEARCH

Exhibit 18: QCOM diversification is underway, though Data Center revenues likely only \$3bn or ~6% by FY29E

Summary of QCOM segment mix forecast



Source: BofA Global Research Estimates

BofA GLOBAL RESEARCH

QCOM's Accelerator Effort: AI200 and AI250

QCOM's upcoming AI200 (2026) and AI250 (2027) rack-scale inference accelerators lean on Hexagon NPU IP and power efficiency expertise from mobile while using lower cost LPDDR memory. QCOM has a ~200MW commitment from HUMAIN beginning late 2026.

We remain constructive on QCOM's engineering capabilities but with inference now the most contested part of the accelerator TAM, a 2026 start provides QCOM with no ecosystem advantage vs. NVDA, AMD, AVGO/GOOGL, and MRVL/AWS while other hyperscalers are already developing high-cost ASICs for their respective clusters.

Data Center CPU re-entrance to leverage Arm-compatible Oryon core

QCOM's current CPU push rests on Arm-compatible Oryon core acquired via Nuvia (2021) which has only shipped at scale in Snapdragon X PC silicon. Porting a laptop-class core to a competitive server SKU poses significant execution challenges while most hyperscalers, neoclouds, and sovereigns already have established server CPU partners.



Role of CPUs in Agentic AI

For much of the AI training era, the compute narrative was dominated by GPUs and custom accelerators. That was appropriate: training large frontier models is highly parallel, matrix-heavy, and accelerator-bound.


In that regime, **CPUs were necessary but largely secondary** — responsible for feeding accelerators, preprocessing data, tokenization, scheduling, storage access, and orchestration.

As a result, AI accelerator spend grew far faster than CPU spend particularly between 2022 and 2025, and accelerators had become the overwhelming majority of data center compute spend, while CPUs represented a much smaller share. The AI accelerator market had grown at +139% CAGR due a tremendous increase in training demand, vs. server CPUs only +14% CAGR.

By 2025, AI accelerators have become 85% of total compute spend and CPUs just 15%, versus accelerators at just 26% (CPUs 74%) in 2021 when the training market had just begun to emerge.

Exhibit 19: AI accelerator market grew at a +139% CAGR between 2022-2025 from a tremendous rise of AI training demand, vs. CPUs at +14% CAGR

CPUs vs. AI Accelerators TAM (\$bn)




Source: BofA Global Research estimates, Mercury Research

BofA GLOBAL RESEARCH

Exhibit 20: AI accelerators as a % of total data center compute spend has risen to 85% by 2025 (CPUs 15%), from just 26% in 2021 (CPUs 74%) before the rise of AI compute demand

CPUs vs. AI Accelerators TAM (%)



Source: BofA Global Research estimates, Mercury Research

BofA GLOBAL RESEARCH

CPUs in training

Overall in training, CPUs are responsible for decompression, tokenization, batching, and ultimately feeding the GPUs. CPUs' main job is to make sure the GPUs/accelerators are utilized as much as possible by constantly preparing and queuing batches, making sure the expensive accelerators are fully in work without bottlenecks.

CPUs control the orchestration and preprocessing of data, i.e. they are typically not involved in the actual heavy processing that AI training requires (mostly done in parallel processing by GPUs/accelerators).

CPUs essentially perform a secondary or supporting role for the GPUs/accelerators in AI training, hence explaining the decline in the % of value in the overall system (vs. GPUs).



CPU in inference

In AI inference, the role of CPUs has become more broad and extensive.

The AI inference process can be broken into three phases: load phase, prefill phase, and decode phase.

1. In **load or initialization**, models are loaded from the disks into the memory, with performance largely dependent on overall disk (I/O) and CPU speed. Weight loading, quantization, and compilation all occur during the load phase, with the latter two (quantization and compilation) much better to be done offline vs. online to avoid runtime overhead.

Overall, the role/speed of CPU is key during the load phase of AI inference.

2. In **prefill**, time to first token (TTFT) and KV cache creation are the two main roles, with both prefill time and first token latency heavily dependent on overall compute/large matrix multiplication abilities, i.e. GPU/accelerators and networking rather than CPUs.


However, the CPUs do provide a role in tokenization, routing, batching, and memory setup.

3. In **decode**, performance is heavily impacted by KV cache reuse, so optimizing KV cache and/or increasing the bandwidth/capacity of system memory is critical, as such with high-bandwidth SRAM or HBM memory, inference storage context memory (i.e. low-latency storage), or even CXL memory (pooling of system memory) methodology.

Here too, increasingly the role of CPUs is increasing in decode because of scheduling, token-by-token control flow, sampling, guardrails, logits processing, memory management, and KV cache routing.

Exhibit 21: LLM inference process can also be broken into three phases: load phase, prefill phase, and decode phase

LLM Inference Process in Three Phases



Source: Intel

BofA GLOBAL RESEARCH

CPUs in Agentic AI

In the next phase of AI – production inference, agentic AI, enterprise AI – is structurally even more CPU-intensive.

Agentic AI does not simply ask a model one question and receive one answer. It plans, retrieves context, calls tools, manages state, invokes APIs, interacts with databases, routes tasks across models, evaluates intermediate outputs, applies guardrails, and loops until a task is completed.

The practical implication is that CPUs become **more than “host processors”**. They become the **control plane** of AI inference.

In a conventional LLM inference request, the GPU may dominate the prefill and decode compute. But in agentic workflows, a single user request can fan out into many CPU-bound subtasks.

Intel’s agentic AI commentary describes these systems as multi-step workflows where CPUs manage orchestration, multiple agents, tool integration, state across workflows, and parallel decision trees. Ciena similarly argues that agentic AI shifts infrastructure beyond a “GPU-only” narrative because orchestration, memory handling, and tool execution increase the role of CPUs and networks in distributed AI workflows.

