BofA SECURITIES

US Semiconductors

Accelerating the Frontier: Agentic race expands AI TAM to $2.2Tn

Industry Overview

Agentic AI race accelerates TAM to $2.2Tn by CY30E

We raise our CY26-30E AI data center systems addressable market to $2.2Tn at a +40% annual pace from $1.8Tn/+33% prior. The virtuous flywheel of agentic tool adoption by consumers and enterprises, competition between public and private AI labs, and limited chip supply will likely keep AI spending robust for the next several years. Any pacing or addition of safety guardrails to frontier models is likely to increase not reduce compute requirements and shift more demand to faster growing inference from training. Meanwhile, valuations for top chip stocks remain compelling relative to their projected EPS growth, with SOX trading at a 21x forward PE, 12% below its 24x median multiple since the onset of ChatGPT and well below 40%+ projected EPS growth potential.

Top picks into Q4: NVDA, INTC, MRVL, MU, LRCX

Historically, CQ4 (and CQ1) have been the two best seasonal quarters to own chip stocks, with 300-500bps of median outperformance vs. SPX from 2010-25. Our top picks with near-term catalysts include: 1) NVDA (enhanced buybacks, multiple GTC tradeshow events), 2) INTC (agentic CPU strength, foundry win anticipation), 3) MU (start of buybacks from Dec-9), 4) MRVL (Analyst Day on Oct 6, XPU ramp) and 5) LRCX (share gain potential across memory, logic). We also continue to like AMD, AMAT, ADI and ON. Staying watchful on AVGO where we are Buy-rated on compelling valuation with potential for enhanced buybacks when the company reports FQ4 in December.

Key takeaways/enhancements to our data center model

Our data center model(s) presents a comprehensive overview across hyperscaler capex, GW deployment plans, and silicon allocation across compute, memory, networking, storage and optical chips. Key takeaways include: **1)** CY30E AI DC TAM at $2.2Tn, +40% CAGR from CY26, driven by AI accelerator TAM of ~$1.5Tn and AI networking TAM of ~$400bn; **2)** DC memory sales expected to reach ~$1.3Tn by CY30E vs. total memory TAM of ~$2.0Tn; **3)** Networking and optics to grow at 30-40%+ CAGR driven by increasing scale-up/out domain and adoption of CPO, NPO, AEC, etc.; **4)** CPU TAM of $210bn, with AI CPUs representing ~$180bn. We also now present a "bottom-up" (chip vendor perspective) and "top-down" perspective of the **35-45 GW+** of annual compute deployments spread across major public, neoclouds and sovereign deployments.

Cloud capex could reach $2-3Tn by 2030E, at 35-45GW/yr

Our latest cloud capex tracker points to top US hyperscale/neocloud and China cloud aggregate capex of ~$1.0Tn in 2026E (~2x YoY) and ~$1.4Tn in 2027E (+37% YoY). By the end of the decade, we expect spending to reach ~$2-3Tn as suggested by NVDA CEO Jensen Huang. This implies annual data center deployment of 35-50 GW through the end of the decade (at ~$50-60bn/GW), in line with our bottom-up/top-down ~35-45 GW+ assumption (see GW Installment Analysis on pg. 4-5). We flag that near-term capex is increasingly backed by multi-year customer cloud commitments and backlog of ~$2.3Tn at just the top four hyperscalers, while aggregate debt financing of ~$370bn remains modest vs. annual FCF potential that greatly exceeds it.

29 September 2026

Equity
United States
Semiconductors

Vivek Arya
Research Analyst
BofAS
+1 646 855 1755
vivek.arya@bofa.com

Duksan Jang
Research Analyst
BofAS
+1 646 556 4825
duksan.jang@bofa.com

Michael Mani
Research Analyst
BofAS
+1 646 855 2232
michael.mani@bofa.com

Liam Pharr
Research Analyst
BofAS
+1 646 855 3146
liam.pharr@bofa.com

See inside for Glossary

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 24 to 27. Analyst Certification on page 23. Price Objective Basis/Risk on page 20. 13022625

Timestamp: 29 September 2026 10:36PM EDT

Contents

AI 2030 TAM Outlook ~$2.2Tn 3

Gigawatt (GW) Installment Analysis 4

Bottom Up GW Analysis by GPU XPU Vendor 4

Top Down GW Analysis by DC Operator 5

$210bn+ Server CPU TAM 6

Market share outlook analysis 7

Traditional vs. AI-ready Node Composition vs. Agentic Racks 8

CPUs ~7.8% of Data Center Systems TAM 10

Cloud/Carrier Capex Overview 11

Capex as a % of cost of revs, shift to an Intangible IP 13

Capex sustainability: demand signals increasingly positive 14

Capex is increasingly backed by commitments and backlog 15

Global Semis Forecast Update 16

Glossary: 19

AI 2030E TAM Outlook ~$2.2Tn

Of the overall global IT spend of ~$6.4Tn today, we believe **data center systems** to represent $839bn across both AI and non-AI. By CY30E, we expect TAM to grow toward ~$2.6Tn, growing at +32% CAGR CY26-30E and outpacing overall IT spend's +11% CAGR.

AI Data Center Systems: ~$2.2Tn by CY30E

For **AI data center systems** specifically, we now see TAM growing to ~$1.8Tn by CY30E from $564bn in CY26E, with AI servers representing ~75-80% of TAM at $1.4Tn, followed by networking at ~15-20% of TAM at $310bn, and storage at ~5% of TAM at $85bn.

- Within AI servers, we expect **AI accelerators** to represent ~$1.5Tn TAM by CY30E from just $199bn in CY25 and $382bn in CY26E, or the vast majority of server spend.
- We forecast overall **server CPU** to grow to $210bn in CY30E from ~$61bn in CY26E (+36% CAGR), while AI CPUs specifically could grow to ~$180bn in CY30E from $42bn in CY26E (+44% CAGR).
- **AI connectivity** is expected to grow at a +41% CAGR to $163bn in CY30E from $41bn in CY26E, driven by **copper connectivity** growing at +52% to $28bn and **optical connectivity** growing at +33% to $111bn.

For **non-AI data center systems**, we see TAM growth of a modest +7% CAGR.

Exhibit 1: We forecast AI Data Center Systems TAM to reach ~$2.2Tn by CY30E, +40% CAGR, with AI accelerators representing ~$1.5Tn within that
Data Center Systems TAM Breakout – AI vs. non-AI

TAM ($bn)	2022	2023	2024	2025	2026E	2027E	2028E	2029E	2030E	CAGR '26-'30
Overall IT Spend (1)	$4,594.2	$4,693.0	$5,037.6	$5,580.0	$6,391.1	$7,351.9	$8,173.4	$8,974.6	$9,761.2	11%
Data Center Systems TAM ($bn) (2)	$227.1	$237.5	$333.5	$505.5	$839.4	$1,327.3	$1,814.7	$2,231.1	$2,561.7	32%
YoY (%)		5%	40%	52%	66%	58%	37%	23%	15%	
Memory	$57.2	$30.2	$70.3	$110.0	$494.3	$766.0	$885.0	$1,090.6	$1,308.7	28%
DRAM (incl. HBM)	$37.5	$19.2	$45.5	$82.3	$362.1	$579.5	$681.7	$838.5	$1,006.3	29%
NAND	$19.7	$11.0	$24.8	$27.7	$132.3	$186.5	$203.2	$252.0	$302.4	23%
Memory % of DC (2)	*25%*	*13%*	*21%*	*22%*	*59%*	*58%*	*49%*	*49%*	*51%*	
AI Data Center Systems TAM ($bn)	$22.8	$61.7	$160.9	$278.3	$569.2	$1,026.1	$1,495.6	$1,889.6	$2,206.0	40%
YoY (%)		171%	161%	73%	105%	80%	46%	26%	17%	
AI % of Overall IT Spend (1)	0.5%	1.3%	3.2%	5.0%	8.9%	14.0%	18.3%	21.1%	22.6%	
AI % of Data Center Systems TAM (2)	10.0%	26.0%	48.3%	55.0%	67.8%	77.3%	82.4%	84.7%	86.1%	
AI Servers	$16.1	$48.9	$131.6	$228.2	$445.2	$805.5	$1,172.3	$1,470.2	$1,706.6	40%
AI CPUs (headnode + agentic) (excl-trad)	$1.0	$1.5	$5.1	$18.5	$42.4	$83.2	$114.4	$144.0	$180.4	44%
AI Accelerators	$14.3	$45.0	$120.2	$198.8	$382.4	$686.9	$1,008.5	$1,266.9	$1,460.5	40%
HBM	$1.6	$4.2	$17.4	$34.5	$77.4	$152.8	$198.9	$229.8	$276.9	38%
HBM (% of Accelerators)	*11%*	*9%*	*14%*	*17%*	*20%*	*22%*	*20%*	*18%*	*19%*	
Other (DDR/SSD/motherboard/power/etc.)	$0.8	$2.3	$6.3	$10.9	$20.4	$35.4	$49.4	$59.3	$65.6	34%
AI Networking	$5.6	$9.6	$21.2	$34.6	$95.6	$171.7	$252.1	$329.4	$394.3	43%
AI Switching (Back/Front-End) (incl. OCS)	$2.2	$3.8	$8.4	$12.5	$31.0	$60.1	$86.7	$128.3	$159.6	51%
AI SmartNIC	$0.9	$2.0	$3.3	$4.8	$23.7	$45.3	$71.9	$84.9	$96.2	42%
AI Connectivity/Other Networking	$2.5	$3.8	$9.5	$17.3	$41.2	$67.9	$98.9	$132.0	$163.3	41%
Optical	$1.9	$3.4	$8.6	$15.2	$35.8	$56.0	$75.2	$93.8	$111.1	33%
CPO	*$0.0*	*$0.0*	*$0.0*	*$0.0*	*$0.2*	*$0.9*	*$3.2*	*$9.4*	*$14.1*	*206%*
NPO	*$0.0*	*$0.0*	*$0.0*	*$0.0*	*$0.1*	*$0.6*	*$2.2*	*$6.4*	*$10.7*	*204%*
Electrical/Copper	$0.6	$0.4	$0.9	$2.1	$5.2	$10.3	$18.3	$22.4	$27.5	52%
DAC	*$0.5*	*$0.4*	*$0.6*	*$0.8*	*$1.3*	*$2.0*	*$2.6*	*$2.9*	*$3.2*	*25%*
ACC	*$0.0*	*$0.0*	*$0.0*	*$0.0*	*$0.4*	*$1.4*	*$2.5*	*$3.8*	*$7.4*	*110%*
AEC	*$0.1*	*$0.1*	*$0.2*	*$1.2*	*$3.5*	*$6.9*	*$13.2*	*$15.8*	*$16.9*	*49%*
AI Storage	$1.1	$3.2	$8.2	$15.5	$28.4	$48.9	$71.2	$90.0	$105.0	39%
Non-AI Data Center TAM ($bn)	$204.3	$175.8	$172.6	$227.3	$270.2	$301.2	$319.1	$341.6	$355.7	7%
YoY (%)		-14%	-2%	32%	19%	11%	6%	7%	4%	

Source: BofA Global Research estimates, Gartner, Mercury Research, IDC, LightCounting, 650 Group

BofA GLOBAL RESEARCH

Exhibit 2: We expect NVDA to generally maintain ~70%+ AI accelerator share over time
AI accelerator (GPU/ASIC/XPU) potential sales power by vendor

AI Accelerator Sales Power	2024	2025	2026E	2027E	2028E	2029E	2030E
Nvidia ($bn)	**$102.2**	**$162.4**	**$298.9**	**$515.0**	**$721.1**	**$894.2**	**$1,001.5**
YoY		59%	84%	72%	40%	24%	12%
Implied Share	85%	82%	78%	75%	72%	71%	69%
AMD ($bn)	**$5.0**	**$6.6**	**$14.5**	**$40.6**	**$64.2**	**$84.3**	**$102.3**
YoY		34%	118%	180%	58%	31%	21%
Implied Share	4%	3%	4%	6%	6%	7%	7%
Broadcom ($bn)	**$9.3**	**$16.2**	**$49.8**	**$102.6**	**$187.3**	**$224.8**	**$251.8**
YoY		75%	208%	106%	83%	20%	12%
Implied Share	8%	8%	13%	15%	19%	18%	17%
Marvell ($bn)	**$0.7**	**$1.3**	**$1.7**	**$3.0**	**$6.8**	**$11.3**	**$15.1**
YoY		87%	25%	82%	128%	65%	34%
Implied Share	1%	1%	0%	0%	1%	1%	1%
Others/TBD ($bn)	**$3.1**	**$12.3**	**$17.6**	**$25.7**	**$29.0**	**$52.4**	**$89.9**
YoY		296%	43%	46%	13%	81%	72%
Implied Share	3%	6%	5%	4%	3%	4%	6%
Total Market ($bn)	**$120.2**	**$198.8**	**$382.4**	**$686.9**	**$1,008.5**	**$1,266.9**	**$1,460.5**
YoY		65%	92%	80%	47%	26%	15%

