

Advanced Micro Devices

The AI datacenter's credible second source: EPYC's server war nearly won, a rack-scale AI system in production, and a multiple that demands flawless delivery

September 28, 2026 · NASDAQ: AMD

Reference price: \$630.63 (September 25, 2026 close)

72/100

DalalBytes score

Positive

Verdict

\$750

FY2029 base target*

\$1.03T

Market cap

Headline verdict

Positive, priced for perfection.

AMD is the only company in the world fielding both a full x86 CPU line and a credible NVIDIA-class AI GPU line, and it has just crossed a \$1 trillion market value on the strength of that position. The data center business grew 107% year over year in Q2 2026, EPYC owns 46% of server CPU revenue dollars, and Meta, OpenAI, and Anthropic have committed to up to 14 gigawatts of Instinct deployments. The stock is up about 187% year to date. The verdict is positive because the fundamentals behind the move are real: rack-scale systems in production, a process lead on 2nm server CPUs, and a software stack that has finally reached functional parity for mainstream inference. But the multiple, roughly 47 times FY2027 consensus earnings and 25 times trailing sales, prices in flawless multi-year execution of customer commitments that are still "up to" deals, not booked backlog. DalalBytes score: 72/100. FY2029 estimates (DalalBytes): bear \$300, base \$750, bull \$1,050.

\$630.63

REFERENCE PRICE

\$154.78-\$639

52-WK RANGE

160x / 47x

TTM GAAP / FY27 FWD P/E

72/100

SCORE

\$750

BASE TARGET*

+19%

BASE UPSIDE*

* FY2029 targets are DalalBytes estimates, not facts. See the Valuation section for the full bear/base/bull framework.

Business overview

Advanced Micro Devices is a fabless semiconductor company selling into three reporting segments: Data Center (EPYC server CPUs, Instinct AI accelerators, Pensando networking), Client and Gaming (Ryzen processors, Radeon graphics, semi-custom console APUs), and Embedded (Xilinx FPGAs and adaptive SoCs). Founded in 1969, the company spent most of its life as Intel's also-ran. That ended under Lisa Su, who took over in October 2014 when the company was worth about \$2 billion and spent a decade rebuilding it around the Zen CPU architecture, the Xilinx acquisition, and the datacenter pivot. On September 21, 2026, AMD joined the \$1 trillion club, the fourth US chipmaker to do so after NVIDIA, Broadcom, and Micron.

Why it matters: the AI compute buildout has a single point of failure. NVIDIA holds roughly two-thirds of AI datacenter share, and every hyperscaler, sovereign, and neocloud building AI infrastructure wants a second source. AMD is the industry's designated second source, the only merchant supplier with a complete stack of its own: CPU (EPYC), GPU (Instinct), networking (Pensando), and now rack-scale integration (ZT Systems). The next two years will test whether "designated" becomes "earned."

Technology: the full-stack stack


EPYC SERVER CPUS

EPYC is AMD's most underappreciated asset in an AI narrative. Mercury Research reported that in Q1 2026 AMD captured 46.2% of total x86 server CPU revenue, an all-time high, on about 33% unit share. Selling a third of the units and taking nearly half the money means AMD owns the premium high-core-count sockets while Intel keeps the volume. By Q2 2026, AMD's unit share reached 34.5%, and head-to-head against Xeon SP, AMD stood at 46.4%, effectively parity. Q1 2026 was the first quarter AMD's Data Center revenue (\$5.8 billion) exceeded Intel's datacenter AI group (\$5.1 billion). AMD doubled its server CPU TAM forecast to \$120 billion by 2030.