Exhibit 22: CPUs excel in document processing/classification, indexing, and captioning, among other AI workloads

Role of CPUs in AI Inference

AI Inference Workload	Good fit for...		
	CPUs	CPUs + PCIe-Based GPU	GPU Clusters
Document processing and classification	✓		
Data mining and analytics	✓		✓
Scientific simulations	✓		
Translation	✓		
Indexing	✓		
Content moderation	✓		
Predictive maintenance	✓		✓
Virtual assistants	✓	✓	
Chatbots	✓	✓	
Expert agents	✓	✓	
Video captioning	✓	✓	
Fraud detection		✓	✓
Decision-making		✓	✓
Dynamic pricing		✓	✓
Audio and video filtering		✓	✓
Financial trading			✓
Telecommunications and networking			✓
Autonomous systems			✓

Source: AMD

BofA GLOBAL RESEARCH




This is in-line with NVDA's announcement of its latest Vera CPU-only rack, to launch alongside the Vera Rubin Platform in 2H26.

The NVDA Vera CPU rack integrates up to 256 Vera CPUs in a dense, liquid-cooled rack to provide scalable, energy-efficient capacity. Per NVDA, a single rack can test, execute, and validate results from the Vera Rubin NVL72 (compute) and LPX (ultra-low-latency) racks. Ultimately, these Vera CPU racks provide the foundation for large-scale agentic AI and reinforcement learning.

Exhibit 24: NVDA will begin offering standalone Vera CPU racks, alongside the Vera Rubin Platform beginning in 2H26, more directly addressing previously unaddressed or overlooked workloads in AI inference

Nvidia Vera Rubin SuperPOD with standalone Vera CPU rack




Source: Nvidia

BofA GLOBAL RESEARCH

When stretched out to the entire 40-rack pod cluster, each Vera Rubin POD contains 5 different rack-scale systems (as seen above): 1) 1,152 Rubin GPUs (across 16 racks), 2) 576 Vera CPUs (as part of the GPU compute racks), and 3) potentially another 512 Vera CPUs (as part of standalone CPU racks, assuming 2 racks per pod). This would total 1,088 CPUs to 1,152 GPUs, reaching the roughly **1-to-1 ratio in agentic AI systems**.

Exhibit 25: Each Vera Rubin pod contains 1,152 Rubin GPUs, 512 Vera CPUs (as part of compute), and potentially another 576 Vera CPUs (standalone)

Nvidia Vera Rubin SuperPOD, full pod of 40 racks



Source: Nvidia

BofA GLOBAL RESEARCH

Exhibit 26: Stocks mentioned in this report

Stocks mentioned

BofA Ticker	Bloomberg Ticker	Price	PO	Rating
AMD	AMD US EQUITY	US\$ 452.4	\$560.00	C-1-9
ARM	ARM US EQUITY	US\$ 307.43	\$335.00	C-2-9
INTC	INTC US EQUITY	US\$ 107.04	\$135.00	C-1-9
NVDA	NVDA US EQUITY	US\$ 200.42	\$350.00	C-1-7
QCOM	QCOM US EQUITY	US\$ 191.2	\$165.00	B-3-7

Source: BofA Global Research

BofA GLOBAL RESEARCH

Glossary

- ACIE — AI Clouds, Industrial, Enterprise
- AI — Artificial Intelligence
- AI200 — Qualcomm AI200
- AI250 — Qualcomm AI250
- AMX — Advanced Matrix Extensions
- AMD — Advanced Micro Devices
- API — Application Programming Interface
- A/P — Advanced Packaging
- ARM — Arm / Arm Architecture
- ASIC — Application-Specific Integrated Circuit
- ASP — Average Selling Price
- AWS — Amazon Web Services
- BF16 — Brain Floating Point 16-bit
- CAGR — Compound Annual Growth Rate
- CAPEX — Capital Expenditure
- CCG — Client Computing Group
- CDNS — Cadence Design Systems
- COGS — Cost of Goods Sold
- CPU — Central Processing Unit
- CPO — Co-Packaged Optics
- CUDA — Compute Unified Device Architecture
- CXL — Compute Express Link
- CY — Calendar Year
- DB — Database
- DC — Data Center
- DCAI — Data Center and AI
- DIMM — Dual Inline Memory Module
- DRC — Design Rule Check
- DTCO — Design Technology Co-Optimization
- E — Estimate
- EBIT — Earnings Before Interest and Taxes
- EBT — Earnings Before Taxes
- EDA — Electronic Design Automation



- EMIB — Embedded Multi-die Interconnect Bridge
- EMIB-T — Embedded Multi-die Interconnect Bridge-T
- EM-IR — Electromigration / IR Drop
- EPS — Earnings Per Share
- ETL — Extract, Transform, Load
- EV/S — Enterprise Value to Sales
- FP16 — 16-bit Floating Point
- FY — Fiscal Year
- GDS — Graphic Data System
- GenAI — Generative Artificial Intelligence
- GM — Gross Margin
- GOOGL — Alphabet / Google
- GPU — Graphics Processing Unit
- GTC — GPU Technology Conference
- GW — Gigawatt
- HBM — High Bandwidth Memory
- HF — High Frequency
- HPC — High Performance Computing
- IaaS — Infrastructure as a Service
- IDC — International Data Corporation
- IDM — Integrated Device Manufacturer
- INTC — Intel
- INT8 — 8-bit Integer
- I/O — Input / Output
- IoT — Internet of Things
- IP — Intellectual Property
- KV cache — Key-Value Cache
- LLM — Large Language Model
- LPDDR — Low-Power Double Data Rate
- LPX — NVIDIA LPX Rack Label
- MGX — NVIDIA MGX Platform
- MRDIMM — Multiplexed Rank Dual Inline Memory Module
- MSFT — Microsoft
- MW — Megawatt
- N/A — Not Available / Not Applicable
- NCI — Non-Controlling Interest
- NCI Adj. — Non-Controlling Interest Adjustment
- NoSQL — Non-Structured Query Language
- Non-GAAP — Non-Generally Accepted Accounting Principles
- NPU — Neural Processing Unit
- NVDA — NVIDIA
- NVL — NVLink
- NVL72 — NVLink 72
- OS — Operating System