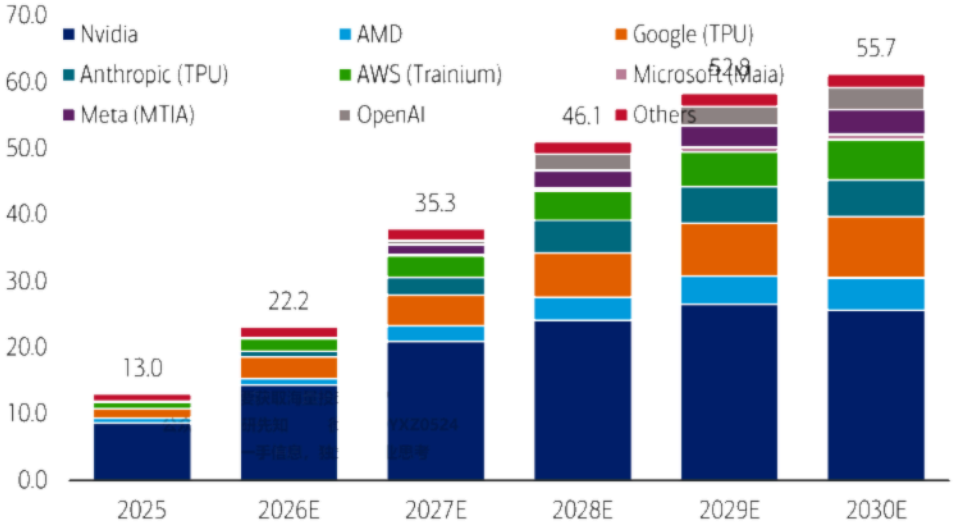
Source: BofA Global Research estimates

Gigawatt (GW) Installment Analysis

Bottom-up GW analysis by GPU/XPU vendor

Off of ~41GW of installed global data center capacity through 2025, we expect the industry to add ~22GW incremental capacity in 2026E, ~35GW in 2027E, and ultimately ~212GW between 2026-2030E, based on individual chip vendor bottom-up forecasts.

Exhibit 3: Bottom-up GW analysis: ~212GW incremental capacity from 2026 to 2030
Estimated annual GW data center deployment based on bottom-up XPU vendor sales



Source: BofA Global Research estimates, company data

Top-down GW analysis by DC operator

Our operator-by-operator analysis suggests global AI/cloud data center capacity could increase from ~41GW in CY25 to ~219GW by CY30E, with annual additions accelerating from ~22GW in CY26E to ~45GW by CY30E. Hyperscalers remain the largest contributors, while neocloud and sovereign deployments become increasingly material over time.

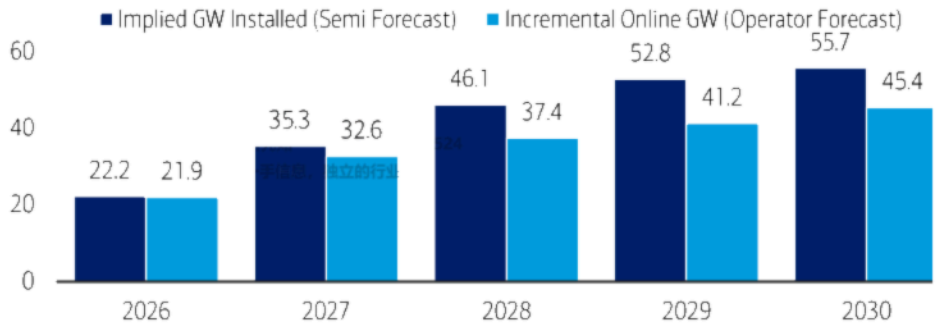
Exhibit 4: Top-down GW analysis: ~41GW 2025 capacity to ~5x to ~219GW by 2030E
Estimated cumulative GW online per operator across hyperscalers, neoclouds, and sovereigns

GW	2025	2026	2027	2028	2029	2030
Hyperscalers						
Alphabet	5.5	8.9	14.7	20.7	26.6	32.4
Amazon	16.2	22.6	31.3	40.5	49.5	58.1
Meta	5.0	7.6	11.2	15.0	18.8	22.8
Microsoft	10.3	13.7	19.1	24.7	29.8	34.4
Oracle	1.6	4.1	6.6	9.3	11.6	13.5
Total Hyperscaler GW	**38.6**	**56.9**	**83.0**	**110.2**	**136.3**	**161.2**
Incremental Hyperscaler GW		*18.3*	*26.1*	*27.2*	*26.1*	*24.9*
Neoclouds						
CoreWeave	0.9	1.9	3.0	4.2	6.1	8.0
Nebius	0.2	0.7	1.5	2.4	3.7	5.0
SPCX/xAI (NVDA classifies as hyperscaler)	0.7	1.5	3.0	5.4	12.4	24.7
IREN	0.0	0.3	0.8	1.2	1.7	2.3
Lambda	0.1	0.1	0.2	0.4	0.7	1.0
Fluidstack	0.1	0.4	1.0	1.1	1.3	1.4
Other Neoclouds	0.3	0.6	1.7	3.9	5.3	6.7
Total Neocloud GW	**2.2**	**5.5**	**11.2**	**18.7**	**31.1**	**49.1**
Incremental Neocloud GW		*3.3*	*5.7*	*7.5*	*12.5*	*17.9*
Sovereign						
HUMAIN (Saudi)	0.0	0.1	0.4	0.8	1.3	1.8
G42 (UAE)	0.1	0.3	0.5	0.8	1.1	1.3
SK Telecom (South Korea)	-	-	-	1.5	2.5	3.0
Other	0.0	0.1	0.4	0.9	1.8	3.1
Total Sovereign GW	**0.1**	**0.5**	**1.3**	**3.9**	**6.6**	**9.2**
Incremental Sovereign GW		*0.4*	*0.8*	*2.7*	*2.7*	*2.6*
Total GW	**40.9**	**62.8**	**95.4**	**132.8**	**174.0**	**219.4**
Incremental GW		**21.9**	**32.6**	**37.4**	**41.2**	**45.4**

Source: BofA Global Research (Semiconductors, Software, Internet, Aerospace & Defense), Company Commentary, Med a Reports

Our ~219 GW operator estimate is modestly below the ~253GW (41GW + 212GW) implied by our semiconductor bottom-up model, a difference of ~15%. We attribute the gap primarily to the semi model capturing broader industry deployments not yet allocated to named operators, such as many on-prem/sovereign projects.

Exhibit 5: Operator forecasts lower than implied GW installations from semiconductor forecast
Implied GW installed from semiconductor bottom-up forecast vs. top-down GW analysis



Source: BofA Global Research (Semiconductors, Software, Internet, Aerospace & Defense), Company Commentary, Med a Reports

$210bn+ 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 ~$210bn by CY30E vs. ~$61bn in CY26E (+36% CAGR).

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

- For **ARM merchant CPUs**, we 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 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 6: We see total server CPU TAM reaching $210bn+ by CY30E from $61bn in CY26E, growing at +36% CAGR
Server CPU TAM (by revenue, units, ASP)

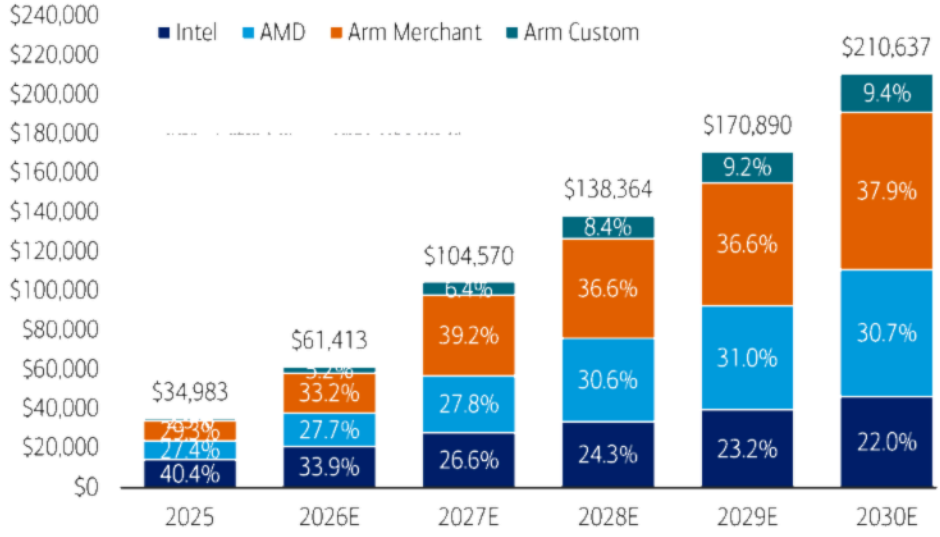
Server CPU TAM	2023	2024	2025	2026E	2027E	2028E	2029E	2030E	CY26-30 CAGR
Revenue ($bn)									
Total	**$19.5**	**$22.5**	**$35.0**	**$61.4**	**$104.6**	**$138.4**	**$170.9**	**$210.6**	**36.1%**
INTC	**$13.2**	**$13.6**	**$14.1**	**$20.8**	**$27.8**	**$33.6**	**$39.6**	**$46.3**	**22.2%**
AMD	**$5.6**	**$7.1**	**$9.6**	**$17.0**	**$29.1**	**$42.4**	**$52.9**	**$64.6**	**39.6%**
ARM	**$0.7**	**$1.8**	**$11.3**	**$23.6**	**$47.7**	**$62.4**	**$78.4**	**$99.7**	**43.4%**
ARM Merchant	$0.0	$0.7	$10.3	$20.4	$41.0	$50.7	$62.6	$79.9	40.7%
NVDA		*$0.7*	*$10.3*	*$20.4*	*$39.7*	*$47.4*	*$54.2*	*$59.7*	*30.8%*
ARM					*$0.9*	*$2.2*	*$6.0*	*$15.0*	
QCOM					*$0.4*	*$1.1*	*$2.4*	*$5.3*	
ARM Custom (Graviton, Axion, Cobalt, etc.)	$0.7	$1.0	$1.0	$3.2	$6.7	$11.7	$15.8	$19.8	57.6%
Units (mn)									
Total	**24.5**	**24.2**	**29.9**	**38.0**	**51.8**	**63.4**	**72.2**	**82.0**	**21.2%**
INTC	**18.5**	**16.9**	**18.5**	**20.2**	**23.4**	**25.6**	**27.5**	**29.3**	**9.7%**
AMD	**4.8**	**5.5**	**7.2**	**10.4**	**15.3**	**19.8**	**22.4**	**24.9**	**24.3%**
ARM	**1.2**	**1.8**	**4.1**	**7.4**	**13.2**	**18.0**	**22.4**	**27.8**	**39.3%**
ARM Merchant	0.0	0.2	2.6	4.8	8.5	10.7	13.5	17.7	38.6%
NVDA	*0.0*	*0.2*	*2.6*	*4.8*	*7.9*	*9.5*	*10.6*	*11.5*	*24.3%*
ARM	*0.0*	*0.0*	*0.0*	*0.0*	*0.4*	*0.8*	*2.0*	*4.6*	
QCOM	*0.0*	*0.0*	*0.0*	*0.0*	*0.2*	*0.4*	*0.8*	*1.6*	
ARM Custom (Graviton, Axion, Cobalt, etc.)	1.2	1.6	1.5	2.6	4.7	7.3	8.9	10.1	40.6%
ASP ($)									
Total	**$797**	**$931**	**$1,170**	**$1,615**	**$2,020**	**$2,182**	**$2,367**	**$2,568**	**12.3%**
INTC	**$716**	**$804**	**$761**	**$1,028**	**$1,190**	**$1,311**	**$1,442**	**$1,579**	**11.3%**
AMD	**$1,164**	**$1,299**	**$1,322**	**$1,635**	**$1,907**	**$2,140**	**$2,364**	**$2,601**	**12.3%**
ARM	**$589**	**$998**	**$2,748**	**$3,190**	**$3,624**	**$3,472**	**$3,504**	**$3,582**	**2.9%**
ARM Merchant		$4,000	$4,000	$4,246	$4,835	$4,735	$4,647	$4,515	1.5%
NVDA		*$4,000*	*$4,000*	*$4,246*	*$5,000*	*$5,000*	*$5,100*	*$5,202*	*5.2%*
ARM					*$2,409*	*$2,681*	*$2,955*	*$3,251*	
QCOM					*$2,409*	*$2,681*	*$2,955*	*$3,251*	
ARM Custom (Graviton, Axion, Cobalt, etc.)	$589	$649	$666	$1,234	$1,427	$1,609	$1,773	$1,951	12.1%
Value Share (%)									
Total	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**	
INTC	**68%**	**61%**	**40%**	**34%**	**27%**	**24%**	**23%**	**22%**	
AMD	**29%**	**32%**	**27%**	**28%**	**28%**	**31%**	**31%**	**31%**	
ARM	**4%**	**8%**	**32%**	**38%**	**46%**	**45%**	**46%**	**47%**	
ARM Merchant	0%	3%	29%	33%	39%	37%	37%	38%	
NVDA		*3%*	*29%*	*33%*	*38%*	*34%*	*32%*	*28%*	
ARM					*1%*	*2%*	*4%*	*7%*	
QCOM					*0%*	*1%*	*1%*	*2%*	
ARM Custom (Graviton, Axion, Cobalt, etc.)	4%	5%	3%	5%	6%	8%	9%	9%	

