06 Oct 2026 05:00:00 ET | 21 pages

US Semiconductors

Meta Muse Personal Agent Implications – AMD to $800

Atif Malik [AC]

Papa Sylla

Kelsey Chia, CFA

CITI'S TAKE

Following the launch of several personal AI agents recently, including Muse, Dots, Instinct, GrokBot, and GeminiSpark, we raise our 2030 CPU TAM to $300B. Unlike chatbots, which remain largely idle between user interactions, always-on agents like Muse can operate continuously, drive a significant increase in inference related compute, memory, and networking consumption across both enterprise and consumer workloads. We raise AMD TP to $800 as the primary CPU beneficiary given, Meta, is one of the largest customers of AMD's server business. We reiterate Buy on NVDA w/$315 TP on GPU demand uplift.

Raising CPU TAM to $300B by 2030 — We update our CPU model post the launch of Muse, Meta's personal AI agent. Fundamentally, we view agentic AI as a potential orders-of-magnitude driver of compute demand relative to traditional chatbots. **We believe the CPU TAM expands from $29B in 2025 to $300B in 2030, or a 60% CAGR.** We still expect AMD to be the key beneficiary of the CPU renaissance and Intel as a secondary beneficiary.

Sizing the GPU Revenue for a Consumer Agent Per User. Meta's Muse is the first consumer AI agent launched at true social-network scale, and its compute footprint is now one of the more consequential unknowns in AI infrastructure. Our base case is that each user requires about 0.0020 GB200-class GPUs. At 100M daily users, our framework implies roughly 200,000–390,000 Blackwell-class GPUs, about 2,800–5,500 NVL72 racks, and approximately $7–19B of one-time Nvidia revenue, including scale-out networking. **On a per-user basis, each incremental million agentic daily users is worth roughly $70–190M to Nvidia if new servers are used.**

Memory: CPUs have a large attach rate of LP, DDR, and SSDs. Micron expects server units to grow in the high teens in CY26/27 and another vector of logic silicon to grow and take advantage of AI trends. Micron recently said that CPUs are creating more of a constraint on DRAM with Meta Muse as an example of how consumers are drawing value from agentic workflows.

Networking: Agentic AI materially increases networking traffic because agents are increasingly accessing infrastructure, creating a new acceleration opportunity using DPUs and Spectrum-X Ethernet.

See Appendix A-1 for Analyst Certification, Important Disclosures and Research Analyst Affiliations.

Data Summary

															Current Fiscal Year		Next Fiscal Year	
						Rating		Short-Term	Target Price						EPS		EPS	
Company	Ticker	Ccy	Price	Mkt Cap (M)	Date & Time	Old	New	View	Old	New	ESPR (%)	Div Yld (%)	ETR (%)	Last Rpt Yr	Old	New	Old	New
Advanced Micro Devices	AMD	US$	631.75	1,031,316	05 Oct 16:00	1	nc	-	575.00	800.00	26.6	0.0	26.6	Dec-25	7.89	nc	18.23	18.84
NVIDIA Corp	NVDA	US$	238.90	5,757,490	05 Oct 16:00	1	nc	-	315.00	nc	31.9	0.0	31.9	Jan-26	9.26	nc	15.17	nc

1 = Buy, 2 = Neutral, 3 = Sell, H = High Risk
ESPR = Expected Share Price Return, ETR = Expected Total Return, nc = no change
Source: Citi Research
^Catalyst Watch

Meta Muse Implications

CPU TAM to $300B by 2030: We update our CPU model following the launch of several personal AI agents, including Muse, Dots, Instinct, GrokBot and GeminiSpark. Fundamentally, we view agentic AI as a potential orders-of-magnitude driver of compute demand relative to traditional chatbots. Unlike chatbots, which remain largely idle between user interactions, AI agents like Muse can operate continuously for hours or even run persistently in the background, autonomously executing multi-step tasks. Moreover, using cybersecurity as a prime example, where always-on agents could generate substantially higher token volumes, driving a significant increase in inference and compute consumption across both enterprise and consumer workloads.