- OpM — Operating Margin
- PCIe — Peripheral Component Interconnect Express
- PDK — Process Design Kit
- PE — Price to Earnings
- P/E — Price to Earnings
- PEG — Price / Earnings to Growth
- pf-EPS — Pro Forma Earnings Per Share
- PO — Price Objective
- POD — NVIDIA POD / SuperPOD Cluster
- PPA — Performance, Power, Area
- QCOM — Qualcomm
- QCT — Qualcomm CDMA Technologies
- QTL — Qualcomm Technology Licensing
- RAG — Retrieval-Augmented Generation
- RAM — Random Access Memory
- RAS — Reliability, Availability, Serviceability
- RL — Reinforcement Learning
- RTX — NVIDIA RTX
- RTL — Register Transfer Level
- RTL-to-GDS — Register Transfer Level to Graphic Data System
- SAM — Serviceable Available Market
- SBC — Stock-Based Compensation
- SerDes — Serializer / Deserializer
- SLC — System-Level Cache
- Slurm — Simple Linux Utility for Resource Management
- SMT — Simultaneous Multi-threading / Spatial Multi-threading
- SOP — Sum of Parts
- SOTP — Sum of the Parts
- SPX — NVIDIA SPX Rack Label
- SQL — Structured Query Language
- SRAM — Static Random Access Memory
- STX — NVIDIA STX Rack Label
- TAM — Total Addressable Market
- TCO — Total Cost of Ownership
- TPU — Tensor Processing Unit
- TTFT — Time to First Token
- UALink — Universal Accelerator Link
- UCle — Universal Chiplet Interconnect Express
- U/P — Underperform
- XPU — Accelerator Processing Unit / Generic AI Accelerator
- YoY — Year over Year
- 2.5D — Two-and-a-Half-Dimensional Packaging
- 3D — Three-Dimensional Packaging
- 14A — Intel 14 Angstrom Process Node



- 14A-E — Intel 14 Angstrom Enhanced Process Node
- 18A — Intel 18 Angstrom Process Node
- 18A-P — Intel 18 Angstrom Performance Process Node

Price objective basis & risk

Advanced Micro Devices, Inc (AMD)

Our \$560 PO is based on 42x our 2027E non-GAAP EPS. Our PO basis is now towards middle/upper range of historical 13x-58x, but is well supported by AMD's 50%+ annual EPS CAGR potential and its AI CPU/GPU share gain potentials, modestly offset by slower growth in cyclical PC/embedded/console markets.

Downside risks: 1) Execution on first rack-scale product (MI400 Series), 2) Timing/Magnitude of Middle East AI Projects, 3) Lumpy nature of consumer and enterprise spending that could create delays in acceptance and success of new products, 4) High reliance on one outsourced manufacturing partner, 5) Maturity of current game console cycle.

Upside risks are greater share gain potential in the PC and server processor market against competitors.

Arm Holdings (ARM)

Our \$335 PO is based on conceptual sum-of-parts valuation that values the IP business at \$229 (2.0x CY30E PEG, discounted back 2 years) and the chiplet business at \$123 (35x CY30E PE, discounted back 2 years), with both multiples in-line with competitive peer range/average. Our PO implies 95x total company CY28E PE which is still within ARM's historical 35x-147x range, given longer-term data center content and silicon/chiplet (AGI CPU) opportunities, offset by high SoftBank dependence and opex ramp.

Upside risks: 1) better than expected smartphone/PC/server unit shipment, 2) faster adoption of higher-royalty rate v9 and CSS IP at customers, 3) faster share gains and ramp of AGI CPU chiplet lineup, 4) further proliferation of AI data centers and hyperscale-specific custom products, 5) Qualcomm/Nuvia content expansion post-settlement, 6) small floating volume.

Downside risks: 1) historically cyclical nature of semiconductor units, 2) high exposure to mature smartphone market, 3) competition against established x86 and other custom Arm-based CPU offerings in the data center, 4) emerging competition from RISC-V in low-end consumer markets, 5) rising geopolitical tensions and deterioration of Arm China relationship, 6) ongoing Qualcomm/Nuvia litigation, 7) small trading float.

Intel (INTC)

Our \$135 price objective is based on 25x CY30E EPS power of \$6+, discounted back two years, given the company's key server CPU and external foundry wafer/packaging opportunities that extend far out. Our PO implies 9.1x our total company 2028E EV/S which is well above historical 1.7x-4x range but we view as appropriate given increasing server CPU and external foundry opportunity, offset by near-/medium-term manufacturing ramp-up uncertainties.

Upside risks to our price objective are 1) key external foundry packaging/wafer deals that could significantly boost sales/utilization, 2) greater than expected yields/ramps at 18A and upcoming 14A nodes, resulting in greater GM/utilization profile, 3) stronger than expected PC market from Windows 10 refresh or AI uplift, 4) geopolitical tensions boosting sentiment for domestic manufacturing asset.



Downside risks to our price objective are 1) lower than yield/ramp at Intel Foundry, particularly for its new 18A and upcoming 14A nodes, 2) lack of material external foundry customer in wafer processing, 3) weaker-than-expected trends in a mature PC market, which is largest revenue generator for Intel, 4) accelerated share loss to major CPU competitors.

NVIDIA Corporation (NVDA)

Our \$350 PO is based on 26x CY27E PE ex cash, within NVDA's historical 25x-56x forward year PE range, which we believe is justified by NVDA's leading share in fast-growing AI compute/networking markets, offset by lumpiness in global AI projects, cyclical gaming market, and concerns around access to power.

