Advanced Micro Devices Inc. (AMD): Key takeaways from Advancing AI 2026

James Schneider, Ph.D.
+1(212)357-2929 |
jim.schneider@gs.com
Goldman Sachs & Co. LLC

Anmol Makkar
+1(212)357-1366 |
anmol.makkar@gs.com
Goldman Sachs & Co. LLC

Luya You
+1(212)902-5297 | luya.you@gs.com
Goldman Sachs & Co. LLC

Khalil Fenina
+1(212)357-6392 |
khalil.fenina@gs.com
Goldman Sachs & Co. LLC

Key stock takeaways: We attended CEO Dr. Lisa Su's keynote at the company's Advancing AI event in San Francisco. Key highlights were: (1) AMD announced a strategic partnership with Anthropic for 2GW of Helios racks beginning in 1H27 (along with a $5 bn equity investment from AMD); (2) AMD expanded its partnership with Microsoft to deploy Helios racks, Venice CPUs and networking infrastructure across Azure workloads; (3) AMD raised its total compute/accelerator/server CPU TAM to $2 tn/$1.4 tn/$220 bn by 2030, corresponding to ~50% growth in server CPUs; (4) The company launched ROCm.ai to provide developers with an AI-driven software development experience. We believe investors were positioned constructively heading into the event, and based on our conversations we believe the series of announcements this week met these elevated expectations ahead of the company's earnings report on August 4. We reiterate our Buy rating on AMD, as we see the company well positioned to capture an increasing share of overall compute market, driven by an already strong presence in the CPU market and a broadening aperture of key customers in the AI accelerator market.

Key takeaways include:

- **Partnership with Anthropic:** AMD and Anthropic announced a strategic partnership under which Anthropic will deploy 2 GWs of Instinct MI450 GPUs in Helios rack-scale AI systems, with the initial 1 GW expected to begin deployment in 1H27. The partnership extends beyond hardware, with the companies collaborating to optimize Claude workloads on AMD GPUs, accelerate development of AMD's ROCm software platform, and expand Claude's use across AMD's engineering and product teams. AMD also plans to make a strategic equity investment of up to $5 bn in Anthropic.
- **Partnership with Microsoft:** Microsoft and AMD announced an expanded strategic partnership, with Microsoft planning to deploy Helios racks (with shipments to Microsoft beginning in 2H26), CPUs, networking and software on Microsoft Azure. The Azure deployment will use Helios for inference workloads spanning frontier models, as well as Microsoft's own AI services and other Azure customer applications. Azure will also introduce two new virtual machines powered by AMD's next-generation 2nm Venice CPUs, while expanding the use of AMD Pensando DPUs in Azure networking services.
- **Increased Datacenter and CPU TAM:** Agentic AI is driving a step-function

change in computing demand, requiring both accelerators and CPUs to orchestrate increasingly complex workflows. As a result, the AI infrastructure market is expanding rapidly, with the datacenter AI accelerator TAM projected to grow from $200 bn to $1.4 tn in 2030 (40% CAGR), datacenter CPU TAM from $26 bn to $220 bn (50% CAGR), and AMD's total compute TAM from $365 bn to $2 tn by 2030. AMD also expects to capture 50% of the datacenter CPU market by 2030.

- **Helios launch:** AMD unveiled its next-generation Helios AI rackscale platform, built on CDNA 5 architecture with GPUs featuring up to 40 PFLOPS (FP4), 20 PFLOPS (FP8), 432 GB of HBM4 memory, and 23.3 TB/s memory bandwidth. The system combines 75 GPUs connected via UALink over Ethernet, a 96-core EPYC CPU, Salina DPU, Volcano 800G AI NIC, and six networking trays to deliver higher-scale AI performance. AMD stated that Helios is already in full production, with shipments beginning at the end of 3Q, and volume ramp in 4Q.
- **Partnership with Cerebras**: AMD and Cerebras have partnered to combine AMD's Helios systems with Cerebras' Wafer-Scale Engine, creating a high-performance AI inference solution that delivers ultra-low latency, higher efficiency, and up to 5x better tokens-per-watt. The solution is expected to launch via Cerebras Cloud in 2H26.
- **ROCm.ai**: AMD launched ROCm.ai, which integrates with tools such as Cursor, Claude, Codex and Gemini to provide developers with an AI-driven software development experience. AMD also introduced Hyperloom, an optimization layer that automatically tunes system configurations and optimizes end-to-end AI workloads. AMD stated that over 14,000 models have been optimized through Hyperloom, delivering an average 3.3X performance improvement compared with ROCm 7.

Price target. We are Buy rated on AMD. Our 12-month target price of $640 is based on a 32X P/E multiple applied to our normalized EPS estimate of $20.00. Key downside risks include: (1) slower-than-expected adoption of agentic AI; (2) slower-than-anticipated deployment of AMD GPUs related to the Meta deal; (3) share erosion of x86 architecture in enterprise AI and (4) lack of operating leverage.

AMD	12m Price Target: **$640.00**	Price: **$539.69**	Upside: **18.6%**

Buy

Market cap: $890.5bn
Enterprise value: $877.8bn
3m ADTV: $16.2bn
United States
Americas Semiconductors & IT Services
M&A Rank: 3

GS Forecast

	12/25	12/26E	12/27E	12/28E
Revenue ($ mn)	34,639.0	50,572.5	86,013.3	109,024.8
EBITDA ($ mn)	6,984.0	13,179.7	26,559.4	36,357.5
EBIT ($ mn)	6,130.0	12,344.7	25,715.4	35,493.5
EPS ($)	2.64	6.20	13.20	17.94
P/E (X)	57.7	87.1	40.9	30.1
EV/EBITDA (X)	34.3	66.2	32.8	23.6
FCF yield (%)	2.7	0.5	0.7	2.6
Dividend yield (%)	–	–	–	–
Net debt/EBITDA (X)	(1.2)	(1.0)	(1.2)	(2.0)

	3/26	6/26E	9/26E	12/26E
EPS ($)	0.84	1.32	1.53	2.51

Source: Company data, Goldman Sachs Research estimates, FactSet. Price as of 23 Jul 2026 close.