Positive Impact on CPUs & Memory Demand: Unlike standard chatbots, Muse operates continuously as an autonomous agent and draws from virtual machines rather than GPUs. Muse is built upon Meta's Spark model and is designed to be the center of user's daily life. Prior to the explosion of agentic AI, a CPU was primarily limited to two main applications: 1.) traditional workflows and 2.) acting as the head node for AI applications.

In a head node, the CPU merely does the "management" by sending the user request to the GPU and sending the output back to the user. In this context, the GPU is doing heavy lifting such as running matrix multiplications to train LLMs or doing the inferencing when you query a chatbot.

The third new application for CPUs is agentic AI, and this new application has seen explosive growth within the past year as AI has moved from chatbot inferencing to agentic autonomous workflows like Claude Code and Meta Muse. At Citi, we define agentic CPUs as those that are handling the orchestration, reasoning loops, data processing, and security components of autonomous AI agents.

In Agentic AI, CPUs have become the new bottleneck as CPU-to-GPU ratios move from 1:8 in model training, to 1:4 in inferencing, to 1:1 or greater in agentic AI. We believe the ratio will increasingly skew towards CPUs over time as agentic AI scales. We believe the CPU TAM expands from $29B in 2025 to $300B in 2030, or a 60% CAGR.

We expect general purpose CPUs to grow at a 26% CAGR to $64B in 2030, AI head nodes to grow at a 57% CAGR to $78B in 2030, and agentic CPUs to grow at a 247% CAGR to $157B in 2030. We model general purpose CPUs to represent 21% of the TAM in 2030, AI head nodes to represent 26% of the TAM in 2030, and agentic CPUs to represent 52% of the TAM by 2030. For reference, AMD forecasted a TAM of $220B by 2030, and Nvidia forecasted a TAM of $200B. For memory, CPUs have a large attach rate of LP, DDR, and SSDs. Micron expects server units to grow in the high teens in CY26/27 and another vector of logic silicon to grow and take advantage of AI trends. Micron recently said that CPUs are creating more of a constraint on DRAM with Meta Muse as an example of how consumers are drawing value from agentic workflows.

Figure 1. Citi Updated CPU TAM Model

TAM Estimates ($M)	C25	C26E	C27E	C28E	C29E	C30E
General Purpose CPU	$20,544	$27,777	$34,545	$38,989	$48,121	$64,488
% of Total	71%	54%	35%	28%	24%	21%
YoY		35%	24%	13%	23%	34%
AI Headnodes	$8,185	$12,549	$27,856	$36,188	$55,938	$78,086
% of Total	28%	24%	28%	26%	28%	26%
YoY		53%	122%	30%	55%	40%
Agentic CPU Applications	$311	$11,591	$37,637	$62,842	$98,447	$157,426
% of Total	1%	22%	38%	46%	49%	52%
YoY		3627%	225%	67%	57%	60%
Total	**$29,039**	**$51,917**	**$100,038**	**$138,018**	**$202,506**	**$300,000**
YoY		76%	35%	24%	30%	18%

© 2026 Citigroup Inc. No redistribution without Citigroup's written permission.

Source: Citi Research, Mercury Research

We recap pre-quiet period highlights with the AMD IR team:

- **Helios launch execution.** Management is confident where they are with the ramp. Customers have racks and are running code on their racks. There are a few blocking and tackling with ODMs, very well-known are with Sanmina, ZT Systems, Wiwynn, and several other ODMS that are early in ramp. Other ODMs will be participating in a couple of quarters (Q1-Q2 timeframe). There are lot of things they are working through, but they are where they need to be.

- **Supply.** AMD already secured supply to do better than the plan. Includes packaging and CoWoS from TSMC. In an advanced position with substrates. AMD has HBM supply and they have DRAM. They are in a good position in supply. Guided to greater than 100% growth, could more than double from 2026 to 2027. AMD previously guided servers to up 70% YoY, they will ramp Venice very quickly for both 256 agentic parts and 96 core for head nodes. Turin should do well over next few quarters as AMD has strong 3nm and 4nm supply. AI business will more than double YoY.

- **Memory De-specing.** AMD is having some conversations with customers but wouldn't be hard to change from a roadmap perspective.