Downside risks are: 1) weakness in consumer driven gaming market, 2) Competition with major public firms, internal cloud projects and other private companies in AI and accelerated computing markets, 3) Larger than expected impact from restrictions on compute shipments to China, or additional restrictions placed on activity in the region, 4) Lumpy and unpredictable sales in new enterprise, data center, and autos markets, 5) Potential for decelerating capital returns, and 6) Enhanced government scrutiny of NVDA's dominant market position in AI chips.

Qualcomm (QCOM)

Our PO of \$165 is based on 15x CY27E pf-EPS (ex-SBC). Our multiple is a discount to diversified/analog peers at 12x-25x and compute/AI peers at 17x-19x on QCOM's more muted growth outlook and relative lack of leverage to data center infrastructure buildout.

Downside risks: Top customer modem roll-off (\$7-8bn by Fall '27), further share loss at other key Android customer (100% to 75%), additional customer insourcing, QTL renewal risk (1H27), China exposure, DRAM tightness from AI server demand, competition, crowded AI data center CPU/accelerator market.

Upside risks: auto/IoT diversification towards \$22bn FY29E target, AI200/AI250 inference accelerators, Nuvia ARM CPUs, HUMAIN 200MW deployment, premium mix shift, rising on-device AI content.

Analyst Certification

I, Vivek Arya, hereby certify that the views expressed in this research report accurately reflect my personal views about the subject securities and issuers. I also certify that no part of my compensation was, is, or will be, directly or indirectly, related to the specific recommendations or view expressed in this research report.

US - Semiconductors and Semiconductor Capital Equipment Coverage Cluster


Investment rating	Company	BofA Ticker	Bloomberg symbol	Analyst
BUY				
	Advanced Energy Industries	AEIS	AEIS US	Duksan Jang
	Advanced Micro Devices, Inc	AMD	AMD US	Vivek Arya
	Allegro Microsystems	ALGM	ALGM US	Vivek Arya
	Analog Devices Inc.	ADI	ADI US	Vivek Arya
	Applied Materials, Inc.	AMAT	AMAT US	Vivek Arya
	Broadcom Inc	AVGO	AVGO US	Vivek Arya
	Cadence	CDNS	CDNS US	Vivek Arya
	Camtek	CAMT	CAMT US	Michael Mani
	Credo Technology	CRDO	CRDO US	Vivek Arya
	Intel	INTC	INTC US	Vivek Arya
	KLA Corporation	KLAC	KLAC US	Vivek Arya
	Lam Research Corp.	LRCX	LRCX US	Vivek Arya
	M/A-Com	MTSI	MTSI US	Vivek Arya
	Marvell Technology, Inc.	MRVL	MRVL US	Vivek Arya
	Microchip	MCHP	MCHP US	Vivek Arya
	Micron Technology, Inc	MU	MU US	Vivek Arya
	MKS Instruments	MKSI	MKSI US	Michael Mani
	Nova	NVMI	NVMI US	Michael Mani
	NVIDIA Corporation	NVDA	NVDA US	Vivek Arya
	onsemi	ON	ON US	Vivek Arya
	Synopsys	SNPS	SNPS US	Vivek Arya
	Teradyne	TER	TER US	Vivek Arya
	Texas Instruments Inc.	TXN	TXN US	Vivek Arya
NEUTRAL				
	Ambarella	AMBA	AMBA US	Vivek Arya
	Ambiq Micro, Inc.	AMBQ	AMBQ US	Vivek Arya
	Arm Holdings	ARM	ARM US	Vivek Arya
	Astera Labs Inc	ALAB	ALAB US	Vivek Arya
	Coherent Corp	COHR	COHR US	Vivek Arya
	Lumentum Holdings	LITE	LITE US	Vivek Arya
	NXP Semiconductors NV	NXPI	NXPI US	Vivek Arya
UNDERPERFORM				
	Axcelis Technologies	ACLS	ACLS US	Duksan Jang
	GlobalFoundries	GFS	GFS US	Vivek Arya
	Lattice Semiconductor	LSCC	LSCC US	Duksan Jang
	Qualcomm	QCOM	QCOM US	Vivek Arya
	Skyworks Solutions, Inc.	SWKS	SWKS US	Vivek Arya
RVW				
	Wolfspeed Inc	WOLF	WOLF US	Vivek Arya



Disclosures


Important Disclosures

Advanced Micro (AMD) Price Chart




The Investment Opinion System is contained at the end of the report under the heading "Fundamental Equity Opinion Key". Dark grey shading indicates the security is restricted with the opinion suspended. Medium grey shading indicates the security is under review with the opinion withdrawn. Light grey shading indicates the security is not covered. Chart is current as of a date no more than one trading day prior to the date of the report.

Arm Holdings (ARM) Price Chart




The Investment Opinion System is contained at the end of the report under the heading "Fundamental Equity Opinion Key". Dark grey shading indicates the security is restricted with the opinion suspended. Medium grey shading indicates the security is under review with the opinion withdrawn. Light grey shading indicates the security is not covered. Chart is current as of a date no more than one trading day prior to the date of the report.

Intel (INTC) Price Chart



The Investment Opinion System is contained at the end of the report under the heading "Fundamental Equity Opinion Key". Dark grey shading indicates the security is restricted with the opinion suspended. Medium grey shading indicates the security is under review with the opinion withdrawn. Light grey shading indicates the security is not covered. Chart is current as of a date no more than one trading day prior to the date of the report.

NVIDIA (NVDA) Price Chart

The Investment Opinion System is contained at the end of the report under the heading "Fundamental Equity Opinion Key". Dark grey shading indicates the security is restricted with the opinion suspended. Medium grey shading indicates the security is under review with the opinion withdrawn. Light grey shading indicates the security is not covered. Chart is current as of a date no more than one trading day prior to the date of the report.