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 7: We expect INTC to represent ~22% value share, AMD ~31%, ARM merchant ~38%, ARM custom ~9% by CY30E
Server CPU TAM ($mn) and Value Market Share Outlook



Source: BofA Global Research estimates, Mercury Research, IDC

- **INTC (~22% value share):** We see share declining to just ~22% of total server CPU value by CY30E from ~34% in CY26E, though we expect INTC's continued relative strength in traditional enterprise workloads.

  However, we see continued $ revenue growth of +20% CAGR between CY26-30E, particularly as INTC navigates through the industry-wide supply shortages, benefits from ASP expansion. By CY28E, 18A-P-based Coral Rapids should ramp and could begin to narrow the performance gap vs. competition.

- **AMD (~31% value share):** We see share of ~27-31% through CY30E, given AMD's continued performance lead at the high-end (high-frequency, high-core-count AI head node and agentic node), modestly offset by ramps of new 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.

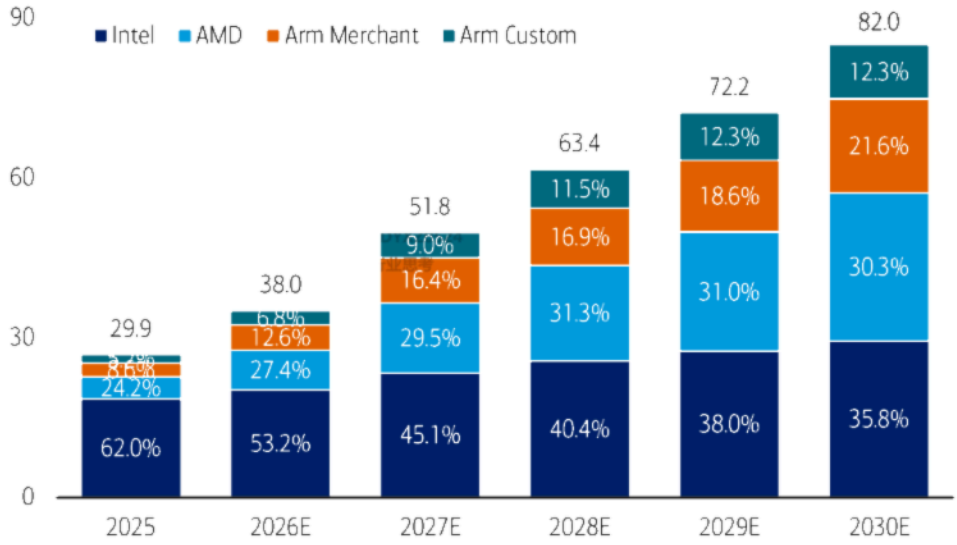
  Similar to INTC, we expect AMD to continue benefiting from the x86 ecosystem's advantages in security and RAS features, important in agentic AI workloads.

- **ARM (merchant; ~38% value share):** We see share rising to ~38% by CY30E from ~28% in CY26E, led by a continued ramp of NVDA CPUs (in both compute/agentic nodes), as well as new ARM/QCOM merchant solutions ramping in CY27-28E.

- **ARM (custom; ~9% value share):** We see share gain toward ~9% by CY30E 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 ~36% share by CY30E.

Exhibit 8: We expect INTC to represent ~36% unit share, AMD ~30%, ARM merchant ~22%, and ARM custom ~12% by CY30E
Server CPU TAM (mn units) and Unit Market Share Outlook



Source: BofA Global Research estimates, Mercury Research, IDC

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

In the following exhibit, we also present the $210bn CY30E server CPU TAM in three 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 (~$180bn) by CY30E, as NVDA recently highlighted.

- **Traditional/IaaS:** 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 four years, we expect this growth rate to continue, at ~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 9: Of the $210bn CY30E CPU TAM, we see ~$30bn for traditional/IaaS, ~$90bn for AI compute/head node, ~$90bn for agentic AI CPU racks
Server CPU TAM Breakout by Workload

		2022	2023	2024	2025	2026	2027	2028	2029E	2030E	CAGR '20-'25	CAGR '26-'30	CAGR '10-'20	Notes
CPU Sales ($mn)	**Total**	**$23.9**	**$19.5**	**$22.5**	**$35.0**	**$61.4**	**$104.6**	**$138.4**	**$170.9**	**$210.6**	**7.2%**	**36.1%**		**Bottom-Up Build**
	Intel	$18.1	$13.2	$13.6	$14.1	$20.8	$27.8	$33.6	$39.6	$46.3	-9.3%	22.2%		
	AMD	$5.6	$5.6	$7.1	$9.6	$17.0	$29.1	$42.4	$52.9	$64.6	44.0%	39.6%		
	Arm	$0.2	$0.7	$1.8	$11.3	$23.6	$47.7	$62.4	$78.4	$99.7	168.1%	43.4%		
	Traditional/IaaS	**$22.9**	**$18.0**	**$17.4**	**$16.5**	**$19.0**	**$21.4**	**$24.0**	**$26.9**	**$30.2**	**-7.5%**	**12.2%**	**11.9%**	
	Intel	$17.3	$12.2	$11.3	$10.2	$12.0	$13.5	$15.1	$17.0	$19.1	-14.8%	12.4%	11.9%	
	AMD	$5.4	$5.1	$5.2	$5.3	$5.9	$6.6	$7.3	$8.2	$9.2	27.8%	11.8%	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**	**$18.5**	**$42.4**	**$83.2**	**$114.4**	**$144.0**	**$180.4**	**124.1%**	**43.6%**		
	Intel	$0.8	$1.0	$2.4	$3.9	$8.8	$14.3	$18.4	$22.6	$27.2	64.9%	32.6%		
	AMD	$0.2	$0.5	$1.9	$4.3	$11.2	$22.5	$35.1	$44.7	$55.5	292.4%	49.3%		
	Arm	$0.0	$0.0	$0.9	$10.2	$22.4	$46.4	$60.9	$76.6	$97.7		44.5%		
	Compute/Head Node	**$1.0**	**$1.5**	**$5.1**	**$18.5**	**$36.0**	**$54.1**	**$68.6**	**$79.2**	**$90.2**	**124.1%**	**25.8%**		
	Intel	$0.8	$1.0	$2.4	$3.9	$7.5	$9.3	$11.1	$12.4	$13.6	64.9%	16.1%		
	AMD	$0.2	$0.5	$1.9	$4.3	$9.5	$14.6	$21.0	$24.6	$27.7	292.4%	30.8%		
	Arm	$0.0	$0.0	$0.9	$10.2	$13.3	$30.1	$36.5	$42.2	$48.9		38.4%		
	Agentic AI Node	**$0.0**	**$0.0**	**$0.0**	**$0.0**	**$6.4**	**$29.1**	**$45.8**	**$64.8**	**$90.2**		**94.1%**		
	Intel	$0.0	$0.0	$0.0	$0.0	$1.3	$5.0	$7.4	$10.2	$13.6		79.1%		
	AMD	$0.0	$0.0	$0.0	$0.0	$1.7	$7.9	$14.0	$20.1	$27.7		101.7%		
	Arm	$0.0	$0.0	$0.0	$0.0	$9.1	$16.2	$24.4	$34.5	$48.9		52.3%		
CPU Units (mn)	**Total**	**36.6**	**24.5**	**24.2**	**29.9**	**38.0**	**51.8**	**63.4**	**72.2**	**82.0**	**-3.4%**	**21.2%**		**Bottom-Up Build**
	Intel	30.7	18.5	16.9	18.5	20.2	23.4	25.6	27.5	29.3	-11.0%	9.7%		
	AMD	5.4	4.8	5.5	7.2	10.4	15.3	19.8	22.4	24.9	27.2%	24.3%		
	Arm	0.4	1.2	1.8	4.1	7.4	13.2	18.0	22.4	27.8	77.8%	39.3%		
	Traditional/IaaS	**35.6**	**23.2**	**20.5**	**20.2**	**21.7**	**23.3**	**25.0**	**26.9**	**28.9**	**-10.6%**	**7.4%**	**9.1%**	**~7-9% CAGR**
	Intel	29.9	17.6	14.8	14.0	15.0	16.1	17.2	18.4	19.7	-15.7%	7.0%	9.3%	
	AMD	5.2	4.4	4.2	4.5	4.9	5.3	5.7	6.2	6.6	15.8%	8.0%	5.9%	
	Arm	0.4	1.2	1.5	1.6	1.8	2.0	2.1	2.3	2.5	48.2%	9.0%	N/A	
	AI (Compute/Agentic)	**1.0**	**1.3**	**3.7**	**9.7**	**16.3**	**28.5**	**38.4**	**45.3**	**53.1**	**110.0%**	**34.3%**		
	Intel	0.8	0.9	2.2	4.5	5.2	7.3	8.4	9.1	9.6	81.0%	16.6%		
	AMD	0.2	0.4	1.3	2.7	5.5	10.0	14.1	16.2	18.2	262.3%	34.7%		
	Arm	0.0	0.0	0.3	2.5	5.6	11.2	15.8	20.0	25.3	456.3%	45.8%		
CPU ASP ($k)	**Total**	**$0.7**	**$0.8**	**$0.9**	**$1.2**	**$1.6**	**$2.0**	**$2.2**	**$2.4**	**$2.6**	**11.0%**	**12.3%**		**Bottom-Up Build**
	Intel	$0.6	$0.7	$0.8	$0.8	$1.0	$1.2	$1.3	$1.4	$1.6	1.8%	11.3%		
	AMD	$1.0	$1.2	$1.3	$1.3	$1.6	$1.9	$2.1	$2.4	$2.6	13.2%	12.3%		
	Arm	$0.5	$0.6	$1.0	$2.7	$3.2	$3.6	$3.5	$3.5	$3.6	50.7%	2.9%		
	Traditional/IaaS	**$0.6**	**$0.8**	**$0.8**	**$0.8**	**$0.9**	**$0.9**	**$1.0**	**$1.0**	**$1.0**	**3.4%**	**4.5%**	**2.5%**	**~3-5% CAGR**
	Intel	$0.6	$0.7	$0.8	$0.7	$0.8	$0.8	$0.9	$0.9	$1.0	1.0%	5.0%	2.4%	
	AMD	$1.0	$1.1	$1.2	$1.2	$1.2	$1.2	$1.3	$1.3	$1.4	10.3%	3.5%	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**	**$1.9**	**$2.6**	**$2.9**	**$3.0**	**$3.2**	**$3.4**	**6.7%**	**7.0%**		
	Intel	$1.0	$1.1	$1.1	$0.9	$1.7	$2.0	$2.2	$2.5	$2.8	-8.9%	13.7%		Double-digit CAGR
	AMD	$1.2	$1.4	$1.5	$1.6	$2.0	$2.3	$2.5	$2.8	$3.0	8.3%	10.8%		Double-digit CAGR
	Arm			$3.2	$4.2	$4.0	$4.1	$3.8	$3.8	$3.9		-0.9%		NVDA share declines
CPU-XPU Ratio	**AI CPU Unit TAM**	**1.0**	**1.3**	**3.7**	**9.7**	**16.3**	**28.5**	**38.4**	**45.3**	**53.1**	**110.0%**	**34.3%**		**Bottom-Up Build**
	XPU Unit TAM	**4.9**	**6.2**	**11.0**	**16.3**	**23.6**	**28.7**	**29.1**	**32.9**	**35.7**	**39.0%**	**10.9%**		**AI Accelerator Model**
	AI CPU/XPU Ratio (%)	**19%**	**20%**	**33%**	**59%**	**69%**	**99%**	**132%**	**138%**	**149%**				