- **View on fast inference.** In house view, 75% of AI will be served by GPUs is consistent, but TAM has grown. There is a niche of fast inference accelerators that is relatively small today and could be an exciting part of the market. Doing a partnership with Cerebras with their WSE in their cloud. AMD does their head nodes, they do FPGAs in WSE, and AMD are investors in Cerebras. They do plan to do internal chiplets and have plans in ultra-fast inference, Talos provides IP and talent for that goal. AMD will have multiple ways to attack this market.

- **Potential TPU/ASIC capabilities.** AMD has semi-custom GPU configurations like with Meta, they have done separate shipments. They have 40 ASIC programs in company today, these are mostly smaller projects - but do include game consoles with Sony, Microsoft.

- **Future MI500/MI600 designs.** MI455 will be doing scale up networking with UAL, and MI500 will do native silicon ethernet partners. AMD will have option to use optics for scale up at end of MI500 products. Will likely be NPO. Copper and Optics will coexist for several generations.

- **Rev per GW.** AMD hasn't given specifics on these numbers, but have acknowledged the range of $15-$20B. Longer term, they expect revenue per GW

will go higher for AMD and the industry. Revenue per GPU will go up significantly, but the number of racks per GW will come down.

- **Venice Servers.** Venice will be most important server program in history of AMD; it will be a successful and fast launch. Not every partner gets allocation at the beginning of their launch given increased demand. Spikes in demand for 256 core agentic part and 96 core AI head node part.

- **$220B TAM for servers.** Ambition to be over 50% dollar share in TAM, which implies AMD will build a $100B server business.

- **Enterprise mix.** Enterprise is a part of the traditional workflows. Feel good there with AMD investments and go-to-market. There is a little bit of share shift happening there with DMR and their 8-channel competitor. As they move through decade it will be smaller portion of TAM, but it is a good business because it is sticky and high OM.

- **Gross margins.** AMD guided to 56% in 2Q and did better than that and guided to 56% in Q3 which includes 1 month of MI shipments and impacts to DT business. MI455 will become a meaningful part of revenue mix in 4Q, will be majority of AI business in Q4, so margin % likely will come down in Q4 (we estimate ~100 bps or so). In 2027, 55% range will move around depending on quarter. GM $ growth will be significantly faster than opex growth. Positive drivers include servers and embedded, and gaming being weaker will be a tailwind to gross margins. AMD hasn't been margin stacking their HBM. HBM capacity and bandwidth are very important for inference workflows.

- **Client and gaming.** Client will be up 3Q, modeling a sub seasonal in 4Q for the industry. Client held better in 1H for the industry, and AMD likely gained share. Embedded has cleared inventory at distributor and OEM level, Embedded will be up 30% YoY in Q3. Embedded is important for GM and OM. Because DC will more than double, Embedded mix will go down.

Meta Muse: Sizing the GPU Revenue for a Consumer Agent Per User

Meta's Muse is the first consumer AI agent launched at true social-network scale, and its compute footprint is now one of the more consequential unknowns in AI infrastructure. Meta discloses neither usage nor serving architecture. We therefore build a bottom-up estimate of GPU demand for a single Muse daily active user, then scale it. Our base case is that each user requires about **0.0020 GB200-class GPUs (in theory)**, so simplistically speaking one GPU serves roughly 500 users and every million users requires about 2,000 GPUs. Our bull case, which assumes two hours of daily agent activity, raises that to about **0.0039 GPUs per user**, or roughly 250 users per GPU. For context, a chat-style user in our framework needs only 12.5-25% as much GPU capacity. This is the core of our view that agents are a structurally heavier inference workload than chatbots.

Figure 2. MUSE User GPU Estimate

MUSE User GPU Estimate				
Users served by one GPU	users	**498**	**377**	**254**
Users served by one NVL72 rack (72 GPUs)	users	35,867	27,176	18,305
GPUs per 1,000 users	GPUs	2.0	2.6	3.9
GPUs per 1 million users	GPUs	**2,007**	**2,649**	**3,933**
Share of GPUs spent writing (decode)	%	70%	70%	70%

© 2026 Citigroup Inc. No redistribution without Citigroup's written permission.