Qualcomm (QCOM) Price Chart

The Investment Opinion System is contained at the end of the report under the heading "Fundamental Equity Opinion Key". Dark grey shading indicates the security is restricted with the opinion suspended. Medium grey shading indicates the security is under review with the opinion withdrawn. Light grey shading indicates the security is not covered. Chart is current as of a date no more than one trading day prior to the date of the report.

Equity Investment Rating Distribution: Technology Group (as of 31 Mar 2026)

Coverage Universe	Count	Percent	Inv. Banking Relationships ^{R1}	Count	Percent
Buy	234	58.50%	Buy	123	52.56%
Hold	90	22.50%	Hold	43	47.78%
Sell	76	19.00%	Sell	23	30.26%

Equity Investment Rating Distribution: Global Group (as of 31 Mar 2026)

Coverage Universe	Count	Percent	Inv. Banking Relationships ^{R1}	Count	Percent
Buy	1993	55.76%	Buy	1186	59.51%
Hold	821	22.97%	Hold	509	62.00%
Sell	760	21.26%	Sell	400	52.63%

^{R1} Issuers that were investment banking clients of BofA Securities or one of its affiliates within the past 12 months. For purposes of this Investment Rating Distribution, the coverage universe includes only stocks. A stock rated Neutral is included as a Hold, and a stock rated Underperform is included as a Sell.



FUNDAMENTAL EQUITY OPINION KEY: Opinions include a Volatility Risk Rating, an Investment Rating and an Income Rating. **VOLATILITY RISK RATINGS**, indicators of potential price fluctuation, are: A - Low, B - Medium and C - High. **INVESTMENT RATINGS** reflect the analyst's assessment of both a stock's absolute total return potential as well as its attractiveness for investment relative to other stocks within its Coverage Cluster (defined below). Our investment ratings are: 1 - Buy stocks are expected to have a total return of at least 10% and are the most attractive stocks in the coverage cluster; 2 - Neutral stocks are expected to remain flat or increase in value and are less attractive than Buy rated stocks and 3 - Underperform stocks are the least attractive stocks in a coverage cluster. An investment rating of 6 (No Rating) indicates that a stock is no longer trading on the basis of fundamentals. Analysts assign investment ratings considering, among other things, the 0-12 month total return expectation for a stock and the firm's guidelines for ratings dispersions (shown in the table below). The current price objective for a stock should be referenced to better understand the total return expectation at any given time. The price objective reflects the analyst's view of the potential price appreciation (depreciation).

Investment rating	Total return expectation (within 12-month period of date of initial rating)	Ratings dispersion guidelines for coverage cluster ^{R2}
Buy	≥ 10%	≤ 70%
Neutral	≥ 0%	≤ 30%
Underperform	N/A	≥ 20%

^{R2}Ratings dispersions may vary from time to time where BofA Global Research believes it better reflects the investment prospects of stocks in a Coverage Cluster.

INCOME RATINGS, indicators of potential cash dividends, are: 7 - same/higher (dividend considered to be secure), 8 - same/lower (dividend not considered to be secure) and 9 - pays no cash dividend. *Coverage Cluster* is comprised of stocks covered by a single analyst or two or more analysts sharing a common industry, sector, region or other classification(s). A stock's coverage cluster is included in the most recent BofA Global Research report referencing the stock.

Price Charts for the securities referenced in this research report are available on the [Price Charts website](#), or call 1-800-MERRILL to have them mailed.

BofAS or one of its affiliates acts as a market maker for the equity securities recommended in the report: Advanced Micro, Arm Holdings, Intel, NVIDIA, Qualcomm.

BofAS or an affiliate was a manager of a public offering of securities of this issuer within the last 12 months: Intel.

The issuer is or was, within the last 12 months, an investment banking client of BofAS and/or one or more of its affiliates: Advanced Micro, Arm Holdings, Intel, Qualcomm.

BofAS or an affiliate has received compensation from the issuer for non-investment banking services or products within the past 12 months: Advanced Micro, Arm Holdings, Intel, NVIDIA, Qualcomm.

The issuer is or was, within the last 12 months, a non-securities business client of BofAS and/or one or more of its affiliates: Advanced Micro, Arm Holdings, Intel, NVIDIA, Qualcomm.

BofAS or an affiliate has received compensation for investment banking services from this issuer within the past 12 months: Advanced Micro, Arm Holdings, Intel, Qualcomm.

BofAS or an affiliate expects to receive or intends to seek compensation for investment banking services from this issuer or an affiliate of the issuer within the next three months: Advanced Micro, Arm Holdings, Intel, Qualcomm.

BofAS together with its affiliates beneficially owns one percent or more of the common stock of this issuer. If this report was issued on or after the 9th day of the month, it reflects the ownership position on the last day of the previous month. Reports issued before the 9th day of a month reflect the ownership position at the end of the second month preceding the date of the report: Advanced Micro, Intel, NVIDIA, Qualcomm.

BofAS or one of its affiliates is willing to sell to, or buy from, clients the common equity of the issuer on a principal basis: Advanced Micro, Arm Holdings, Intel, NVIDIA, Qualcomm.

The issuer is or was, within the last 12 months, a securities business client (non-investment banking) of BofAS and/or one or more of its affiliates: Advanced Micro, Arm Holdings, Intel, NVIDIA, Qualcomm.

BofA Global Research personnel (including the analyst(s) responsible for this report) receive compensation based upon, among other factors, the overall profitability of Bank of America Corporation, including profits derived from investment banking. The analyst(s) responsible for this report may also receive compensation based upon, among other factors, the overall profitability of the Bank's sales and trading businesses relating to the class of securities or financial instruments for which such analyst is responsible.

Other Important Disclosures

From time to time research analysts conduct site visits of covered issuers. BofA Global Research policies prohibit research analysts from accepting payment or reimbursement for travel expenses from the issuer for such visits.