Venice, the sixth-generation EPYC built on the Zen 6 architecture, extends this lead. AMD announced on May 21, 2026 that Venice entered production ramp in Taiwan on TSMC's 2nm (N2) process, the first high-performance computing product in the industry to ramp on N2, with production planned to expand to TSMC Arizona. At Advancing AI 2026 on July 22, 2026, AMD formally launched the family: up to 256 cores and 512 threads on the new SP7 socket, 16 DDR5 memory channels delivering roughly 1.6 TB/s of per-socket bandwidth (up from 614 GB/s on Turin), PCIe Gen 6, and a fifth-generation Infinity Fabric. AMD claims about 70% throughput uplift over Turin (company figures; independent verification pending). The process fact is what matters: Intel has no 2nm-class server

product through at least the end of 2027. The follow-on, Verano, adds native LPDDR5X memory and 112 Gbps CPU-to-GPU bandwidth for rack-scale AI chassis, with Meta as the lead customer. EPYC lead times reportedly stretched past 30 weeks in 2026 with price increases of 10 to 35% on some parts, scarcity pricing power that preceded the September price-increase reports behind the final leg of the trillion-dollar run.

x86 server CPU revenue share, Q1 2026 (AMD all-time high)




Mercury Research via press coverage, Q1 2026. AMD unit share was ~33%.

INSTINCT ACCELERATORS

The Instinct roadmap is now a credible annual cadence. The MI300X/MI325X (CDNA 3, shipped) are deployed at Microsoft, Meta, and hyperscaler fleets and in the El Capitan supercomputer. The MI350X/MI355X (CDNA 4, shipped mid-2025) run on TSMC 3nm with 288 GB of HBM3E at 8 TB/s and about 20 PFLOPS of FP4 on the liquid-cooled MI355X. At MLPerf Inference v6.1, an eight-GPU MI355X led submitted results against NVIDIA's B200 and B300 on GPT-OSS-120B, and a 72-GPU MI355X cluster led NVIDIA's GB200 submission: peer-reviewed, not vendor marketing.

The MI455X "Altair" and the MI400 family (CDNA 5, shipments starting end of Q3 2026, ramping into 2027) carry 40 PFLOPS of FP4, 20 PFLOPS of FP8, 432 GB of HBM4 at AMD's disclosed 19.6 TB/s of bandwidth, about 320 billion transistors on 2nm/3nm processes,

and 300 GB/s scale-up fabric links. The memory-capacity edge is the structural point the market treats as a marketing bullet: 432 GB per GPU is 1.5 times NVIDIA's Rubin and 2.25 times B200-class parts. For inference, memory capacity is the binding constraint on how large a model fits in one node; more memory per GPU means fewer nodes, fewer GPUs, and less interconnect complexity per served model, which compounds into total-cost-of-ownership savings.



MI455X: AMD-disclosed 432 GB HBM4. Rubin/B200: public specs. More memory per GPU = fewer nodes per model.

HELIOS: THE RACK IS THE PRODUCT

Helios is the ZT Systems acquisition made tangible: a liquid-cooled double-wide rack with 72 MI455X accelerators, sixth-generation EPYC Venice CPUs, and Pensando Vulcano networking. Rack totals: about 31 TB of HBM4, about 2.9 exaFLOPS of FP4 compute, 260 TB/s of scale-up bandwidth, 43 TB/s of scale-out bandwidth. In production, with shipments beginning at the end of Q3 2026. Hyperscalers buy racks, not cards, and until ZT Systems AMD was a component vendor competing against NVIDIA's NVL72 rack with a bag of parts. Helios is the only credible non-NVIDIA rack-scale alternative, and its all-Ethernet, open-standards networking lets operators reuse existing tooling instead of adopting proprietary NVLink. Microsoft plans to deploy Helios across Azure AI services starting in the second half of 2026. HPE added Helios to its portfolio in December 2025. Oracle is deploying 50,000 MI400-class GPUs starting Q3 2026 in a deal reported at about \$2 billion.