Source: Citi Research

How one user consumes compute. We model a typical user as running four simple requests a day, such as a quick factual question, plus twelve agentic tasks in our base case, such as finding a restaurant, checking availability against a calendar and drafting an invitation. We estimate an agentic task at about eight model calls. In each call the model reasons briefly, invokes a tool, reads the result and repeats. The model has no memory between calls, so it re-reads the growing conversation each time. A task therefore reads about 104,000 input tokens in total we estimate, even though its final context is only around 21,000. We assume the serving stack caches previously processed text. On that basis, only about 19,000 of those tokens are genuinely new work: the 5,000-token starting context (system instructions, tool definitions, user memory and the request) plus seven tool results of roughly 2,000 tokens each. Output per agentic task is about 4,000 tokens: seven decision steps of about 300 tokens, a 1,000-token final answer and a 900-token allowance for hidden reasoning. Across the day, our base-case user generates about 238,000 new input tokens and 50,000 output tokens.

From tokens to GPUs. We anchor GPU throughput to public benchmarks of DeepSeek R1, the closest disclosed analogue to a frontier mixture-of-experts model, on GB200 NVL72. We assume Muse Spark has around 800B total parameters, of which about 40B are active per token. Reading input (prefill) is compute-bound, and we estimate about 22,500 tokens per second per GPU, consistent with the roughly 26,000 LMSYS has demonstrated. Writing output (decode) is limited by memory bandwidth, not compute, because the GPU must re-read weights and each user's stored context for every token it generates. We assume 2,000 tokens per second per GPU. That sits below short-prompt benchmarks to reflect agent-length contexts. We then apply a 1.5x haircut to both figures to bridge benchmark conditions and a live fleet. On these assumptions, our base-case user consumes about 53 GPU-seconds per day: 37.5 seconds writing and 16 seconds reading. Dividing by the 86,400 seconds in a day, sizing for a 2.5x peak-to-average traffic ratio (Muse is US-only, which concentrates usage) and adding 30% spare capacity yields our 0.0020 GPUs per user. Roughly 70% of that requirement comes from writing rather than reading.

Investment implications. At 100M daily users, our framework implies roughly 200,000–390,000 Blackwell-class GPUs, about 2,800–5,500 NVL72 racks, and approximately $7–19B of one-time Nvidia revenue, including scale-out networking. On a per-user basis, each incremental million agentic daily user is worth roughly $70–190M to Nvidia if new servers are used.

We spoke with the NVIDIA IR team recently and recap key highlights.

- **Agentic AI applications** — Management sees agentic AI as a potentially orders-of-magnitude increase in compute demand versus traditional chatbots. Chatbots are largely dormant between prompts, while agents can run persistently for hours or continuously—particularly in applications like cybersecurity—driving substantially more token generation across both consumer and enterprise workloads.

- **Open vs Closed Models** — NVIDIA continues to expect large open and proprietary model ecosystems to coexist, with open models particularly important for enterprise and sovereign AI. Enterprises can bring open models in-house and customize/fine-tune them on proprietary data without that data leaving the premise, preserving data sovereignty and control. NVIDIA also emphasized open data sets alongside open model weights, as for its own open models, the company releases both the weights and the underlying training data sets, providing greater visibility into how the model was trained and should behave. Management compared the future AI ecosystem to proprietary software coexisting alongside Linux and other open-source infrastructure.

- **Agent Security** — NVIDIA's agent-security architecture adds another potential infrastructure role for BlueField. As detailed in the tech blog, OpenShell provides the open-source policy/guardrail layer, while NVIDIA Sentry runs enforcement and telemetry in silicon on the BlueField DPU, operating out-of-band so the agent cannot manipulate the monitoring layer. This gives enterprises the ability to monitor, control and, if necessary, shut down agent activity.

- **Scale-In / Scale-Context** — Scale-context builds on the previously discussed STX/CMX platforms, with STX providing the rack architecture and CMX using it for large-context-memory workloads. Scale-in is the more incremental opportunity, extending NVIDIA beyond its historical east-west networking focus into north-south traffic such as storage access, security, provisioning and other data-center services. Management argued agentic AI materially increases this traffic because agents—not humans—are increasingly accessing infrastructure, creating a new acceleration opportunity using BlueField 4 and Spectrum-X Ethernet (see technical blog here).