Prices are indicative and for information purposes only. Except as otherwise stated in the report, for any recommendation in relation to an equity security, the price referenced is the publicly traded price of the security as of close of business on the day prior to the date of the report or, if the report is published during intraday trading, the price referenced is indicative of the traded price as of the date and time of the report and in relation to a debt security (including equity preferred and CDS), prices are indicative as of the date and time of the report and are from various sources including BofA Securities trading desks.

The date and time of completion of the production of any recommendation in this report shall be the date and time of dissemination of this report as recorded in the report timestamp.

Recipients who are not institutional investors or market professionals should seek the advice of their independent financial advisor before considering information in this report in connection with any investment decision, or for a necessary explanation of its contents.

Officers of BofAS or one or more of its affiliates (other than research analysts) may have a financial interest in securities of the issuer(s) or in related investments.

Refer to [BofA Global Research policies relating to conflicts of interest](#).

"BofA Securities" includes BofA Securities, Inc. ("BofAS") and its affiliates. Investors should contact their BofA Securities representative or Merrill Global Wealth Management financial advisor if they have questions concerning this report or concerning the appropriateness of any investment idea described herein for such investor. "BofA Securities" is a global brand for BofA Global Research.

Information relating to Non-US affiliates of BofA Securities and Distribution of Affiliate Research Reports:

BofAS and/or Merrill Lynch, Pierce, Fenner & Smith Incorporated ("MLPF&S") may in the future distribute, information of the following non-US affiliates in the US (short name: legal name, regulator): Merrill Lynch (South Africa): Merrill Lynch South Africa (Pty) Ltd., regulated by the Financial Sector Conduct Authority; MLI (UK): Merrill Lynch International, regulated by the Financial Conduct Authority (FCA) and the Prudential Regulation Authority (PRA); BofASE (France): BofA Securities Europe SA is authorized by the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and regulated by the ACPR and the Autorité des Marchés Financiers (AMF). BofA Securities Europe SA ("BofASE") with registered address at 51, rue La Boétie, 75008 Paris is registered under no 842 602 690 RCS Paris. In accordance with the provisions of French Code Monétaire et Financier (Monetary and Financial Code), BofASE is an établissement de crédit et d'investissement (credit and investment institution) that is authorised and supervised by the European Central Bank and the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and regulated by the ACPR and the Autorité des Marchés Financiers. BofASE's share capital can be found at www.bofam.com/BofASEdisclaimer; BofA Europe (Milan): Bank of America Europe Designated Activity Company, Milan Branch, regulated by the Bank of Italy, the European Central Bank (ECB) and the Central Bank of Ireland (CBI); BofA Europe (Frankfurt): Bank of America Europe Designated Activity Company, Frankfurt Branch regulated by BaFin, the ECB and the CBI; BofA Europe (Zurich): Bank of America Europe Designated Activity Company, Zurich Branch, regulated by the Swiss Financial Market Supervisory Authority FINMA, the ECB and CBI; BofA Europe (Madrid): Bank of America Europe Designated Activity Company, Sucursal en España, regulated by the Bank of Spain, the ECB and the CBI; Merrill Lynch (Australia): Merrill Lynch Equities (Australia) Limited, regulated by the Australian Securities and Investments Commission; Merrill Lynch (Hong Kong): Merrill Lynch (Asia Pacific) Limited, regulated by the Hong Kong Securities and Futures Commission (HKSF); Merrill Lynch (Singapore): Merrill Lynch (Singapore) Pte Ltd, regulated by the Monetary Authority of Singapore (MAS); Merrill Lynch (Canada): Merrill Lynch Canada Inc, regulated by the Canadian Investment Regulatory Organization; Merrill Lynch (Mexico): Merrill Lynch Mexico, SA de CV, Casa de Bolsa, regulated by the Comisión Nacional Bancaria y de Valores; BofAS Japan: BofA Securities Japan Co., Ltd., regulated by the Financial Services Agency; Merrill Lynch (Seoul): Merrill Lynch International, LLC Seoul Branch, regulated by the Financial Supervisory Service; Merrill Lynch (Taiwan): Merrill Lynch Securities (Taiwan) Ltd., regulated by the Securities and Futures Bureau; BofAS India: BofA Securities India Limited, regulated by the Securities and Exchange Board of India (SEBI); Merrill Lynch (Israel): Merrill Lynch Israel Limited, regulated by Israel Securities



Authority; Merrill Lynch (DIFC): Merrill Lynch International (DIFC Branch), regulated by the Dubai Financial Services Authority (DFSA); Merrill Lynch (Brazil): Merrill Lynch S.A. Corretora de Títulos e Valores Mobiliários, regulated by Comissão de Valores Mobiliários; Merrill Lynch KSA Company: Merrill Lynch Kingdom of Saudi Arabia Company, regulated by the Capital Market Authority. This information: has been approved for publication and is distributed in the United Kingdom (UK) to professional clients and eligible counterparties (as each is defined in the rules of the FCA and the PRA) by MLI (UK), which is authorized by the PRA and regulated by the FCA and the PRA - details about the extent of our regulation by the FCA and PRA are available from us on request; has been approved for publication and is distributed in the European Economic Area (EEA) by BofASE (France), which is authorized by the ACPR and regulated by the ACPR and the AMF; has been considered and distributed in Japan by BofAS Japan, a registered securities dealer under the Financial Instruments and Exchange Act in Japan, or its permitted affiliates; is issued and distributed in Hong Kong by Merrill Lynch (Hong Kong) which is regulated by HKSF; is issued and distributed in Taiwan by Merrill Lynch (Taiwan); is issued and distributed in India by BofAS India; and is issued and distributed in Singapore to institutional investors and/or accredited investors (each as defined under the Financial Advisers Regulations) by Merrill Lynch (Singapore) (Company Registration No 198602883D). Merrill Lynch (Singapore) is regulated by MAS. Merrill Lynch Equities (Australia) Limited (ABN 65 006 276 795), AFS License 235132 (MLEA) distributes this information in Australia only to 'Wholesale' clients as defined by s.761G of the Corporations Act 2001. With the exception of Bank of America N.A., Australia Branch, neither MLEA nor any of its affiliates involved in preparing this information is an Authorised Deposit-Taking Institution under the Banking Act 1959 nor regulated by the Australian Prudential Regulation Authority. No approval is required for publication or distribution of this information in Brazil and its local distribution is by Merrill Lynch (Brazil) in accordance with applicable regulations. Merrill Lynch (DIFC) is authorized and regulated by the DFSA. Information prepared and issued by Merrill Lynch (DIFC) is done so in accordance with the requirements of the DFSA conduct of business rules. BofA Europe (Frankfurt) distributes this information in Germany and is regulated by BaFin, the ECB and the CBL. BofA Securities entities, including BofA Europe and BofASE (France), may outsource/delegate the marketing and/or provision of certain research services or aspects of research services to other branches or members of the BofA Securities group. You may be contacted by a different BofA Securities entity acting for and on behalf of your service provider where permitted by applicable law. This does not change your service provider. Please refer to the [Electronic Communications Disclaimers](#) for further information.