ROCm VERSUS CUDA

The honest 2026 assessment: ROCm 7.x reaches 90 to 95% of H100 throughput for standard LLM inference workloads (PyTorch, vLLM, SGLang) on MI300X and MI355X. PyTorch support is official, Triton is a first-class target, day-zero model support is established, and all three major clouds offer AMD GPU instances. Meta's 6 GW production


deployment is the strongest enterprise-validation signal the software stack has ever had, and Anthropic's equity arrangement directly funds Claude-driven acceleration of ROCm development. The remaining gaps are real: no ROCm equivalent of TensorRT-LLM or FlashAttention 3 (a 20 to 40% CUDA lead in custom-kernel workloads), and NVIDIA maintains an order of magnitude more internal software CI cluster capacity. The structural shift favoring AMD: inference, not training, is where the market's growth is moving (Bloomberg Intelligence expects inference to be roughly double training by 2032), and the inference abstraction layer increasingly hides the CUDA versus ROCm difference from the end user. Training at hyperscale remains CUDA's fortress; inference at scale is genuinely contestable.

PENSANDO, EMBEDDED, SEMI-CUSTOM


Pensando (\$1.9 billion, 2022) completes the datacenter silicon portfolio: the Vulcano 800 AI NIC (800 Gbps, UALink and Ultra Ethernet, positioned against NVIDIA's ConnectX-8), the Pollara 400 UEC-ready NIC, and the Salina 400G DPU handling front-end networking and KV-cache management inside Helios. The payoff shows up in rack-scale deployments rather than standalone revenue. Embedded (Xilinx) hit \$977 million in Q2 2026, up 19% year over year, its strongest growth in over three years, with more than \$18 billion of record new design wins cited on the Q2 call; Versal AI Edge Gen 2 launched in September 2026 for automotive and robotics. Semi-custom consoles are late cycle (Gaming revenue fell 31% in Q2), but next-generation consoles are rumored to stay AMD-powered on Zen 6 and RDNA 5 for a 2027 refresh: unconfirmed, and a free call option the market prices at roughly zero.

Financials: the data center engine, the multiple problem

Data Center went from 26% of revenue in FY2022 to 48% in FY2025 to 58% in Q2 2026. Q2 2026 revenue was \$11.5 billion, up 50% year over year, with Data Center at \$6.7 billion (up 107%), Client at \$3.1 billion (up 23%), Gaming at \$0.8 billion (down 31%), and Embedded at \$1.0 billion (up 19%). Q3 2026 guidance is about \$13.0 billion (plus or minus \$300 million), up about 41% year over year, with strong double-digit sequential growth in Data Center and Embedded. Q2 Data Center segment operating margin was 31.3%, up from 27.7% in Q1; Embedded carries a 39.5% segment margin. The mix shift is the structural margin driver.




Company-reported revenue. Q2 2026 annualized is 4x the reported quarter, not guidance.



Company-reported segment revenue. YoY growth as reported by AMD.

Non-GAAP gross margin moved from 50 to 53% across 2022 to 2024 to 55% in Q1 2026 and 56% in Q2 2026. Q2 GAAP gross margin was 54%, GAAP EPS \$1.38 (up 156%), non-GAAP EPS \$1.66 (up 246%, above consensus of about \$1.62). Two caveats: the Q2 2025 comparison was polluted by roughly \$800 million of inventory charges tied to the MI308 China export restrictions, and the GAAP-to-non-GAAP gap remains dominated by about \$544 million per quarter of Xilinx amortization plus \$503 million of stock-based compensation. The stock fell about 9% after the Q2 print partly because Q3 non-GAAP gross margin guidance was flat at 56%: the market wanted the mix story to keep compounding, and management would not promise it.




Company-reported. Q1 2026 GAAP operating margin not split out in the materials reviewed.

AMD has never disclosed a discrete Instinct revenue figure. Analyst estimates cited in media put 2026 AI-GPU revenue at \$8 billion to \$12 billion. These are analyst estimates, not company guidance, and the refusal to break out the number is itself a data point: the company wants the optionality of the narrative without the accountability of the number.