- **Fast Inference / LPX** — LPX is expected to begin shipping late in the current quarter and ramp through Q4 and beyond. NVIDIA still views extreme-interactivity/fast inference as a relatively niche portion of AI infrastructure—roughly ~10% nearer term and potentially ~15–20% longer term—with the majority of throughput adequately served by NVL-class GPU racks. However, against NVIDIA's expectation for $3 - 4T of annual AI capex by decade-end, even a modest percentage represents a very large absolute opportunity.

- **Power Availability** — NVIDIA acknowledged that power remains tight and should stay tight but does not currently view it as a constraint on its overall growth outlook. Management argued that power contracting is inherently multi-year, forcing customers and suppliers to plan well ahead, and believes the issue is ultimately "surmountable." Some US regions may not support incremental data-center builds, but globally there is sufficient power and workloads can migrate toward regions where power is available. Management highlighted several factors that might be underappreciated in the power supply: off-grid generation, generation-to-generation energy-efficiency gains and scale-across, which can link geographically separate data centers into a larger logical AI factory. NVIDIA also argued its global footprint and broad customer/project pipeline allow infrastructure to be redirected if an individual project is delayed, making the company relatively resilient in a constrained environment.

Raising TP on AMD to $800

Following the release of Meta Muse and our subsequent CPU TAM raise, we are raising estimates on AMD given our belief that AMD will be the primary beneficiary of the CPU renaissance and our belief that Meta is one of the largest customers of AMD's server business. We raise our TP to $800 from $575 and maintain AMD as our favorite growth compute semis pick in 2026.

Valuation - We lift our TP for AMD to $800 from $575. We value AMD using SOTP analysis with data center GPU/CPU biz at $370/$325 (35/35x 2028 PE), client $32 (15x), gaming $2 (6x), embedded $30 (20x), and net cash per share ~$41. Our target PE multiples are in-line with respective peer group averages. Our prior TP on AMD of $575 used SOTP analysis with data center GPU/CPU biz at $267/$207 (27/25x 2028 PE), client $31 (15x), gaming $2 (6x), embedded $29 (20x), and net cash per share ~$40.

Figure 3. AMD SOTP Valuation

SOTP	2026E	2027E	2028E	P/E
Data Center GPU	$51	$184	$370	35
Data Center CPU	$128	$246	$325	35
Client	$20	$25	$32	15
Gaming	$1	$1	$2	6
Embedded (Xilinx)	$18	$25	$30	20
Net Cash Per Share	$7	$17	$41	
Total			**$800**	

© 2026 Citigroup Inc. No redistribution without Citigroup's written permission.

Source: Citi Research, FactSet

Figure 4. Citi CY26/27/28 estimates are 3%/18%/28% above Street

AMD	Old	New		Old	New	
FY end: December	**2027E**	**2027E**	**Delta**	**2028E**	**2028E**	**Delta**
Total Revenue	**94,500.0**	**94,500.0**	**0%**	**131,000.0**	**133,000.0**	**2%**
% Change Y/Y	84.3%	84.3%		38.6%	40.7%	
% Change Q/Q						
3-year Seasonality	32.0%	32.0%		55.6%	55.6%	
5-year Seasonality	27.2%	27.2%		35.3%	35.3%	
Gross Margins	**54.9%**	**54.9%**	**0 bps**	**56.5%**	**56.5%**	**0 bps**
SG&A	6,825.0	6,825.0	0.0%	7,459.0	7,578.0	1.6%
% Total Revenue	7.2%	7.2%	0.0%	5.7%	5.7%	0.1%
R&D	12,780.0	12,780.0	0.0%	14,357.0	14,587.0	1.6%
% Total Revenue	13.5%	13.5%	0.0%	11.0%	11.0%	0.1%
Operating Income	**32272.5**	**32272.5**	**0.0%**	**52244.0**	**53025.0**	**1.5%**
% Total Revenue	34.2%	34.2%	0%	39.9%	39.9%	0%
Tax Rate	13.0%	13.0%	0 bps	13.0%	13.0%	0 bps
Net Income	28,285.9	28,285.9	**0%**	45,661.1	46,340.6	**1%**
% Total Revenue	29.9%	29.9%	**0 bps**	34.9%	34.8%	**-1 bps**
Diluted Shares	1,720.0	1,720.0	**0%**	1,750.0	1,750.0	**0%**
Proforma EPS - ex-options	**$18.84**	**$18.84**	**0%**	**$29.36**	**$29.79**	**1%**
% Change Y/Y	138.9%	138.9%		55.8%	58.1%	
% Change Q/Q						