This information has been prepared and issued by BofAS and/or one or more of its non-US affiliates. The author(s) of this information may not be licensed to carry on regulated activities in your jurisdiction and, if not licensed, do not hold themselves out as being able to do so. BofAS and/or MLPF&S is the distributor of this information in the US and accepts full responsibility for information distributed to BofAS and/or MLPF&S clients in the US by its non-US affiliates. Any US person receiving this information and wishing to effect any transaction in any security discussed herein should do so through BofAS and/or MLPF&S and not such foreign affiliates. Hong Kong recipients of this information should contact Merrill Lynch (Asia Pacific) Limited in respect of any matters relating to dealing in securities or provision of specific advice on securities or any other matters arising from, or in connection with, this information. Singapore recipients of this information should contact Merrill Lynch (Singapore) Pte Ltd in respect of any matters arising from, or in connection with, this information. For clients that are not accredited investors, expert investors or institutional investors Merrill Lynch (Singapore) Pte Ltd accepts full responsibility for the contents of this information distributed to such clients in Singapore.

General Investment Related Disclosures:

Taiwan Readers: Neither the information nor any opinion expressed herein constitutes an offer or a solicitation of an offer to transact in any securities or other financial instrument. No part of this report may be used or reproduced or quoted in any manner whatsoever in Taiwan by the press or any other person without the express written consent of BofA Securities. This document provides general information only, and has been prepared for, and is intended for general distribution to, BofA Securities clients. Neither the information nor any opinion expressed constitutes an offer or an invitation to make an offer, to buy or sell any securities or other financial instrument or any derivative related to such securities or instruments (e.g., options, futures, warrants, and contracts for differences). This document is not intended to provide personal investment advice and it does not take into account the specific investment objectives, financial situation and the particular needs of, and is not directed to, any specific person(s). This document and its content do not constitute, and should not be considered to constitute, investment advice for purposes of ERISA, the US tax code, the Investment Advisers Act or otherwise. Investors should seek financial advice regarding the appropriateness of investing in financial instruments and implementing investment strategies discussed or recommended in this document and should understand that statements regarding future prospects may not be realized. Any decision to purchase or subscribe for securities in any offering must be based solely on existing public information on such security or the information in the prospectus or other offering document issued in connection with such offering, and not on this document.

Securities and other financial instruments referred to herein, or recommended, offered or sold by BofA Securities, are not insured by the Federal Deposit Insurance Corporation and are not deposits or other obligations of any insured depository institution (including, Bank of America, N.A.). Investments in general and, derivatives, in particular, involve numerous risks, including, among others, market risk, counterparty default risk and liquidity risk. No security, financial instrument or derivative is suitable for all investors. Digital assets are extremely speculative, volatile and are largely unregulated. In some cases, securities and other financial instruments may be difficult to value or sell and reliable information about the value or risks related to the security or financial instrument may be difficult to obtain. Investors should note that income from such securities and other financial instruments, if any, may fluctuate and that price or value of such securities and instruments may rise or fall and, in some cases, investors may lose their entire principal investment. Past performance is not necessarily a guide to future performance. Levels and basis for taxation may change.

This report may contain a short-term trading idea or recommendation, which highlights a specific near-term catalyst or event impacting the issuer or the market that is anticipated to have a short-term price impact on the equity securities of the issuer. Short-term trading ideas and recommendations are different from and do not affect a stock's fundamental equity rating, which reflects both a longer term total return expectation and attractiveness for investment relative to other stocks within its Coverage Cluster. Short-term trading ideas and recommendations may be more or less positive than a stock's fundamental equity rating.

BofA Securities is aware that the implementation of the ideas expressed in this report may depend upon an investor's ability to "short" securities or other financial instruments and that such action may be limited by regulations prohibiting or restricting "shortselling" in many jurisdictions. Investors are urged to seek advice regarding the applicability of such regulations prior to executing any short idea contained in this report.

Foreign currency rates of exchange may adversely affect the value, price or income of any security or financial instrument mentioned herein. Investors in such securities and instruments, including ADRs, effectively assume currency risk.

BofAS or one of its affiliates is a regular issuer of traded financial instruments linked to securities that may have been recommended in this report. BofAS or one of its affiliates may, at any time, hold a trading position (long or short) in the securities and financial instruments discussed in this report.

BofA Securities, through business units other than BofA Global Research, may have issued and may in the future issue trading ideas or recommendations that are inconsistent with, and reach different conclusions from, the information presented herein. Such ideas or recommendations may reflect different time frames, assumptions, views and analytical methods of the persons who prepared them, and BofA Securities is under no obligation to ensure that such other trading ideas or recommendations are brought to the attention of any recipient of this information.