TTM free cash flow is about \$8.4 billion, up 108% year over year (media-computed; AMD does not report a full-year FCF figure). Q1 2026 set a record at \$2.57 billion, but Q2 2026 fell 39% sequentially to \$1.56 billion as capex more than doubled to \$808 million tied to AI infrastructure scaling. H1 2026 operating cash flow was \$5.3 billion against \$1.2 billion of capex, implying about \$4.1 billion of H1 free cash flow. Inventories rose to \$8.5 billion and accounts payable jumped 78% to \$5.4 billion. The balance sheet is a fortress with a construction site next to it: \$13.1 billion of liquidity against \$3.25 billion of debt at June 27, 2026 (about \$9.9 billion net cash), and after the quarter AMD issued \$4.75 billion of senior notes at 4.6% to 5.5% coupons, taking pro forma gross debt to about \$8 billion, still comfortably net cash. No dividend. Roughly \$10 billion of buyback authorization, but zero open-market repurchases in Q2 2026 while the stock tripled: growth capex over shareholder return, at the top of the cycle.

Free cash flow: record TTM, but Q2 fell 39% sequentially on capex



Company-reported quarterly FCF; TTM ~\$8.41Bmedia – computed. Q2 capex doubled to \$808M.


Contingent dilution is the number the bulls underweight. The OpenAI and Meta warrants cover 160 million shares each at a \$0.01 exercise price, vesting only on gigawatt-scale purchase milestones plus share-price targets (the final tranche requires about \$600, which the stock has now reached). None had vested as of June 27, 2026. Full vesting of both would expand the share count by about 19.6%. Add AMD's commitment of up to \$5 billion of equity investment in Anthropic, funded through fiscal 2028, and the pattern is clear: AMD is spending real capital to buy its demand commitments.

Valuation: the crux

160x	~47x	25x
TRAILING GAAP P/E	FY27 CONSENSUS FWD P/E	TRAILING P/S
124.5x	0.63	\$610
TRAILING P/FCF	PEG (MEDIA-CITED)	MEDIAN ANALYST TARGET

At \$630.63 AMD trades at 160 times trailing GAAP earnings, about 95 to 100 times non-GAAP trailing, about 47 times FY2027 consensus non-GAAP EPS of \$13.47, 25 times trailing sales, and 124.5 times trailing free cash flow, against NVIDIA at 43x and

Broadcom at 43.8x. The stock sits just above the median analyst target of about \$610 but below several fresh September targets (BofA \$720, UBS \$670, Stifel \$635, Goldman Sachs \$640, Raymond James \$641, Piper Sandler \$600). The target spread, \$365 to \$1,250, is one of the widest in large-cap tech: nobody is sure how much of the AI commitments convert to revenue.



Financecharts, September 2026. AMD's 124.5x is ~3x NVIDIA and Broadcom.

The valuation debate, stated fairly. Bulls: FY2026 consensus non-GAAP EPS of \$6.44 doubling to \$13.47 in FY2027 collapses the forward multiple fast, and management's \$20 annual EPS target within five years makes 47x forward look like growth pricing, not bubble pricing. Bears: trailing GAAP P/E of 160x and P/FCF of 124.5x are the reality today, and the earnings ramp requires 12-plus gigawatts of deployments, HBM supply that is sold out industry-wide, and ROCm execution, all at once. The honest middle: AMD is priced as if the base case is the floor. Any stumble in the MI450/Helios ramp reprices the stock through the \$365 low-end analyst targets, not the \$610 median.