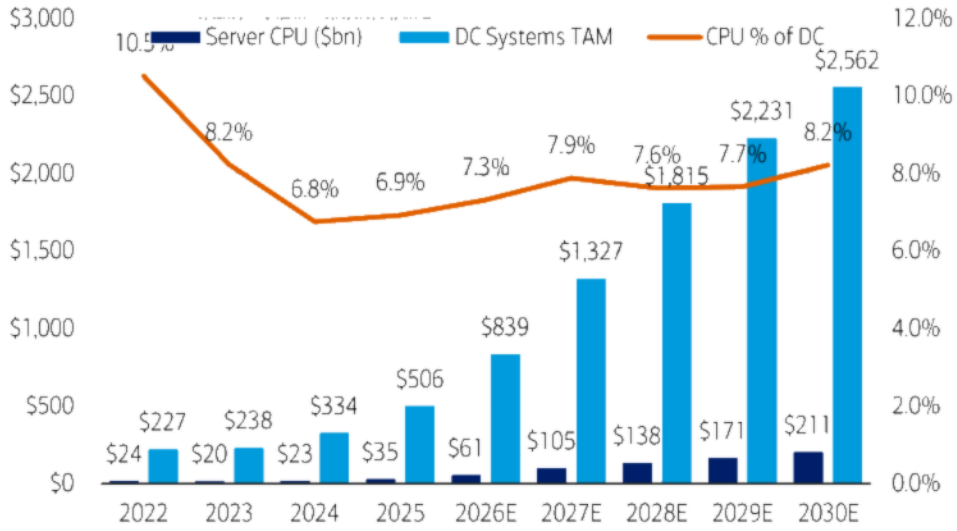
Source: BofA Global Research estimates, Mercury Research, IDC

CPUs ~7-8% of Data Center Systems TAM

At ~$210bn+ TAM, we expect server CPUs to make up 7-8% of overall data center systems TAM over time vs. <7% in 2024-25 when the value share of AI training was at its peak, still a modest and conservative outlook, in our view.

Over time, we expect the value of CPUs in DC systems to modestly expand given their roles in AI inference and agentic AI workloads, generally growing faster than other components (such as AI accelerators, HBM memory, networking, connectivity) which also remain critical to the overall system.

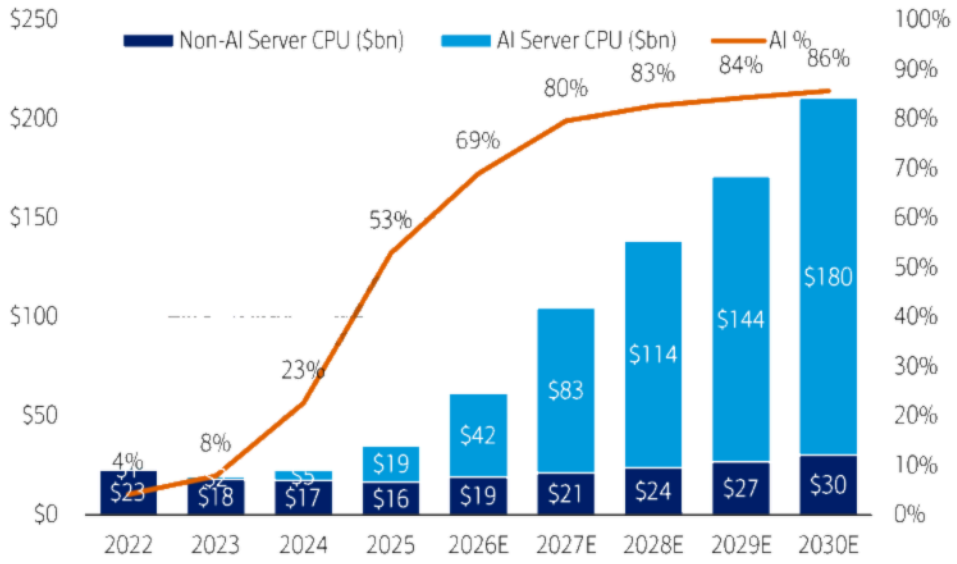
Exhibit 10: We expect server CPUs to make up ~7-8% of total DC TAM value over time, modestly expanding over time as their roles increase in agentic AI workloads
Server CPU as a % of DC Systems TAM



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

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

Exhibit 11: Server CPU TAM could reach $210bn+ by CY30E, with AI CPUs representing ~86%
Server CPU TAM and AI Server CPU as a % of mix



Source: BofA Global Research estimates, Mercury Research, IDC

Cloud/Carrier Capex Overview

Exhibit 12: Consensus/BofA aggregate projections suggest capex across major cloud and neocloud vendors can grow +98%/+37% YoY in CY26/27E
Summary of Cloud Capex Outlook (incl. leases)

CapEx ($mn)	C1Q26	C2Q26	C3Q26E	C4Q26E	C1Q27E	C2Q27E	C3Q27E	C4Q27E	2024	2025	2026E	2027E	2028E
Top US	**169,119**	**217,554**	**276,646**	**293,102**	**293,231**	**320,013**	**341,343**	**366,097**	**273,785**	**478,627**	**956,421**	**1,320,684**	**1,502,621**
Google	35,674	44,924	55,645	65,311	68,031	74,824	81,447	88,507	52,535	91,447	201,554	312,809	360,697
Microsoft (incl. leases)	31,900	41,000	51,568	52,703	54,975	58,526	60,729	61,622	75,600	118,000	177,170	235,852	266,007
Amazon (incl. leases)	44,203	54,208	58,546	64,826	62,967	70,540	73,398	80,905	82,999	131,819	221,783	287,810	305,195
Meta (incl. leases)	19,840	30,116	40,852	49,674	45,727	51,032	56,100	62,551	39,227	70,920	140,483	215,411	232,087
Oracle	18,635	16,493	28,499	20,241	21,342	22,783	25,501	25,130	10,745	35,477	83,868	94,756	104,943
Neoclouds	**18,867**	**30,813**	**41,536**	**40,348**	**40,188**	**42,308**	**44,169**	**47,382**	**12,679**	**30,964**	**131,563**	**174,047**	**233,692**
xAI (SpaceX AI)	6,503	15,803	15,803	16,791	16,750	17,750	18,250	19,250	NA	12,609	54,900	72,000	116,593
CoreWeave (ex-CIP)	6,800	5,600	12,501	8,308	9,879	10,718	11,135	12,110	8,702	9,700	33,209	43,842	44,750
Nebius	2,473	5,657	6,798	9,360	7,699	7,701	8,304	8,989	808	4,066	24,289	32,693	43,074
Other Neoclouds	3,091	3,752	6,433	5,889	5,860	6,139	6,480	7,033	3,170	4,589	19,166	25,512	29,275
Top China	**11,810**	**11,851**	**14,107**	**15,960**	**15,169**	**14,240**	**17,250**	**19,602**	**19,348**	**31,405**	**53,727**	**66,261**	**82,350**
Alibaba	4,203	6,216	6,901	6,027	5,164	7,242	7,961	7,178	9,456	16,836	23,347	27,545	34,797
Tencent	6,938	4,676	6,216	8,749	9,120	5,772	8,031	10,996	8,600	12,851	26,579	33,920	41,263
Baidu	669	959	990	1,184	885	1,225	1,258	1,427	1,292	1,718	3,801	4,796	6,290
Total	**180,929**	**229,404**	**290,753**	**309,062**	**308,401**	**334,253**	**358,593**	**385,699**	**293,133**	**510,032**	**1,010,149**	**1,386,946**	**1,584,971**
NVDA/AMD/AVGO DC	91,821	112,620	131,730	152,631	176,069	201,849	231,421	264,449	141,801	234,896	488,802	873,787	1,235,041
YoY %													
Top US	**91.6%**	**97.4%**	**113.1%**	**95.0%**	**73.4%**	**47.1%**	**23.4%**	**24.9%**	**69.8%**	**74.8%**	**99.8%**	**38.1%**	**13.8%**
Google	107.4%	100.1%	132.3%	134.5%	90.7%	66.6%	46.4%	35.5%	62.9%	74.1%	120.4%	55.2%	15.3%
Microsoft (incl. leases)	49.1%	69.4%	47.8%	40.5%	72.3%	42.7%	17.8%	16.9%	83.5%	56.1%	50.1%	33.1%	12.8%
Amazon (incl. leases)	76.7%	68.4%	66.8%	64.0%	42.5%	30.1%	25.4%	24.8%	57.4%	58.8%	68.2%	29.8%	6.0%
Meta (incl. leases)	53.3%	77.0%	117.0%	124.4%	130.5%	69.5%	37.3%	25.9%	39.4%	80.8%	98.1%	53.3%	7.7%
Oracle	217.9%	81.6%	235.2%	68.2%	14.5%	38.1%	(10.5%)	24.2%	54.9%	230.2%	136.4%	13.0%	10.8%
Neoclouds	**223.9%**	**481.8%**	**384.8%**	**257.9%**	**113.0%**	**37.3%**	**6.3%**	**17.4%**	**NA**	**144.2%**	**324.9%**	**32.3%**	**34.3%**
xAI (SpaceX AI)	153.3%	2,009.9%	232.0%	270.4%	157.6%	12.3%	15.5%	14.6%	NA	NA	335.4%	31.1%	61.9%
CoreWeave (ex-CIP)	257.9%	93.1%	558.0%	176.9%	45.3%	91.4%	(10.9%)	45.8%	NA	11.5%	242.4%	32.0%	2.1%
Nebius	354.7%	1,008.0%	611.5%	355.3%	211.4%	36.1%	22.1%	(4.0%)	NA	403.5%	497.4%	34.6%	31.8%
Other Neoclouds	279.4%	230.1%	575.9%	249.4%	89.6%	63.6%	0.7%	19.4%	NA	44.8%	317.7%	33.1%	14.8%
Top China	**83.3%**	**57.6%**	**64.0%**	**80.6%**	**28.4%**	**20.2%**	**22.3%**	**22.8%**	**144.6%**	**62.3%**	**71.1%**	**23.3%**	**24.3%**
Alibaba	40.4%	33.2%	34.6%	49.0%	22.9%	16.5%	15.4%	19.1%	214.6%	78.0%	38.7%	18.0%	26.3%
Tencent	119.0%	90.9%	105.8%	107.7%	31.4%	23.4%	29.2%	25.7%	157.2%	49.4%	106.8%	27.6%	21.6%
Baidu	137.6%	137.2%	118.3%	104.3%	32.4%	27.8%	27.1%	20.6%	(17.2%)	33.0%	121.2%	26.2%	31.1%
Total	**91.1%**	**94.8%**	**110.0%**	**94.2%**	**70.5%**	**45.7%**	**23.3%**	**24.8%**	**73.3%**	**74.0%**	**98.1%**	**37.3%**	**14.3%**
NVDA/AMD/AVGO DC	94.6%	127.3%	112.3%	100.6%	91.8%	79.2%	75.7%	73.3%	139.1%	65.7%	108.1%	78.8%	41.3%
Cap. intensity (%)													
Top US	**37.8%**	**44.4%**	**56.9%**	**52.5%**	**56.5%**	**56.6%**	**58.6%**	**55.2%**	**18.6%**	**28.4%**	**48.2%**	**56.7%**	**52.3%**
Google	32.5%	37.5%	50.1%	52.9%	56.2%	57.0%	58.5%	58.0%	15.0%	22.7%	43.4%	57.5%	48.0%
Microsoft (incl. leases)	38.5%	45.6%	56.9%	55.1%	56.1%	55.0%	56.5%	54.3%	28.9%	38.6%	49.3%	55.5%	51.9%
Amazon (incl. leases)	24.4%	27.0%	29.0%	26.5%	30.2%	30.8%	31.6%	29.1%	13.0%	18.4%	26.8%	30.4%	28.1%
Meta (incl. leases)	35.2%	49.5%	64.6%	67.6%	67.5%	69.9%	73.7%	70.6%	23.8%	35.3%	55.3%	70.5%	63.7%
Oracle	108.4%	86.0%	147.3%	95.0%	91.4%	86.3%	92.4%	81.4%	19.6%	58.1%	108.9%	87.5%	66.4%
Neoclouds	**235.9%**	**256.3%**	**221.1%**	**143.2%**	**122.4%**	**112.2%**	**104.0%**	**98.0%**	**467.8%**	**409.9%**	**196.4%**	**107.9%**	**83.5%**
xAI (SpaceX AI)	138.5%	202.2%	122.8%	83.9%	73.3%	68.1%	63.6%	60.0%	NA	NA	121.0%	65.6%	61.5%
CoreWeave (ex-CIP)	327.2%	217.5%	352.1%	177.0%	186.0%	176.7%	159.7%	150.4%	454.3%	189.0%	257.5%	166.1%	107.0%
Nebius	619.8%	971.6%	765.1%	655.8%	355.5%	290.3%	247.8%	218.2%	686.7%	761.1%	736.7%	266.1%	166.1%
Other Neoclouds	374.4%	356.6%	434.8%	288.7%	235.1%	211.2%	188.3%	173.4%	467.8%	243.0%	355.1%	197.8%	129.6%
Top China	**17.3%**	**16.0%**	**18.4%**	**19.2%**	**19.5%**	**17.0%**	**19.8%**	**20.8%**	**7.8%**	**11.9%**	**17.8%**	**19.3%**	**21.0%**
Alibaba	12.0%	15.7%	17.1%	13.0%	12.6%	15.8%	16.7%	13.3%	6.9%	11.9%	14.5%	14.6%	15.5%
Tencent	24.5%	15.5%	19.7%	27.5%	28.8%	17.4%	23.4%	31.4%	9.4%	12.3%	21.8%	25.3%	28.2%
Baidu	14.4%	20.8%	21.7%	24.0%	17.7%	24.3%	25.4%	27.0%	7.0%	9.6%	20.3%	23.7%	27.6%
Total	**34.0%**	**39.0%**	**48.4%**	**44.3%**	**46.6%**	**46.1%**	**47.5%**	**45.1%**	**17.0%**	**25.9%**	**41.8%**	**46.3%**	**41.4%**
NVDA/AMD/AVGO DC	50.7%	49.1%	45.3%	49.4%	57.1%	60.4%	64.5%	68.6%	48.4%	46.1%	48.4%	63.0%	77.9%