© 2026 Citigroup Inc. No redistribution without Citigroup's written permission.

Source: Citi Research

Companies Mentioned:

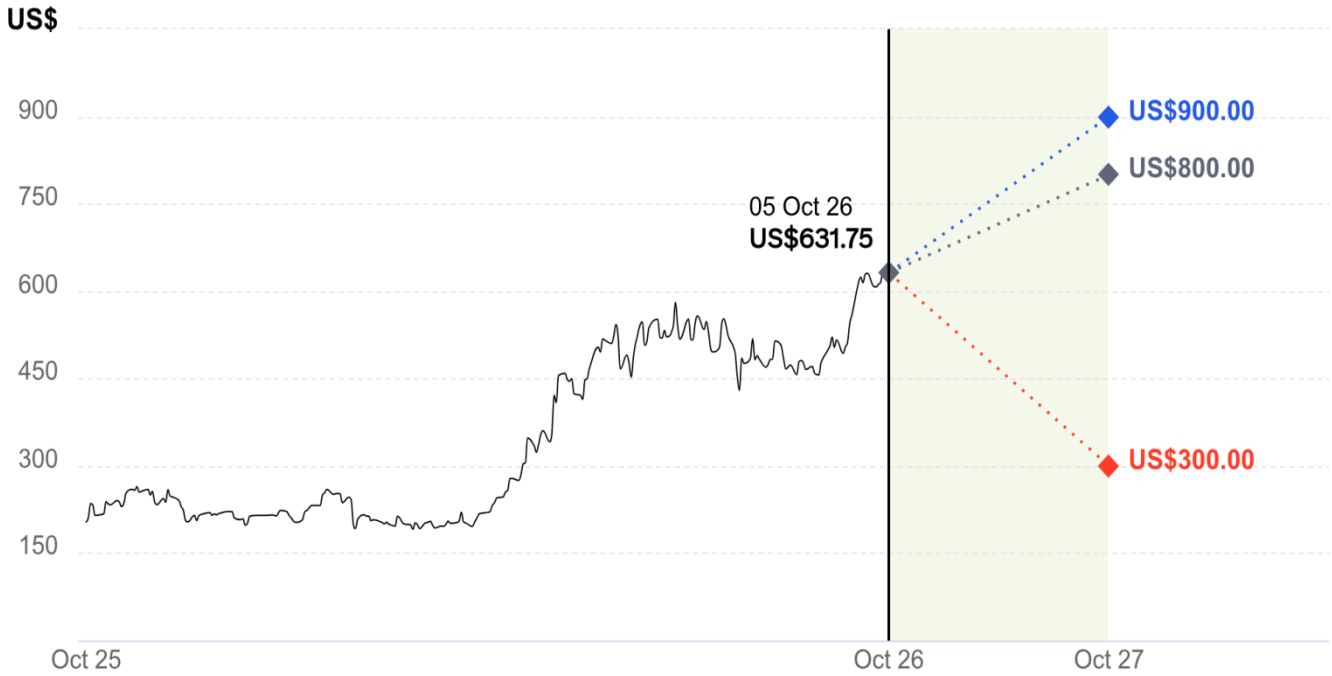
Cerebras Systems Inc (CBRS.O; US$181.55; 1; 05 Oct 26; 16:00) | Intel Corp (INTC.O; US$116.19; 1; 05 Oct 26; 16:00) | Meta Platforms Inc (META.O; US$741.9; 1; 05 Oct 26; 16:00) | Micron Technology Inc (MU.O; US$1063.96; 1; 05 Oct 26; 16:00) | Microsoft Corp. (MSFT.O; US$525.18; 1; 05 Oct 26; 16:00) | TSMC (2330.TW; NT$2575.0; 1; 05 Oct 26; 13:30) | Wiwynn (6669.TW; NT$2150.0; 1; 05 Oct 26; 13:30)