Strategy: buying the second-source position

Every major AMD acquisition since 2022 targets the same gap: the difference between selling chips and selling systems. Xilinx (\$49 billion, all-stock, closed February 2022) brought FPGAs, adaptive SoCs, and high-margin embedded revenue; the integration worked, though the price was full. Pensando (\$1.9 billion, 2022) completed the CPU plus GPU plus networking portfolio, with the payoff inside Helios. Silo AI (\$665 million cash, July 2024) added about 300 AI scientists and engineers feeding the AMD AI Group. ZT Systems (about \$4.9 billion, closed March 2025) may be the deal of the decade: AMD kept roughly 1,000 system design engineers and sold the manufacturing arm to Sanmina for about \$3 billion, cutting the effective cost roughly in half while avoiding channel

conflict, and the team built Helios, the delivery vehicle for the Meta and OpenAI commitments. The 2024 to 2026 AI software spree (Mipsology, Nod.ai, Enosemi, Brium, the Untether AI team, about \$36 million for the 2025 batch) targets the software and interconnect gap with NVIDIA rather than more silicon.

THE THREE ANCHOR COMMITMENTS

Customer	Announced	Commitment	Economics
Meta	Feb 2026	Up to 6 GW Instinct; custom inference-optimized MI450; first GW H2 2026 on Helios + Venice	160M-share warrant at \$0.01; WSJ values the deal over \$100B
OpenAI	Oct 2025	6 GW across Instinct generations; first GW MI450 H2 2026; tied to Stargate	160M-share warrant at \$0.01; AMD expects "tens of billions" in revenue over time
Anthropic	Jul 2026	Up to 2 GW MI450-series in Helios; first GW H1 2027	AMD committed up to \$5B equity investment, funded through FY2028

The honest framing, which the market's enthusiasm skips: these are "up to" commitments, not booked backlog. Citi's translation of roughly \$15 billion of AMD revenue per gigawatt is an analyst model, not disclosed economics. AMD guaranteed a \$300 million Goldman Sachs loan to the neocloud Crusoe for chip purchases, a customer-financing structure that mirrors NVIDIA's circular arrangements. And Anthropic signed a separate six-year, \$45 billion lease for NVIDIA-powered capacity in 2026: AI labs multi-source, and paper commitments are not exclusivity. AMD is buying demand, not just winning it. The warrants' \$600 final-tranche threshold, now reached, means the dilution incentive is live.



Company-announced 'up to' figures. Meta and OpenAI include 160M-share warrants at \$0.01.

Competition. NVIDIA holds roughly 68% of AI datacenter share (down from 92% in 2024); training at hyperscale remains CUDA's fortress, inference at scale is the contested ground. Intel has effectively lost the server CPU war: EPYC's 46.2% of x86 server revenue dollars, the first quarter of AMD Data Center revenue exceeding Intel's DCAI, and a 2nm process lead Intel cannot match through 2027. Broadcom and Marvell play a different game, custom AI silicon co-designed with hyperscalers (Broadcom reports a \$73 billion AI backlog); the irony is that Meta's custom MI450 is a merchant-custom hybrid that looks like the Broadcom model, with AMD as the supplier. Hyperscaler custom silicon is both threat and validation: every gigawatt built in-house is one AMD never bids on, but the scale of the custom programs confirms the inference demand AMD is targeting.

Supply chain and China are the two physical risks. Each MI455X carries 432 GB of HBM4, and 2026 HBM supply is fully booked industry-wide. AMD named Samsung as its primary HBM4 supplier (MOU at Pyeongtaek, covering MI455X HBM4, Venice DDR5, and Helios), with Samsung potentially fabricating future processors as TSMC diversification. But Samsung's HBM4 was roughly a year behind SK Hynix, where NVIDIA holds priority, and TSMC CoWoS allocation tension constrains any 12-gigawatt ramp. China is about 20% of revenue, mostly PC, gaming, and EPYC: the April 2025 license requirement killed MI308 sales to China (\$800 million of Q2 2025 charges), February 2026 brought licenses for MI325 shipments to certain customers, MI450-class parts remain off-limits, and the unpublished 15% revenue-share "expectation" voiced by US officials in August 2025 remains a litigation-flagged overhang.