Source: Company reports, BofA Global Research estimates, Bloomberg

BofA GLOBAL RESEARCH

CY26/27 capex estimates have consistently increased over time

Below, we highlight that Street estimates for global cloud capex have steadily increased over the last 12 months. On an absolute $ basis, the outlook for CY26E capex (excl. leases and neoclouds for like-for-like comp) has risen +73% YoY, with CY27E capex outlook rising +118% YoY.

Including capital leases (as in Exhibit 1) but still excluding neoclouds (for like-for-like comp), the outlook for CY26E capex has risen +67% YoY and CY27E capex +111%.

Exhibit 13: Estimates for 2026 & 2027 have increased by ~65-120% over the last 12 months
Global cloud capex estimate changes over past 12 months (excl. leases and incl. leases) (excl. neoclouds)

		Oct' 25	Nov' 25	Dec' 25	Jan' 26	Feb' 26	Apr' 26	May' 26	Jul' 26	Aug' 26	Sep' 26	YoY
Capex ($mn)	**2025**	416,273	424,034	432,400	431,851	442,182	442,620	443,175	442,796	443,125	442,933	6.4%
(excl. leases)	**2026E**	495,915	559,725	586,274	595,004	668,563	711,130	778,672	805,209	835,982	858,952	73.2%
(excl. neoclouds)	**2027E**	544,738	641,371	689,512	693,544	783,684	833,079	962,771	997,408	1,146,851	1,184,843	117.5%
	2028E			684,268	701,043	796,911	902,738	1,057,661	1,083,521	1,306,628	1,318,459	NA
Change (%)	**2025**	2.6%	1.9%	2.0%	-0.1%	2.4%	0.1%	0.1%	-0.1%	0.1%	0.0%	
	2026E	5.5%	12.9%	4.7%	1.5%	12.4%	6.4%	9.5%	3.4%	3.8%	2.7%	
	2027E	6.1%	17.7%	7.5%	0.6%	13.0%	6.3%	15.6%	3.6%	15.0%	3.3%	
	2028E		NA	NA	2.5%	13.7%	13.3%	17.2%	2.4%	20.6%	0.9%	
		Oct' 25	**Nov' 25**	**Dec' 25**	**Jan' 26**	**Feb' 26**	**Apr' 26**	**May' 26**	**Jul' 26**	**Aug' 26**	**Sep' 26**	**YoY**
Capex ($mn)	**2025**	442,636	466,151	472,039	471,559	478,486	478,927	479,486	478,931	479,260	479,068	8.2%
(incl. leases)	**2026E**	527,549	610,689	631,338	641,057	748,374	749,684	806,601	851,253	858,772	878,586	66.5%
(excl. neoclouds)	**2027E**	575,937	691,332	734,062	739,160	868,655	871,799	1,005,472	1,145,382	1,184,064	1,212,899	110.6%
	2028E		724,269	756,666	764,846	908,762	942,031	1,101,998	1,300,391	1,372,197	1,351,279	NA
Change (%)	**2025**	2.4%	5.3%	1.3%	-0.1%	1.5%	0.1%	0.1%	-0.1%	0.1%	0.0%	
	2026E	5.5%	15.8%	3.4%	1.5%	16.7%	0.2%	7.6%	5.5%	0.9%	2.3%	
	2027E	6.2%	20.0%	6.2%	0.7%	17.5%	0.4%	15.3%	13.9%	3.4%	2.4%	
	2028E		NA	4.5%	1.1%	18.8%	3.7%	17.0%	18.0%	5.5%	-1.5%	

Source: BofA Global Research estimates, Bloomberg

On an operating cash flow basis, cloud capex (excl. neoclouds) is expected to reach ~105-110% of total operating cash flow in CY26-28E. While extremely elevated versus historical 35-50% between 2018 and 2024, we flag US telco capex had also reached 65-70% during the peak of the 4G/5G infra buildout phase, and AI infra investment has significantly more upside potential than telco infra buildout.

Exhibit 14: Telco capex has reached as high as 65-70% during the peak of its 4G/5G infra buildout phase of 2014-2017, vs. cloud capex expected to reach ~105-110% over the next few years on expected AI infra buildout
Telco vs. Cloud Capex as a % of Cash Flow from Ops

Cloud CFO	**2018**	**2019**	**2020**	**2021**	**2022**	**2023**	**2024**	**2025**	**2026E**	**2027E**	**2028E**	**CAGR '19-'24**	**CAGR '25-'28**	**CAGR '15-'25**
Cash Flow from Ops ($bn)	$213	$251	$339	$347	$337	$441	$540	$659	$825	$1,124	$1,239	16.5%	23.4%	19.3%
Capex w/ Leases ($bn)	$78	$86	$112	$148	$172	$169	$280	$479	$879	$1,213	$1,351	26.8%	41.3%	32.8%
Free Cash Flow ($bn)	**$135**	**$172**	**$206**	**$203**	**$165**	**$272**	**$260**	**$180**	**($54)**	**($89)**	**($113)**	**8.6%**	**-185.6%**	**7.9%**
Cloud Capex / CFO (%)	**36.5%**	**34.1%**	**33.0%**	**42.8%**	**50.9%**	**38.4%**	**51.9%**	**72.7%**	**106.5%**	**107.9%**	**109.1%**			
FCF Margin (%, RHS)	**19.3%**	**20.7%**	**20.1%**	**15.5%**	**11.7%**	**17.5%**	**15.1%**	**9.1%**	**-2.2%**	**-3.0%**	**-2.9%**			

公众号：独家的一手信息，独立的行业思考

Telco CFO	**2018**	**2019**	**2020**	**2021**	**2022**	**2023**	**2024**	**2025**	**2026E**	**2027E**	**2028E**	**CAGR '19-'24**	**CAGR '25-'28**	**CAGR '15-'25**
Cash Flow from Ops ($bn)	$87	$95	$97	$88	$91	$99	$102	$105	$108	$115	$125	1.3%	5.8%	2.8%
Capex w/ Leases ($bn)	$43	$44	$45	$48	$57	$46	$46	$48	$49	$49	$49	1.0%	1.1%	1.2%
Free Cash Flow ($bn)	**$44**	**$51**	**$52**	**$40**	**$34**	**$53**	**$55**	**$58**	**$59**	**$66**	**$75**	**1.6%**	**9.4%**	**4.3%**
Telco Capex / CFO (%)	**49.8%**	**46.2%**	**46.4%**	**54.8%**	**62.4%**	**46.8%**	**45.5%**	**45.4%**	**45.4%**	**42.4%**	**39.5%**			
FCF Margin (%)	**12.7%**	**14.3%**	**14.1%**	**11.2%**	**10.1%**	**15.7%**	**16.4%**	**16.3%**	**16.1%**	**17.8%**	**19.8%**			

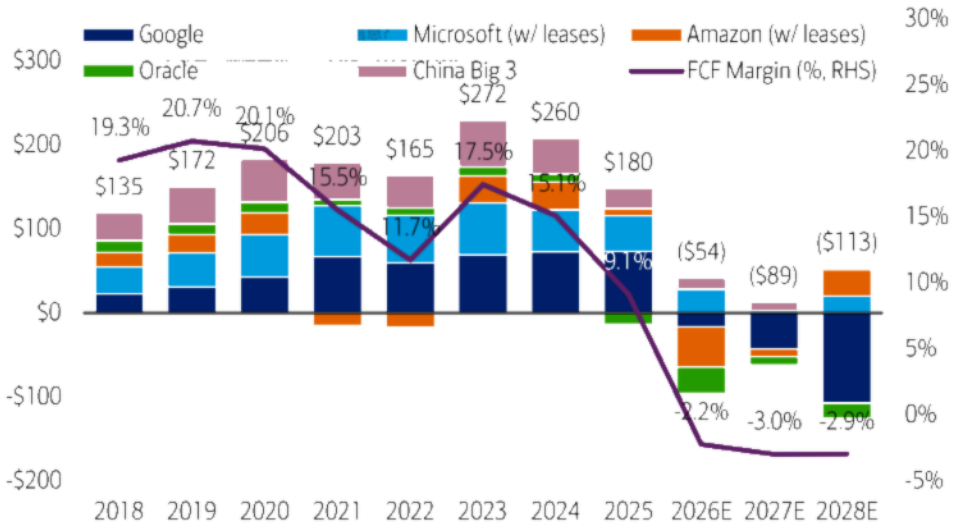
Source: BofA Global Research estimates, Bloomberg

Moreover, we expect cloud/hyperscale (excl. neoclouds) aggregate FCF margin to reach ~2% in 2026E and decline to ~3% range in 2027-28E (vs. ~5-6% outlook in August). This

FCF margin prospect is much lower than the historical 15-20% level for hyperscale vendors and is driven by planned accelerated infra buildout over the next few years.

While extremely low (and negative), we view this near-term trend as still sustainable given they are nearly fully funded by operating cash flow (i.e. with cash generated by business), and we expect a material improvement in OCF/FCF in the out years (2029+) as prior AI investments proliferate and begin generating ROI/revenue/profit.

Exhibit 15: We expect cloud/hyperscale aggregate FCF margin to reach -2% in 2026E and -3% in 2027-28E given accelerated infra buildout in the next few years
FCF and FCF Margin Outlook for Each Hyperscaler



Source: BofA Global Research estimates, Bloomberg

Capital raises: low cost of debt, sufficient future FCF

We flag the top 5 hyperscalers have now collectively raised **~$370bn of capital** so far in 2026. Note these are **long-dated debt** (3-40 years mostly) and permanent equity capital rather than short-term funding. Given the long duration of this capital, we view them as providing balance sheet flexibility to fund multi-year AI infra buildouts, rather than reflecting near-term liquidity needs.

Importantly, we expect the **aggregate annual FCF of hyperscalers to well exceed** prior average of **~$200-250bn+/yr** (2020-2024) once these AI infra are fully at work, and do not view the recent fundraising activities as a signal for financial stress or inadequate free cash flow generation.