In the event that the recipient received this information pursuant to a contract between the recipient and BofAS for the provision of research services for a separate fee, and in connection therewith BofAS may be deemed to be acting as an investment adviser, such status relates, if at all, solely to the person with whom BofAS has contracted directly and does not extend beyond the delivery of this report (unless otherwise agreed specifically in writing by BofAS). If such recipient uses the services of BofAS in connection with the sale or purchase of a security referred to herein, BofAS may act as principal for its own account or as agent for another person. BofAS is and continues to act solely as a broker-dealer in connection with the execution of any transactions, including transactions in any securities referred to herein.

Copyright and General Information:

Copyright 2026 Bank of America Corporation. All rights reserved. iQDatabase® is a registered service mark of Bank of America Corporation. This information is prepared for the use of BofA Securities clients and may not be redistributed, retransmitted or disclosed, in whole or in part, or in any form or manner, without the express written consent of BofA Securities. This document and its content is provided solely for informational purposes and cannot be used for training or developing artificial intelligence (AI) models or as an input in any AI application (collectively, an AI tool). Any attempt to utilize this document or any of its content in connection with an AI tool without explicit written permission from BofA Global Research is strictly prohibited. BofA Global Research utilizes AI, including machine learning and other technologies, to enhance the services we provide to our clients. These technologies assist our analysts in various aspects of their work, including but not limited to data analysis, content extraction, content creation, data aggregation and summarization and identifying relevant information from diverse sources. All AI-driven processes are subject to review by BofA Global Research employees. BofA Global Research information is distributed simultaneously to internal and client websites and other portals by BofA Securities and is not publicly-available material. Any unauthorized use or disclosure is prohibited. Receipt and review of this information constitutes your agreement not to redistribute, retransmit, or disclose to others the contents, opinions, conclusion, or information contained herein (including any investment recommendations, estimates or price targets) without first obtaining express permission from an authorized officer of BofA Securities.

Materials prepared by BofA Global Research personnel are based on public information. Facts and views presented in this material have not been reviewed by, and may not reflect information known to, professionals in other business areas of BofA Securities, including investment banking personnel. BofA Securities has established information barriers between BofA Global Research and certain business groups. As a result, BofA Securities does not disclose certain client relationships with, or compensation received from, such issuers. To the extent this material discusses any legal proceeding or issues, it has not been prepared as nor is it intended to express any legal conclusion, opinion or advice. Investors should consult their own legal advisers as to issues of



law relating to the subject matter of this material. BofA Global Research personnel's knowledge of legal proceedings in which any BofA Securities entity and/or its directors, officers and employees may be plaintiffs, defendants, co-defendants or co-plaintiffs with or involving issuers mentioned in this material is based on public information. Facts and views presented in this material that relate to any such proceedings have not been reviewed by, discussed with, and may not reflect information known to, professionals in other business areas of BofA Securities in connection with the legal proceedings or matters relevant to such proceedings.

This information has been prepared independently of any issuer of securities mentioned herein and not in connection with any proposed offering of securities or as agent of any issuer of any securities. None of BofAS any of its affiliates or their research analysts has any authority whatsoever to make any representation or warranty on behalf of the issuer(s). BofA Global Research policy prohibits research personnel from disclosing a recommendation, investment rating, or investment thesis for review by an issuer prior to the publication of a research report containing such rating, recommendation or investment thesis.

Any information relating to sustainability in this material is limited as discussed herein and is not intended to provide a comprehensive view on any sustainability claim with respect to any issuer or security.

Any information relating to the tax status of financial instruments discussed herein is not intended to provide tax advice or to be used by anyone to provide tax advice. Investors are urged to seek tax advice based on their particular circumstances from an independent tax professional.

The information herein (other than disclosure information relating to BofA Securities and its affiliates) was obtained from various sources and we do not guarantee its accuracy. This information may contain links to third-party websites. BofA Securities is not responsible for the content of any third-party website or any linked content contained in a third-party website. Content contained on such third-party websites is not part of this information and is not incorporated by reference. The inclusion of a link does not imply any endorsement by or any affiliation with BofA Securities. Access to any third-party website is at your own risk, and you should always review the terms and privacy policies at third-party websites before submitting any personal information to them. BofA Securities is not responsible for such terms and privacy policies and expressly disclaims any liability for them.

All opinions, projections and estimates constitute the judgment of the author as of the date of publication and are subject to change without notice. Prices also are subject to change without notice. BofA Securities is under no obligation to update this information and BofA Securities ability to publish information on the subject issuer(s) in the future is subject to applicable quiet periods. You should therefore assume that BofA Securities will not update any fact, circumstance or opinion contained herein.

Subject to the quiet period applicable under laws of the various jurisdictions in which we distribute research reports and other legal and BofA Securities policy-related restrictions on the publication of research reports, fundamental equity reports are produced on a regular basis as necessary to keep the investment recommendation current.

Certain outstanding reports or investment opinions relating to securities, financial instruments and/or issuers may no longer be current. Always refer to the most recent research report relating to an issuer prior to making an investment decision.

In some cases, an issuer may be classified as Restricted or may be Under Review or Extended Review. In each case, investors should consider any investment opinion relating to such issuer (or its security and/or financial instruments) to be suspended or withdrawn and should not rely on the analyses and investment opinion(s) pertaining to such issuer (or its securities and/or financial instruments) nor should the analyses or opinion(s) be considered a solicitation of any kind. Sales persons and financial advisors affiliated with BofAS or any of its affiliates may not solicit purchases of securities or financial instruments that are Restricted or Under Review and may only solicit securities under Extended Review in accordance with firm policies.

Neither BofA Securities nor any officer or employee of BofA Securities accepts any liability whatsoever for any direct, indirect or consequential damages or losses arising from any use of this information.