Bull/Bear: Advanced Micro Devices (AMD.O)

US$ 900.00 ▲42% Upside

US$ 800.00 ▲27% Upside

US$ 300.00 ▼53% Downside



Spread 95pp
Current Price and expected returns (upside/downside) as of 05 Oct 2026

BULL Assumptions

- Stronger than expected AI GPU ramp
- Better than expected AI GPU margins

BASE Assumptions

- SOTP analysis

BEAR Assumptions

- Aggressive pricing environment with Intel and Nvidia
- AI GPU ramp fails to materialize

Advanced Micro Devices

Company description

AMD is a fabless semiconductor company that specializes in PC and server microprocessors, graphics processing units (GPUs), semi-custom processors, and embedded processors. AMD is the #2 supplier of microprocessors for desktop PCs, notebook PCs, and servers with approximately 30% total market share. Based in Sunnyvale, CA, AMD has roughly 31,000 employees worldwide. AMD derives approximately 85% of revenue from the Data Center and Client segments and 15% of revenue from the Embedded and Gaming segments. AMD's largest customers include gaming console vendors, PC OEMs, hyperscalers, and LLM companies. AMD competes with Intel in the microprocessor segment and NVIDIA in the AI GPU segment. As a fabless semiconductor company, AMD outsources its wafer fabrication to third party foundries. The company has internal back-end assembly and test operations in Malaysia and China.

Investment strategy

We are Buy rated on AMD given increasing conviction on AMD's AI business and server CPU tailwinds. Namely, we believe Meta will be a significantly larger customer of AMD's AI products than the street is expecting.

Valuation

Our TP on AMD is $800. We value AMD using SOTP analysis with data center GPU/CPU biz at $370/$325 (35/35x 2028 PE), client $32 (15x), gaming $2 (6x), embedded $30 (20x), and net cash per share ~$41. Our target PE multiples are in-line with respective peer group averages.

Risks

Competition: AMD competes directly with Intel in the microprocessor market. Consequently, any fluctuations in market share between AMD and Intel could result in risk to our estimates. AMD competes directly with NVIDIA in the graphics and AI GPU market. Consequently, any fluctuations in market share between AMD and NVIDIA could result in risk to our estimates.

PC End Market: AMD derives roughly 30% of sales from the PC industry, which is highly dependent on IT spending. Therefore, any major uptick/downtick in IT spending could result in upside/downside to our estimates and rating for AMD.

Valuation: While AMD is currently trading slightly above its historical range, revisions to consensus revenue estimates could result in high volatility to AMD stock.

Tender Offer: Any public or private tender offer for all or part of AMD's business could drive the stock higher and cause us to re-assess our rating.

If the impact on the company from any of these factors proves to be greater than we anticipate, the stock will likely have difficulty achieving our target

price. However, should they be less than anticipated, the stock could trade above our target price.

NVIDIA Corp

Company description

Nvidia is a leader in graphics processing units or GPUs. It is engaged in creating graphics and networking chips, which are used in gaming, data centers, and automotive end markets. Nvidia was founded by Jen-Hsun Huang (current CEO) and partners in January 1993 and is headquartered in Santa Clara, CA.

Investment strategy

We like Buy-rated NVDA on secular AI growth opportunities.

Valuation

Our price target for NVDA of $315 is based on ~21x P/E on C27E earnings power. Our 21x P/E multiple is in-line with 1 year average.

Risks

Downside risks to the attainment of our target price include: 1) competition on gaming could drive the stock lower if Nvidia loses market share; 2) slower-than-expected adoption of new platforms can drive lower data center and gaming sales; 3) lumpiness in auto and data center markets can add volatility to the stock/multiple; and 4) cryptomining impact on gaming sales.