Exhibit 16: We see a total of ~$370bn of capital raise collectively by top 5 hyperscalers YTD
Top 5 Hyperscaler Capital Raises in 2026 YTD

Company	Date	Type	USD ($bn)
Amazon	Sep-26	Sterling bonds	$5.8
Amazon	Jul-26	Bond	$25.0
Amazon	Jun-26	Loan Facility	$17.5
Amazon	Jun-26	Canadian Bonds	$10.0
Amazon	May-26	Swiss-franc bonds	$3.5
Amazon	Mar-26	Euro bonds	$16.6
Amazon	Mar-26	U.S. bonds	$37.0
Alphabet	Aug-26	Bond	$25.0
Alphabet	Jun-26	Senior Notes	$20.3
Alphabet	Jun-26	Equity Offering	$90.0
Alphabet	Feb-26	Global Bonds	$31.5
Meta	Jul-26	Project-Finance Bonds ($25bn, 20% owned)	$12.5
Meta	May-26	Bonds	$25.0
Oracle	Aug-26	ATM Equity Program	$20.0
Oracle	Feb-26	Mandatory Convertible Preferred	$5.0

Exhibit 16: We see a total of ~$370bn of capital raise collectively by top 5 hyperscalers YTD
Top 5 Hyperscaler Capital Raises in 2026 YTD

Company	Date	Type	USD ($bn)
Oracle	Feb-26	Bonds	$25.0
Microsoft	N/A	N/A	N/A
Total	**2026**	**Capital Raise**	**~$370.0**

Source: Company data, BofA Global Research estimates

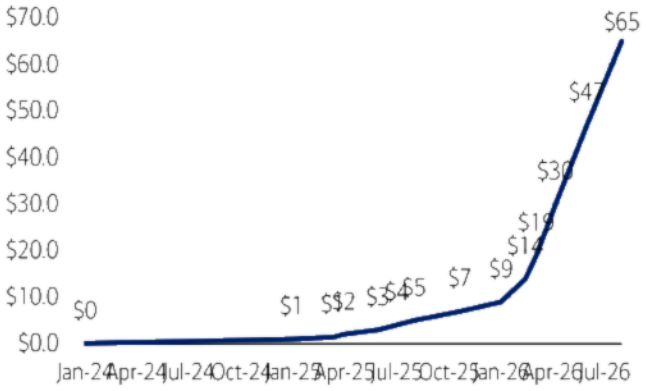
Capex sustainability: demand signals increasingly positive

We see the AI capex cycle as sustainable as hyperscaler spending is increasingly supported by: 1) frontier lab revenue acceleration, 2) contracted cloud/backlog growth, and 3) token-intensive agentic workloads that convert usage into compute demand. A key risk remains execution (power, data center construction, financing, and component supply) but we see surging demand outpacing the capacity being built resulting in sustainably strong demand in the mid-term.

Frontier lab growth is becoming large enough to absorb compute

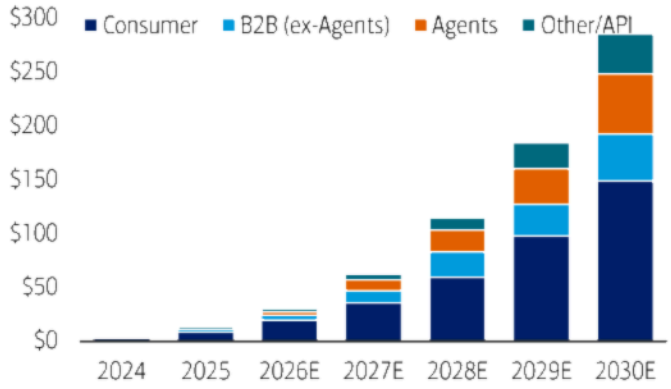
OpenAI and Anthropic provide the clearest evidence that AI demand is scaling faster than traditional enterprise software adoption curves. OpenAI reportedly expects revenue to increase from $13bn in 2025 to $30bn in 2026, $62bn in 2027, and $283bn by 2030, while projecting cumulative compute spend of ~$665bn through 2030. OpenAI's higher revenue and compute projections strengthen the AI demand thesis, though lower margins and higher cash burn remain a profitability challenge for the sector. BofA's Justin Post sees agentic AI contributing roughly 20% of 2030 revenue, supported by agent packages priced from $1,500 to $20,000 per user per month.

Exhibit 17: Anthropic's ARR has inflected to >$65bn as of Jul'26
Anthropic's annualized revenue run rate ($bn)



Source: Anthropic, CNBC, Inc.com, CNBC, PYMNTS, Bloomberg

Exhibit 18: Total OpenAI revenue is expected to reach $284bn in 2030E, with $56bn coming from agents
OpenAI segment revenue, 2024-30E



Source: Company reports, The Information

Anthropic is a near-term proof point for AI workload monetization with annualized revenue run rate surpassing $19bn, up $17bn YoY and $10bn from the end of 2025. Anthropic's ARR increase from $9bn in December to $19bn in March could suggest a $2.5bn+ QoQ revenue increase, and that AWS could see up to $1bn QoQ revenue contribution related to Anthropic depending on workload share and training cost capture. This creates a direct bridge from frontier lab ARR to cloud capacity demand and therefore to accelerator, HBM, networking, custom silicon, and power/thermal content.

Token velocity extends the cycle beyond training

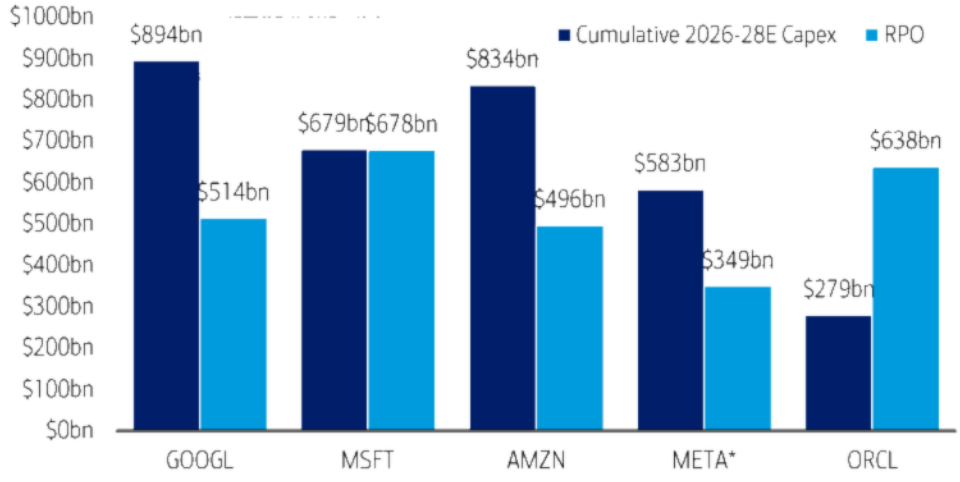
Agentic AI workflows can generate 10x to 1,000x more tokens per workflow than a chat interaction, and that token consumption is becoming the primary driver of compute demand. Tokens also change cloud monetization from seat-based economics to usage-based economics, giving cloud platforms a clearer pathway to monetize agent adoption.

Efficiency gains should not be viewed simply as a capex headwind; lower cost per token can expand usage, similar to prior compute elasticity cycles. Inference, post-training, synthetic data generation, agent orchestration, and enterprise workflow automation can sustain demand even after the initial training-cluster buildout matures.

Capex is increasingly backed by commitments and backlog

Hyperscaler capex is increasingly backed by customer commitments (via RPO) and backlog, giving further demand visibility and sustainability into the out years.

Exhibit 19: Much of hyperscaler capex is supported by RPO
Cumulative CY26-28E capex estimates vs. current RPO across hyperscalers



Source: BofA Global Research, Company Reports, Bloomberg

Alphabet said Google Cloud revenue grew +82% YoY to ~$25bn, with backlogs now reaching $514bn. More than half of the backlog is expected to convert over the next two years, with the vast majority expected to be recognized in 2027.

Microsoft disclosed nearly $678bn of commercial RPO, up +8% QoQ and +84% YoY, with ~30% expected to be recognized within 12 months. Sequential growth was driven by customers other than frontier-model developers. OpenAI represented ~45% (~$250bn) of Azure RPO as of Dec'25, though no update since. Azure revenue grew +43% in F4Q26 (Jun'26) and surpassed $100bn for FY26. MSFT added 31 DCs and 1 GW of capacity in JunQ and remains on track to double overall capacity within two years.

Amazon disclosed AWS backlog of $496bn, up >100% YoY, while AWS revenue accelerated toward ~$169bn annualized run-rate. AWS's AI and chips businesses each surpassed $25bn/yr, also growing at >100% YoY. Both Anthropic and OpenAI have made multi-year, multi-gigawatt Trn commitments.

Meta does not yet have a cloud resale model (it may in the future), but it disclosed $349bn of multi-year contractual (supplier) commitment balance in Q2, underscoring that its internal AI compute demand remains strong.

Oracle disclosed $638bn of RPO in F4Q26, up +363% YoY and up +$85bn QoQ. Mgmt indicated that ~12% is expected to be recognized within 12 months and another ~34% during months 13 to 36, highlighting the long-duration nature of the backlog. Most of the FQ3/FQ4 RPO increases were related to large-scale AI contracts under which customers either prepaid Oracle to purchase GPUs or supplied the GPUs directly, with the prepaid hardware contracts reaching ~$75bn. Importantly, this structure reduces Oracle's upfront funding requirements.

Global Semis Forecast Update

We model CY26E semis/core semis (ex-memory) growth of +113%/+30% YoY, led by growth in memory, data center, and greatly improved outlook in industrial and automotive, offset by continued unit headwinds in PCs/smartphones/consumer.

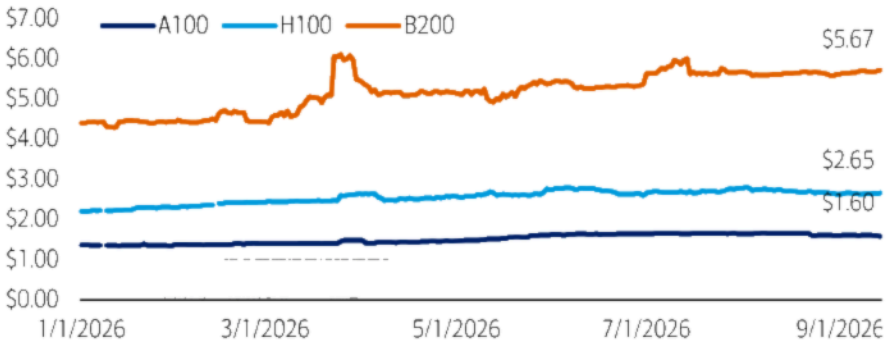
By end market, we now model (1) compute and storage up +50% YoY (continued server strength); (2) wireless comms down -8% YoY (smartphone unit headwind); (3) auto sales up +12% on weak units but quickly improving content; (4) Industrial up +32% YoY on improved end demand and inventory dynamics since 2H25; (5) consumer down -7% YoY; and (6) wired comms up +29% YoY on data center related infra buildout.