Leadership is the pillar with the least uncertainty. Lisa Su took over in October 2014 with the company worth about \$2 billion; twelve years later it is worth \$1 trillion, through the Zen bet, the datacenter pivot, the Xilinx acquisition, and a repeated under-promise and over-deliver pattern. The board's March 2026 performance award (up to about \$450

million if the stock reaches \$600 by March 2031, requiring her to remain in a full-time leadership role) signals no succession plan because there is no succession need yet. Key lieutenants: Forrest Norrod (Data Center Solutions), Jean Hu (CFO), Mark Papermaster (CTO), Vamsi Boppana (AI Group), Salil Raje (Adaptive and Embedded).

Risks

- Customer concentration: two customers anchor essentially the entire forward AI-GPU story, and both warrants dilute about 19.6% if fully vested.
- Execution on the H2 2026 ramp: Helios shipments beginning end of Q3 2026 and the first Meta and OpenAI gigawatts must land on schedule; any slip reprices the forward earnings the multiple depends on.
- HBM and CoWoS supply: sold-out industry HBM, Samsung as primary on the flagship, TSMC packaging shared with NVIDIA and Broadcom.
- ROCm adoption: the single biggest technical risk. If inference abstraction stalls or NVIDIA locks workloads via libraries, the hardware advantage does not convert.
- China and export controls: 20% of revenue exposed to license roulette, with escalation risk up the stack.
- Valuation: the price assumes the base case. The \$365 low-end analyst target is the map of what a stumble looks like.


What the market may be mispricing

- The server CPU engine is underweighted in the AI narrative: EPYC at 46% of server CPU revenue dollars with 30-week lead times and price increases is a scarcity asset the market prices as a footnote to the GPU lottery ticket.
- The memory-capacity edge compounds: 432 GB per GPU is not a spec bullet; it is fewer nodes per model, less interconnect, lower TCO, the exact arithmetic inference buyers run.
- ZT Systems may be the deal of the decade: roughly \$1.9 billion net of the Sanmina sale bought the rack-integration capability that made the Meta and OpenAI deals physically deliverable.
- The warrants cut both ways: they align the two anchor customers with AMD's success and put a \$600 floor of incentive under the stock; they also represent 320 million shares of contingent dilution and evidence that AMD paid for its demand.
- Samsung as second foundry and memory source reduces single-vendor TSMC and

HBM risk in a way the market has not priced, though it bets the flagship on the HBM laggard.

Valuation and targets: FY2029

Reference: \$630.63, September 25, 2026. Bear, \$300: the MI450/Helios ramp stumbles or slips into 2027, ROCm adoption disappoints, China restrictions escalate, and the multiple compresses to growth-stock norms; roughly \$12 of FY2029 non-GAAP EPS at 25x. Base, \$750: Meta, OpenAI, and Anthropic deployments ramp through 2027 and 2028, data center revenue roughly doubles in 2027 per guidance, EPYC extends its server lead, and non-GAAP EPS reaches about \$21 in FY2029 on a 35 to 36x multiple; the earnings ramp the forward multiple already assumes, delivered. Bull, \$1,050: AMD becomes the entrenched second source, inference share gains compound, the memory-capacity edge sustains pricing power, and FY2029 non-GAAP EPS of about \$23 to \$24 earns a 44 to 45x growth multiple; the \$20 annual EPS target management articulated, delivered early.



The base case is the consensus earnings ramp actually happening. At \$630.63 the stock is priced for it; the upside from here is execution compounding, and the downside is the multiple meeting reality.

Bottom line

Own it for the execution; size it for the valuation.

AMD earned the trillion-dollar milestone with the strongest strategic position it has ever held: the server CPU war nearly won, a rack-scale AI system in production, the three most important AI labs as anchor customers, and a software stack that finally contests inference. The financials back it up: 50% revenue growth, a 56% non-GAAP gross margin, a net-cash balance sheet, and record free cash flow. The price, however, has already banked the next three years of that story. Positive, with the multiple as the margin of safety's replacement.