Exhibit 20: We now see total semis sales of $3.4Tn (+19% CAGR CY26-30E), vs. $3.2Tn prior (+18% CAGR), with core semis now at $1.4Tn (+16% CAGR) vs. $1.4Tn prior (+16% CAGR)
BofA Semiconductor forecasts by end market: New vs. Old

	New			Old		
Semiconductor Sales Summary	**2026**	**2030**	**CAGR**	**2026**	**2030**	**CAGR**
Total Semis Sales ($bn)	**$1,676.9**	**$3,383.1**	**19%**	**$1,676.9**	**$3,203.7**	**18%**
YoY	113.2%	15.4%		113.2%	11.0%	
Memory	$937.4	$2,032.4	**21%**	$937.4	$1,853.1	**19%**
YoY	326.9%	19.4%		326.9%	11.9%	
Core Semis (ex-memory, $bn)	**$739.5**	**$1,350.7**	**16%**	**$739.5**	**$1,350.7**	**16%**
YoY	30.4%	9.9%		30.4%	9.9%	
Servers (traditional/AI)	**$359.6**	**$848.7**	24%	**$359.9**	**$848.7**	24%
YoY	73.8%	12.8%		74.0%	12.8%	
PC	**$56.0**	**$68.7**	5%	**$55.8**	**$68.7**	5%
YoY	(8.8%)	2.5%		(9.3%)	2.5%	
Smartphone	**$71.1**	**$77.8**	2%	**$71.1**	**$77.8**	2%
YoY	(9.2%)	2.2%		(9.2%)	2.2%	
Automotive	**$60.0**	**$90.2**	11%	**$60.0**	**$90.2**	11%
YoY	12.0%	10.2%		12.0%	10.2%	
Industrial	**$65.7**	**$99.2**	11%	**$65.7**	**$99.2**	11%
YoY	32.0%	5.6%		32.0%	5.6%	
Other (storage, wired, wireless infra, consumer)	**$127.0**	**$166.1**	7%	**$127.0**	**$166.1**	7%
YoY	8.6%	5.5%		8.6%	5.5%	
Wafer Fab Equipment Spend	**2026**	**2030**	**CAGR**	**2026**	**2030**	**CAGR**
Total WFE ($bn)	**$155.9**	**$359.8**	**23%**	**$155.9**	**$359.8**	**23%**
YoY	33.4%	17.6%		33.4%	17.6%	
WFE Intensity (% of semis sales)	9.3%	10.6%		9.3%	11.2%	
Memory WFE ($bn)	**$61.3**	**$144.4**	24%	**$61.3**	**$144.4**	24%
YoY	45.7%	16.5%		45.7%	16.5%	
WFE Intensity (% of semis sales)	6.5%	7.1%		6.5%	7.8%	
Non-Memory WFE ($bn)	**$94.6**	**$215.3**	23%	**$94.6**	**$215.3**	23%
YoY	26.5%	18.3%		26.5%	18.3%	
WFE Intensity (% of semis sales)	12.8%	15.9%		12.8%	15.9%	

Source: BofA Global Research estimates, Mercury Research, Gartner, Omdia, SIA

Exhibit 21: NVDA B200 spot rental pricing is still ~$5.70 vs. prior $6.10 peak from March 2026, suggesting continued strong broad-based demand
NVDA A100, H100, B200 GPU Rental Pricing Index



Source: Bloomberg, Silicon Data

Exhibit 22: We model semis/core semis sales up +113%/+30% YoY in CY26E
Summary of BofA Semiconductor forecasts by end market

Revenue ($mn)	2020	2021	2022	2023	2024	2025	2026E	2027E	2028E	2029E	2030E	CAGR '16-26	CAGR '21-26	CAGR '26-30
Total Semis	**$451**	**$572**	**$594**	**$528**	**$633**	**$787**	**$1,677**	**$2,323**	**$2,611**	**$2,931**	**$3,383**	**17.3%**	**24.0%**	**19.2%**
YoY%	7.7%	27.0%	3.8%	(11.0%)	19.7%	24.3%	113.2%	38.5%	12.4%	12.2%	15.4%			
Memory	**$128**	**$170**	**$150**	**$94**	**$170**	**$220**	**$937**	**$1,402**	**$1,525**	**$1,702**	**$2,032**	**28.4%**	**40.6%**	**21.3%**
YoY%	13.5%	33.4%	(11.9%)	(37.4%)	81.3%	28.9%	326.9%	49.6%	8.7%	11.7%	19.4%			
Core Semis (ex-memory)	**$323**	**$402**	**$444**	**$435**	**$463**	**$567**	**$739**	**$921**	**$1,087**	**$1,229**	**$1,351**	**10.9%**	**13.0%**	**16.3%**
YoY%	5.6%	24.4%	10.5%	(2.1%)	6.4%	22.6%	30.4%	24.5%	18.0%	13.1%	9.9%			
Compute and Storage	**$108**	**$128**	**$154**	**$167**	**$205**	**$297**	**$445**	**$599**	**$735**	**$851**	**$950**	**18.3%**	**28.3%**	**20.9%**
YoY%	11.2%	19.1%	20.3%	8.4%	22.5%	45.3%	49.8%	34.6%	22.6%	15.9%	11.6%			
PCs	**$56**	**$68**	**$59**	**$52**	**$57**	**$61**	**$56**	**$57**	**$64**	**$67**	**$69**	**3.3%**	**(3.8%)**	**5.2%**
YoY%	16.5%	22.3%	(12.4%)	(11.8%)	9.3%	7.2%	(8.8%)	1.8%	13.0%	4.0%	2.5%			
Servers (silicon only)	**$27**	**$32**	**$63**	**$80**	**$119**	**$207**	**$360**	**$512**	**$639**	**$753**	**$849**	**34.5%**	**62.5%**	**23.9%**
YoY%	2.2%	16.9%	97.3%	27.5%	48.7%	73.9%	73.8%	42.3%	24.9%	17.8%	12.8%			
Wireless Communications	**$85**	**$104**	**$111**	**$94**	**$93**	**$91**	**$83**	**$83**	**$87**	**$90**	**$92**	**2.4%**	**(4.4%)**	**2.5%**
YoY%	11.1%	22.3%	6.8%	(15.8%)	(0.5%)	(2.4%)	(8.4%)	(0.6%)	5.0%	3.3%	2.5%			
Smartphone	**$71**	**$88**	**$91**	**$77**	**$80**	**$78**	**$71**	**$70**	**$74**	**$76**	**$78**	**2.0%**	**(4.3%)**	**2.3%**
YoY%	12.3%	23.7%	3.2%	(15.4%)	3.6%	(2.0%)	(9.2%)	(1.2%)	5.0%	3.2%	2.2%			
Wireless Infrastructure	**$14**	**$16**	**$20**	**$16**	**$13**	**$13**	**$12**	**$13**	**$13**	**$14**	**$14**	**4.9%**	**(5.0%)**	**4.0%**
YoY%	5.0%	14.7%	27.1%	(17.5%)	(20.0%)	(5.0%)	(3.0%)	3.0%	5.0%	4.0%	4.0%			
Automotive	**$34**	**$46**	**$49**	**$59**	**$56**	**$54**	**$60**	**$66**	**$74**	**$82**	**$90**	**6.3%**	**5.4%**	**10.7%**
YoY%	(9.1%)	34.6%	5.4%	22.3%	(6.4%)	(3.7%)	12.0%	10.8%	11.8%	10.2%	10.2%			
Global Automotive Units (mn)	**74.6**	**77.2**	**82.3**	**90.5**	**92.5**	**93.1**	**91.3**	**91.7**	**93.1**	**95.0**	**96.9**	**(0.2%)**	**3.4%**	**1.5%**
YoY%	(16.1%)	3.5%	6.7%	9.9%	2.2%	0.6%	(1.9%)	0.4%	1.5%	2.0%	2.0%			
Auto semi content ($/LV) / Inv. Adj.	**$459**	**$597**	**$590**	**$657**	**$602**	**$576**	**$657**	**$725**	**$798**	**$862**	**$931**	**6.5%**	**1.9%**	**9.1%**
YoY%	8.4%	30.0%	(1.1%)	11.3%	(8.5%)	(4.3%)	14.2%	10.3%	10.1%	8.0%	8.0%			
Industrial & Other	**$43**	**$54**	**$62**	**$59**	**$45**	**$50**	**$66**	**$78**	**$88**	**$94**	**$99**	**6.7%**	**4.0%**	**10.8%**
YoY%	(2.6%)	26.5%	15.3%	(4.8%)	(24.3%)	10.8%	32.0%	18.0%	13.0%	7.2%	5.6%			
Automation	**$10**	**$13**	**$15**	**$15**	**$12**	**$13**	**$18**	**$21**	**$25**	**$26**	**$28**	**9.2%**	**7.0%**	**11.9%**
YoY%	(3.0%)	29.5%	17.0%	(1.0%)	(15.0%)	5.6%	35.0%	20.0%	15.0%	7.0%	6.0%			
Power/Energy	**$6**	**$7**	**$9**	**$9**	**$8**	**$8**	**$12**	**$16**	**$19**	**$22**	**$23**	**10.9%**	**10.4%**	**18.3%**
YoY%	(0.5%)	25.4%	22.0%	5.0%	(18.0%)	4.1%	50.0%	35.0%	20.0%	12.0%	8.0%			
Consumer	**$35**	**$48**	**$48**	**$30**	**$31**	**$33**	**$31**	**$30**	**$30**	**$32**	**$34**	**0.7%**	**(8.5%)**	**2.4%**
YoY%	6.5%	37.5%	0.1%	(37.1%)	3.4%	6.6%	(7.4%)	(3.4%)	1.3%	6.0%	5.8%			
Wired Communications	**$18**	**$21**	**$20**	**$25**	**$33**	**$42**	**$54**	**$65**	**$73**	**$80**	**$85**	**11.7%**	**20.7%**	**11.9%**
YoY%	1.0%	17.0%	(8.3%)	26.9%	32.8%	28.1%	29.3%	19.6%	12.2%	9.2%	6.9%			
Ethernet/Network switch	**$5**	**$5**	**$6**	**$9**	**$13**	**$17**	**$25**	**$31**	**$36**	**$40**	**$43**	**15.0%**	**34.9%**	**15.3%**
YoY%	(3.8%)	11.9%	15.0%	45.0%	40.0%	35.0%	42.0%	28.0%	15.0%	11.0%	8.0%			
Optical Equipment	**$4**	**$5**	**$5**	**$6**	**$9**	**$12**	**$15**	**$19**	**$21**	**$24**	**$26**	**22.7%**	**26.8%**	**13.4%**
YoY%	9.8%	11.8%	5.0%	30.0%	35.0%	35.0%	32.0%	20.0%	15.0%	11.0%	8.0%			

Source: BofA Global Research estimates, Mercury Research, Gartner, Omdia, SIA

In 2026E, we expect semis/core semis sales to grow +113%/+30% YoY. Memory is expected to be even stronger, now up nearly +327% YoY (after +81%/+29% YoY growth in CY24/25).

By device type, we see CY26E memory sales driven by both DRAM (+328% YoY) and NAND (+341% YoY).

Ex-memory, microprocessors (MPUs) strong (+32% YoY) on hyperscaler consumption modestly offset by weak PC units, and logic particularly strong (+41% YoY) on AI accelerator-related demand. We also model industrial-centric markets (MCUs, Analog) seeing recovery following inventory digestion throughout CY24-25. In CY26E, we model other markets (discretes, optos, sensors) to generally return to healthy growth (+13% YoY) as well, particularly driven by data center-related optoelectronic demands.

Exhibit 23: We model memory, logic, microcomponents driving semiconductor growth
Summary of BofA Semis Forecast by device type

Revenue ($bn)	2020	2021	2022	2023	2024	2025	2026E	2027E	2028E	2029E	2030E	CAGR '16-26	CAGR '21-26	CAGR '26-30
Total Semis	**$451**	**$572**	**$594**	**$528**	**$633**	**$787**	**$1,677**	**$2,323**	**$2,611**	**$2,931**	**$3,383**	**17.3%**	**24.0%**	**19.2%**
YoY%	7.7%	27.0%	3.8%	(11.0%)	19.7%	24.3%	113.2%	38.5%	12.4%	12.2%	15.4%			
Memory	**$128**	**$170**	**$150**	**$94**	**$170**	**$220**	**$937**	**$1,402**	**$1,525**	**$1,702**	**$2,032**	**28.4%**	**40.6%**	**21.3%**
YoY%	13.5%	33.4%	(11.9%)	(37.4%)	81.3%	28.9%	326.9%	49.6%	8.7%	11.7%	19.4%			
DRAM	**$65**	**$92**	**$78**	**$47**	**$88**	**$134**	**$573**	**$917**	**$1,025**	**$1,158**	**$1,396**		**44.1%**	**24.9%**
YoY%	4.6%	40.9%	(15.7%)	(39.1%)	85.9%	52.2%	328.2%	60.1%	11.7%	13.1%	20.5%			
NAND	**$59**	**$73**	**$67**	**$42**	**$78**	**$81**	**$357**	**$477**	**$492**	**$535**	**$627**		**37.3%**	**15.1%**
YoY%	26.4%	25.3%	(8.2%)	(36.9%)	84.0%	3.7%	341.4%	33.4%	3.1%	8.8%	17.2%			
Core Semis (ex-memory)	**$323**	**$402**	**$444**	**$435**	**$463**	**$567**	**$739**	**$921**	**$1,087**	**$1,229**	**$1,351**	**10.9%**	**13.0%**	**16.3%**
YoY%	5.6%	24.4%	10.5%	(2.1%)	6.4%	22.6%	30.4%	24.5%	18.0%	13.1%	9.9%			
Analog	**$56**	**$74**	**$89**	**$81**	**$80**	**$87**	**$98**	**$112**	**$124**	**$132**	**$136**	**7.4%**	**5.8%**	**8.6%**
YoY%	3.2%	33.1%	20.1%	(8.8%)	(2.0%)	8.8%	13.3%	13.8%	10.9%	6.3%	3.6%			
Microcomponents	**$70**	**$81**	**$79**	**$77**	**$79**	**$85**	**$111**	**$150**	**$188**	**$222**	**$254**	**6.2%**	**6.6%**	**23.1%**
YoY%	4.9%	15.5%	(1.5%)	(3.5%)	3.7%	7.0%	30.6%	35.3%	25.1%	18.2%	14.6%			
Microprocessors	**$51**	**$56**	**$51**	**$45**	**$53**	**$59**	**$82**	**$119**	**$154**	**$185**	**$213**	**6.8%**	**7.9%**	**26.8%**
YoY%	6.5%	10.4%	(9.5%)	(11.2%)	18.0%	11.1%	39.0%	44.3%	29.7%	20.0%	15.0%			
Microcontrollers	**$18**	**$23**	**$27**	**$29**	**$24**	**$23**	**$26**	**$29**	**$31**	**$35**	**$39**	**4.7%**	**3.0%**	**10.8%**
YoY%	3.3%	25.0%	20.8%	7.6%	(19.2%)	(2.4%)	13.3%	10.7%	7.8%	11.3%	13.3%			
Logic	**$118**	**$154**	**$176**	**$178**	**$213**	**$301**	**$424**	**$525**	**$618**	**$701**	**$773**	**16.6%**	**22.4%**	**16.2%**
YoY%	11.1%	30.2%	14.2%	1.4%	19.1%	41.5%	40.7%	24.0%	17.6%	13.4%	10.2%			
Others (Discretes, optos, sensors)	**$79**	**$93**	**$100**	**$98**	**$91**	**$95**	**$107**	**$134**	**$157**	**$175**	**$188**	**5.6%**	**2.9%**	**15.0%**
YoY%	0.3%	17.4%	7.1%	(1.2%)	(7.5%)	4.0%	13.1%	25.1%	17.5%	10.9%	7.4%			

Source: BofA Global Research estimates, Mercury Research, Gartner, Omdia, SIA

Looking at our forecasts versus BofA Global Semis bottom-up estimates, we are generally in line with the growth trajectory for CY26E/CY27E ex-NVDA (which now realizes significant revenue from system sales growth and software). Modest outperformance versus industry reflects continued industry trends of industry consolidation and share gains by top vendors.

Exhibit 24: Our Semis industry outlook generally aligns with bottom-up company estimates
Summary of BofA Semis forecasts versus bottom-up estimates

Semis ex-mem Forecast ($bn)	2025	2026E	2027E	2028E	2029E	2030E
BofA Semi Ex-Memory Forecast	$567	$739	$921	$1,087	$1,229	$1,351
YoY (%)	*22.6%*	*30.4%*	*24.5%*	*18.0%*	*13.1%*	*9.9%*
BofA Coverage Bottom-Up Sales	$539	$767	$984	$1,208	$1,890	$2,116
YoY (%)	*26.8%*	*42.4%*	*28.2%*	*22.8%*	*56.4%*	*12.0%*
BofA Coverage Bottom-Up Ex-NVDA Sales	$323	$407	$519	$650	$804	$944
YoY (%)	*9.6%*	*26.2%*	*27.3%*	*25.3%*	*23.7%*	*17.4%*

Source: BofA Global Research estimates, SIA, company reports

Glossary:

ACC: Active copper cable
AEC: Active electrical cable
AGI: Arm AGI CPU product designation; not expanded in the report
AI: Artificial intelligence
ARR: Annualized revenue run rate
ASIC: Application-specific integrated circuit
ASP: Average selling price
ATM: At-the-market equity offering
AWS: Amazon Web Services
bps: Basis points
CAGR: Compound annual growth rate
capex: Capital expenditure
CIP: CoreWeave software
CPO: Co-packaged optics
CPU: Central processing unit
CQ: Calendar quarter
CPX: NVIDIA Vera CPX product designation
DAC: Direct-attach copper cable
DC: Data center
DDR: Double data rate memory
DRAM: Dynamic random-access memory
FQ: Fiscal quarter
GPU: Graphics processing unit
GTC: NVIDIA GTC conference, originally GPU Technology Conference
GW: Gigawatt
HBM: High-bandwidth memory
I/O: Input/output
IaaS: Infrastructure as a service
IDC: International Data Corporation
IT: Information technology
LV: Light vehicle
MCU: Microcontroller unit
MPU: Microprocessor unit
NAND: NOT-AND flash memory
NIC / SmartNIC: Network interface card / programmable network interface card
NPO: Near-packaged optics
OCS: Optical circuit switching
PC: Personal computer
RAS: Reliability, availability and serviceability
ROI: Return on investment
RPO: Remaining performance obligation
SIA: Semiconductor Industry Association
SOX: PHLX Semiconductor Sector Index
SPX: S&P 500 Index ticker
SSD: Solid-state drive
TAM: Total addressable market
TBD: To be determined
TPU: Tensor processing unit
Trn: AWS Trainium
WFE: Wafer fabrication equipment
XPU: General term for an accelerator or other processing unit

ADI: Analog Devices
AMAT: Applied Materials
AMD: Advanced Micro Devices
ARM: Arm Holdings

AVGO: Broadcom
CRDO: Credo Technology Group
INTC: Intel
IREN: IREN
KLAC: KLA Corporation
LRCX: Lam Research
MRVL: Marvell Technology
MSFT: Microsoft
MU: Micron Technology
NVDA: NVIDIA
ON: onsemi
QCOM: Qualcomm
SPCX: SpaceX
xAI: xAI

Exhibit 25: Stocks mentioned
Prices and ratings for stocks mentioned in this report

BofA Ticker	Bloomberg ticker	Company name	Price	Rating
AMD	AMD US	Advanced Micro	US$ 607.57	C-1-9
ADI	ADI US	Analog Devices	US$ 398.22	B-1-7
AMAT	AMAT US	Applied Materials	US$ 512.01	C-1-7
AVGO	AVGO US	Broadcom	US$ 355.1	C-1-7
INTC	INTC US	Intel	US$ 115.93	C-1-9
LRCX	LRCX US	Lam Research	US$ 323.86	C-1-7
MRVL	MRVL US	Marvell	US$ 263.27	C-1-7
MU	MU US	Micron	US$ 1065.08	C-1-7
NVDA	NVDA US	NVIDIA	US$ 227.21	C-1-7
ON	ON US	ON Semiconductor	US$ 75.95	C-1-9

Source: BofA Global Research

BofA GLOBAL RESEARCH

Price objective basis & risk

Advanced Micro Devices, Inc (AMD)

Our $720 PO is based on 30x our 2028E non-GAAP EPS, including SBC. Our PO basis is in the middle of historical 13x-58x range, 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.

Analog Devices Inc. (ADI)

Our $500 PO is based on 23x CY2028E P/E. Our PO basis is within the 14x-40x historical range and at a slight premium to analog semi peers. This is justified, in our view, based on ADI's best-in-class profitability and differentiated sources of both cyclical and secular growth across AI, industrial, auto, and comms.

Downside risks to our price objective: 1) Economic downturn, which could reduce demand for automotive, industrial products, impacting gross margins, especially given recent capital expenditures and higher fixed cost footprint. 2) Inability to realize the planned cost synergies from the Maxim combination. 3) Competition from larger vendors that have lower-cost production facilities. 4) Exposure to US-China tensions/tariffs via Chinese automotive market.

Applied Materials, Inc. (AMAT)
Our PO of $650 is based on 27x our CY28E P/E estimate. Our PO basis is near the middle of AMAT's historical 10-42x trading range justified given the potential for WFE outgrowth in 2026/27E and remains at a discount to other large peers given more balanced growth profile and lower profitability.

Downside risks to our PO are: ongoing US government probe that we are unable to size the financial impact at this time, slower-than-expected capital spending cycle, delay in memory capacity adds, market share loss in deposition, implant, thermal, CMP, etch, or process control segments, merger & integrations risk, and macro headwinds.

Broadcom Inc (AVGO)
Our $530 price objective for Broadcom is based on 30x CY27E P/E, in the middle of its 11x-41x historical range, still in-line with 0.5x-2x PEG framework for high-growth compute vendors, and justified given double-digit EPS growth and best-in-semis profitability, FCF generation, and returns.

Downside risks to our price objective are: 1) semiconductor cycle risks, including sensitivity to fundamental or sentiment shifts in AI theme, 2) high exposure to Apple and Google with potential design out risks, 3) competitive risks in networking, smartphone, storage, enterprise software markets, including rising NVDA competition in networking, 4) frequent acquirer of assets, which increases financial and integration risks, and 5) large $60bn net-debt.

Intel (INTC)
Our $145 price objective is based on 29x 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 8.9x 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.

Lam Research Corp. (LRCX)
We assign a $385 PO based on 29x CY28E PE. Our PO basis is near the middle of the historical 9x-50x trading range justified by ongoing memory and leading-edge foundry/logic WFE cycle, high-teens EPS CAGR over time, etch/deposition product leadership, rising etch/deposition intensity, share gains, growing foundry/logic exposure

over memory, improving prospects of NAND greenfield along with continued upgrade spending,, and robust FCF generation, offset by near-term concerns around cost inflation and tariffs.

Upside risks are tech inflections, F/L share gains, NAND upgrades.

Downside risks are slower than expected capital spending cycle, delay in memory capacity adds, market share loss in etch or clean segments, merger & integrations risk, macro headwinds, customer consolidation and China.

Marvell Technology, Inc. (MRVL)
Our $365 price objective is based on 33x CY28E PE (excl. SBC), modestly higher than 26x historical median but within 14x-53x range and is justified by an improvement in visibility for major customer ASIC projects, broadening AI portfolio across connectivity/switch/compute, and cyclical industry risks offset by ASIC upside, networking strength, and AEC/CPO/scale-up share gains.

Upside risks: 1) Faster than anticipated ramp/visibility in major custom ASIC projects, 2) Continued growth in DSP-based pluggable market, versus new LPO/LRO techs, 3) Share gains in emerging AEC/CPO/scale-up switch markets against incumbents.

Downside risks: 1) Loss of visibility in key custom ASIC projects, particularly in the next-gen 3nm/2nm chips at AWS and Microsoft, 2) Competition in AI compute, with merchant vendors continuing to proliferate and ASIC incumbent AVGO winning many of the large hyperscaler/AI customers, 3) cyclical industry risks including potential slowdown in legacy storage, enterprise networking, carrier markets.

Micron Technology, Inc (MU)
Our $1550 PO is based on a sum-of-parts valuation that values: (1) traditional cyclical memory business at $1,040/sh at 3x CY28E P/B, toward the high-end of MU's long-term range 0.8x-3.1x as we are potentially in a memory upcycle, and (2) AI HBM business at 31x CY28E PE, in-line with AI compute peer group median.

Downside risks: (1) larger than expected memory ASP decline, (2) greater competition from China newcomers, (3) share loss to large competitors, (4) softening of demand across major end markets such as data center, smartphones, or PCs.

NVIDIA Corporation (NVDA)
Our $350 PO is based on 22x CY27E PE ex cash, within NVDA's historical 15x-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.

onsemi (ON)
Our PO of $120 PO is based on 20x 2028E P/E. Our PO basis is in the middle of ON's 9x-28x trading range and auto/industrial peers trading 14x-29x, justified, in our view, given potential for strong EPS and FCF recovery along with ongoing content and share gains in areas like silicon carbide and image sensing which should drive strong secular growth longer-term. We also think this valuation credits the company's restructuring initiatives that should streamline its manufacturing and opex base.

Downside risks to our PO are 1) macro/cyclical risks, given high exposure to automotive and industrial markets, make ON susceptible to any potential global trade tensions/tariffs, 2) prolonged COVID-19 headwinds limiting pace of automotive/industrial recovery, impacting utilization levels, 3) difficulty in ramping 300mm fabrication facility limiting gross margin improvement, 4) sustained elevated capex levels relative to peers.

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.

Special Disclosures

BofA Securities is currently acting as exclusive financial advisor to a consortium comprising IMM Investment and Stonebridge Capital in connection with its proposed acquisition of 20% stake in a spin-off company, SK Horizon Co., Ltd., (a data center and subsea cable business carved out from SK Broadband Company Limited), which was announced on 27th August 2026.

BofA Securities is currently acting as financial advisor to One Investment Management Ltd in connection with the proposed acquisition with PayPay Corporation, a subsidiary of SoftBank Group Corp., of shares of T&D Financial Life Insurance Company, a consolidated subsidiary of T&D Holdings, Inc., which was announced on June 4, 2026.

BofA Securities is currently acting as joint financial advisor to ABB Ltd in connection with the sale of its Robotics Division to SoftBank Group Corp, which was announced on October 8, 2025.